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

SERVICES OVERVIEW GUIDE
RELEASE 16.0.R4

3HE 14138 AAAB TQZZA 01

Issue: 01

November 2018

Nokia — Proprietary and confidential.
Use pursuant to applicable agreements.

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

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

© 2018 Nokia.

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

Table of Contents

1 Getting Started

1.1 About This Guide

This guide describes subscriber services, and mirroring support provided by Nokia’s family of routers and presents examples to configure and implement various protocols and services.

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

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

- 7450 ESS
- 7750 SR
- 7950 XRS
- VSR

Table 1 lists the available chassis types for each SR OS router.

Table 1 **Supported SR OS Router Chassis Types**

7450 ESS	7750 SR	7950 XRS
• 7450 ESS-7/12 running in standard mode (not mixed-mode)	• 7450 ESS-7/12 running in mixed-mode (not standard mode) • 7750 SR-a4/a8 • 7750 SR-c4/c12 • 7750 SR-1e/2e/3e • 7750 SR-7/12 • 7750 SR-12e • 7750 SR-7s/14s • 7750 SR-1	• 7950 XRS-16c • 7950 XRS-20/40

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

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

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

1.2 Services Configuration Process

Table 2 lists the tasks associated with configuring subscriber services.

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

Table 2 **Configuration Process**

Area	Task	Section
Service configuration	Configure global service entities	Configuring Global Service Entities with CLI
	Create and configure subscriber (customer) accounts and service distribution points (SDPs)	Common Configuration Tasks
	Service management for customer accounts and SDPs	Global Service Entity Management Tasks

2 Services Overview

2.1 Introduction

A service is a globally unique entity that refers to a type of connectivity service for either Internet or VPN connectivity. Each service is uniquely identified by a service ID and an optional service name within a service area. The Nokia service router 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 Nokia router services can provide Layer 2 bridged service or Layer 3 IP-routed connectivity 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.

Distributed services use service distribution points (SDPs) to direct traffic to another Nokia router through 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).

2.1.1 Service Types

The Nokia routers offer 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.
  - ATM VLL (Apipe) — A 7750 SR point-to-point ATM service between users connected to 7750 SR nodes on an IP/MPLS network.
  - Frame-Relay (Fpipe) — A7750 SR point-to-point Frame Relay service between users connected to 7750 SR nodes on the IP/MPLS network.

  - IP Pipe (Ipipe) — Provides 7750 SR and 7450 ESS IP connectivity between a host attached to a point-to-point access circuit (FR, ATM, PPP) with routed IPv4 encapsulation and a host attached to an Ethernet interface.

  Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information about VLL services.

- Virtual Private LAN Service (VPLS) — A Layer 2 multipoint-to-multipoint VPN. VPLS includes Hierarchical VPLS (H-VPLS) which is an enhancement of VPLS which extends Martini-style signaled or static virtual circuit labeling outside the fully meshed VPLS core.

  Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information about VPLS.

- Internet Enhanced Service (IES) — A direct Internet access service where the customer is assigned an IP interface for Internet connectivity.

  Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for more information about IES.

- Virtual Private Routed Network (VPRN) — A Layer 3 IP multipoint-to-multipoint VPN service as defined in RFC 2547bis.

  Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for more information about VPRN services.

- Circuit Emulation Service (Cpipe) — 7750 SR circuits encapsulated in MPLS use circuit pipes (Cpipes) to connect to the far end circuit. Cpipes support either SAP-spoke SDP or SAP-SAP connections.

2.1.2 Service Policies

Common to all Nokia service router 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 Nokia service router service enhancements. The types of policies that are common to the router’s connectivity services are:

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

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

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

  Filter policies, also referred to as access control lists (ACLs), control the traffic allowed in or out of a SAP based on MAC or IP match criteria. Associating a filter policy on a SAP is optional. Filter policies are identified by a unique filter policy ID. A filter policy must be created before it can be applied to a SAP. A single ingress and single egress filter policy can be associated with a SAP.
- Scheduler policies define the hierarchy and operating parameters for virtual schedulers. Schedulers are divided into groups based on the tier each scheduler is created under. A tier is used to give structure to the schedulers within a policy and define rules for parent scheduler associations.
- Accounting policies define how to count the traffic usage for a service for billing purposes.

  The 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.

2.1.2.1 Multipoint Shared Queuing

Multipoint shared queuing is supported only on Nokia service router routers.

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

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

2.1.2.1.1 Ingress Queuing Modes of Operation

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

2.1.2.1.2 Ingress Service Queuing

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

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

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

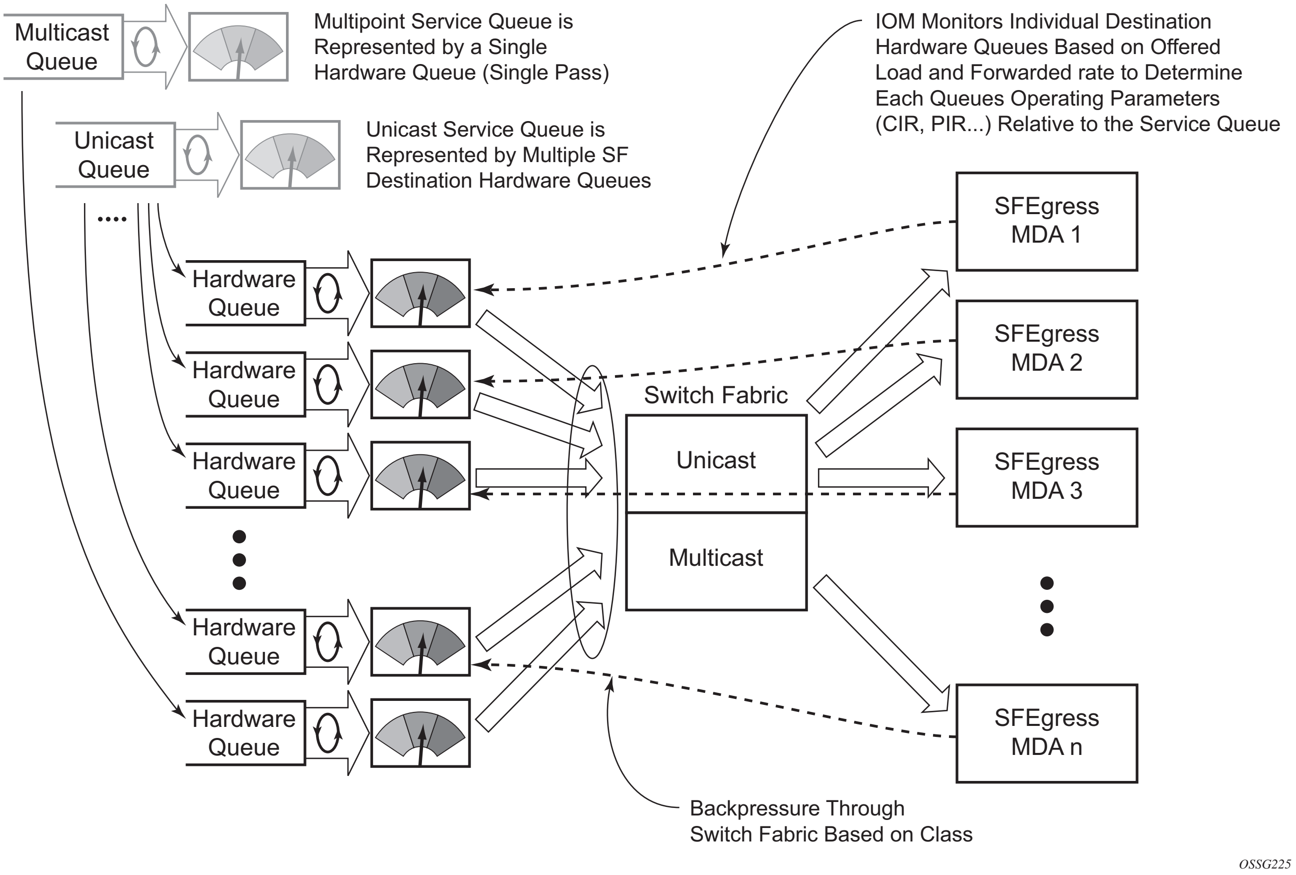
Multicast Queue
Multipoint Service Queue is Represented by a Single Hardware Queue (Single Pass)
IOM Monitors Individual Destination Hardware Queues Based on Offered Load and Forwarded rate to Determine Each Queues Operating Parameters (CIR, PIR...) Relative to the Service Queue
Unicast Queue
Unicast Service Queue is Represented by Multiple SF Destination Hardware Queues
SFEgress MDA 1
Hardware Queue
SFEgress MDA 2
Hardware Queue
Switch Fabric
Unicast
Hardware Queue
SFEgress MDA 3
Multicast
Hardware Queue
Hardware Queue
SFEgress MDA n
Backpressure Through Switch Fabric Based on Class
OSSG225

2.1.2.1.3 Ingress Shared Queuing

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

When the first SAP or SLA profile instance is configured for shared queuing on an ingress forwarding plane, the system allocates eight hardware queues per available destination forwarding plane, one queue per forwarding class. Twenty four hardware queues are also allocated for multipoint shared traffic. The shared queue parameters that define the relative operation of the forwarding class queues are derived from the Shared Queue policy defined in the QoS CLI node. Figure 2 demonstrates shared unicast queuing. SAP or SLA profile instance multipoint queuing is not affected by enabling shared queuing. Multipoint queues are still created as defined in the ingress SAP QoS policy and ingress multipoint packets only traverse the ingress forwarding plane a single time, as demonstrated in Figure 3.

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

Figure 2 **Unicast Service Queuing With Shared Queuing Enabled**

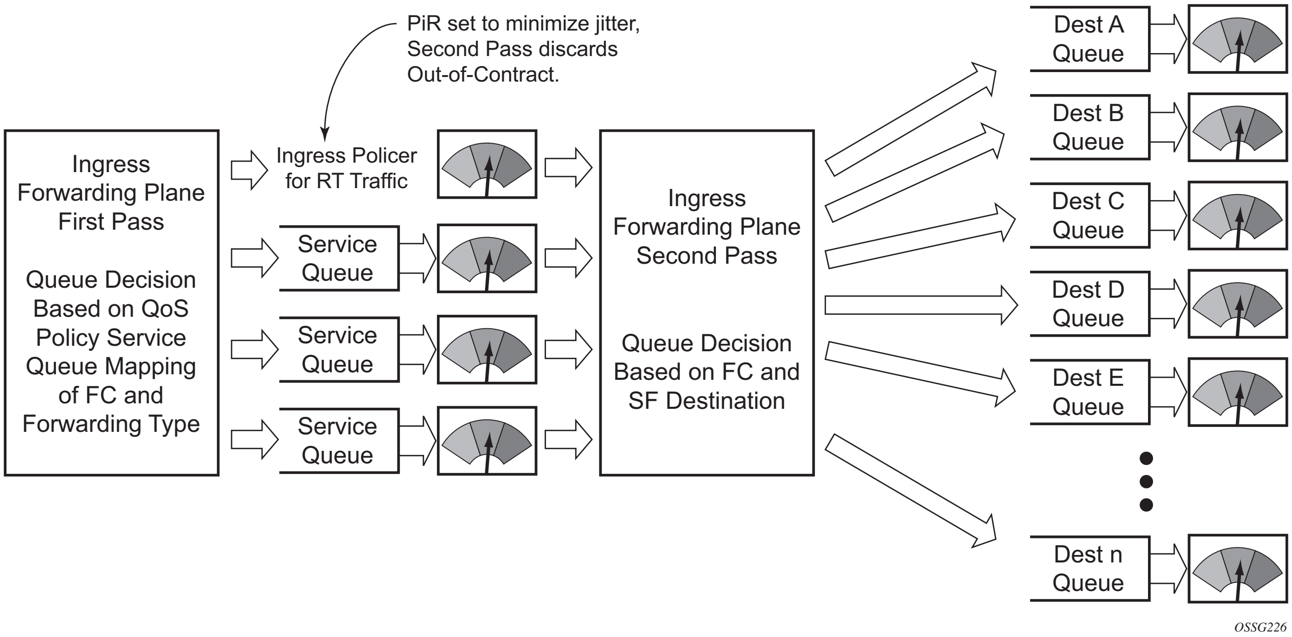
PiR set to minimize jitter,
Second Pass discards
Out-of-Contract.
Ingress Forwarding Plane First Pass
Queue Decision Based on QoS Policy Service Queue Mapping of FC and Forwarding Type
Ingress Policer for RT Traffic
Service Queue
Service Queue
Service Queue
Ingress Forwarding Plane Second Pass
Queue Decision Based on FC and SF Destination
Dest A Queue
Dest B Queue
Dest C Queue
Dest D Queue
Dest E Queue
Dest n Queue
OSSG226

Figure 3 **Multipoint Queue Behavior with Shared Queuing Enabled**

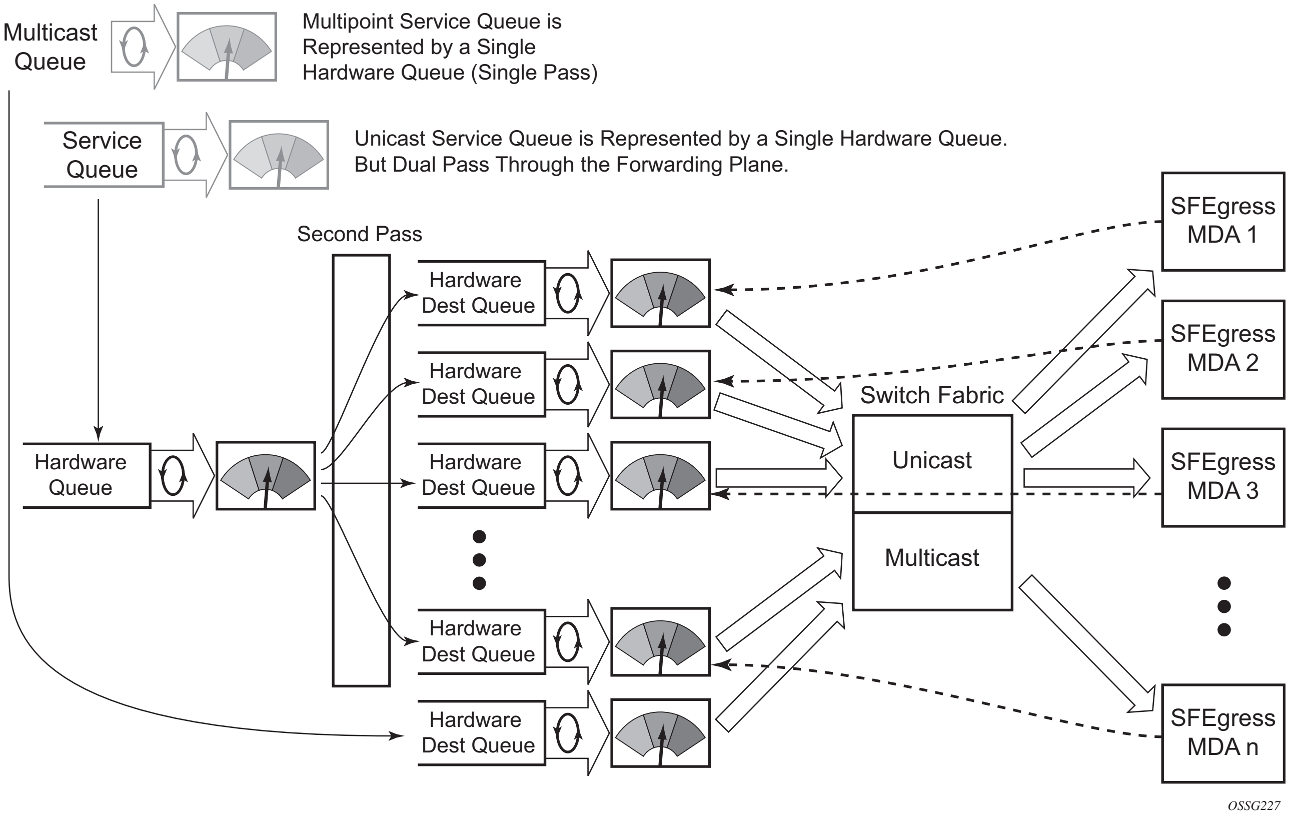

2.1.2.1.4 Ingress Multipoint Shared Queuing

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

The benefit of defining multipoint shared queuing is the savings of the multipoint queues per service. By using the unicast queues in the first pass and then the aggregate shared queues in the second pass, per service multipoint queues are not required. The predominate scenario where multipoint shared queuing may be required is with subscriber managed QoS environments using a subscriber per SAP model. Usually, ingress multipoint traffic is minimal per subscriber and the extra multipoint queues for each subscriber reduces the overall subscriber density on the ingress forwarding plane. Multipoint shared queuing eliminates the multipoint queues sparing hardware queues for better subscriber density. Figure 4 demonstrates multipoint shared queuing.

One disadvantage of enabling multipoint shared queuing is that multipoint packets are no longer managed per service (although the unicast forwarding queues may provide limited benefit in this area). Multipoint packets in a multipoint service (VPLS, IES and VPRN) use significant resources in the system, consuming ingress forwarding plane multicast bandwidth and egress replication bandwidth. Usually, the per service unicast forwarding queues are not rate limited to a degree that allows adequate management of multipoint packets traversing them when multipoint shared queuing is enabled. It is possible to minimize the amount of aggregate multipoint bandwidth by setting restrictions on the multipoint queue parameters in the QoS node’s shared queue policy. Aggregate multipoint traffic can be managed per forwarding class for each of the three forwarding types (broadcast, multicast or unknown unicast – broadcast and unknown unicast are only used by VPLS).

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

Figure 4 Multipoint Shared Queuing Using First Pass Unicast Queues

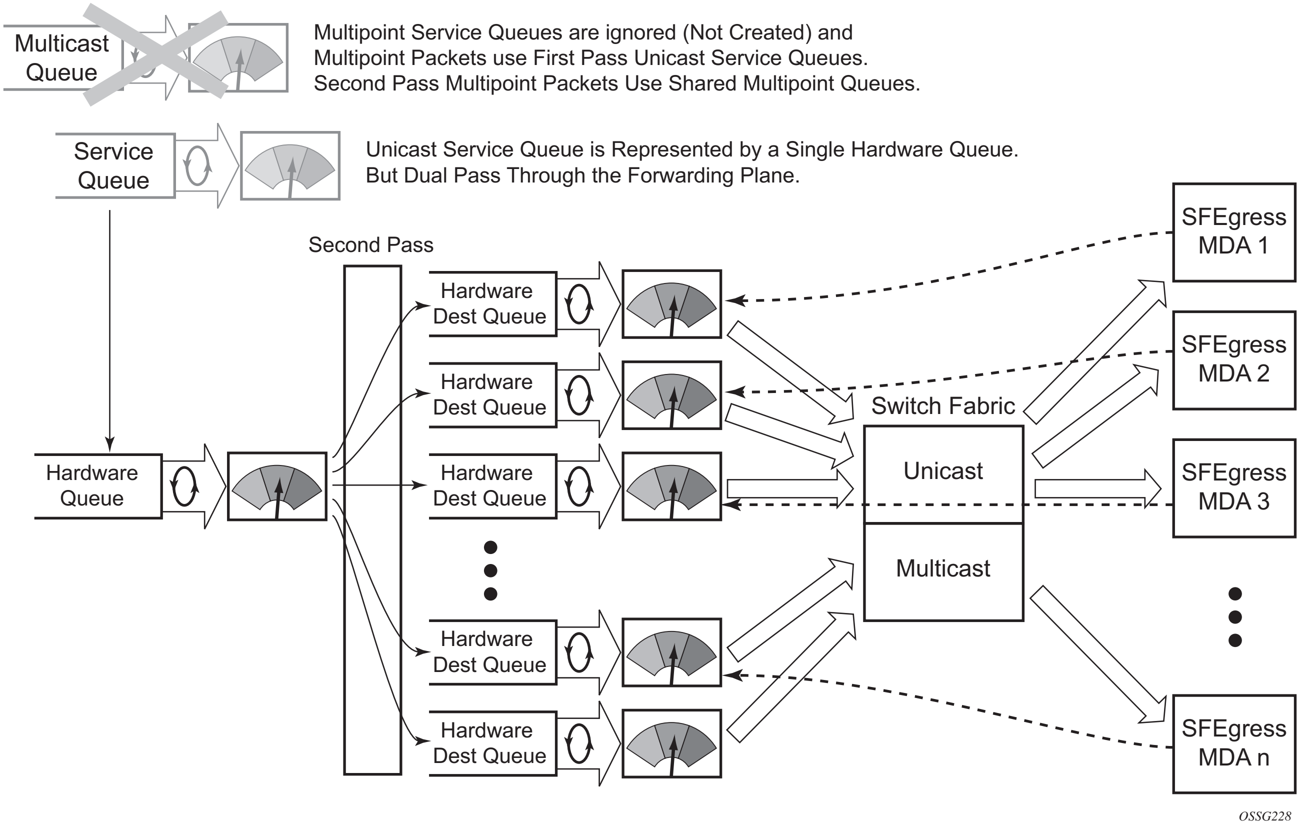
Multicast Queue
Multipoint Service Queues are ignored (Not Created) and
Multipoint Packets use First Pass Unicast Service Queues.
Second Pass Multipoint Packets Use Shared Multipoint Queues.
Service Queue
Unicast Service Queue is Represented by a Single Hardware Queue.
But Dual Pass Through the Forwarding Plane.
Second Pass
Hardware Queue
Hardware Dest Queue
Hardware Dest Queue
Hardware Dest Queue
Hardware Dest Queue
Hardware Dest Queue
Switch Fabric
Unicast
Multicast
SFEgress MDA 1
SFEgress MDA 2
SFEgress MDA 3
SFEgress MDA n
OSSG228

2.2 Nokia Service Model

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

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

- Many services can be bound to a single customer.
- Many services can be bound to a single tunnel.
- Tunnel configurations are independent of the services they carry.
- Changes are made to a single logical entity rather than multiple ports on multiple devices. It is easier to change one tunnel rather than several services.
- The operational integrity of a logical entity (such as a service tunnel and service end points) can be verified rather than dozens of individual services improving management scaling and performance.
- On 7450 ESS and 7750 SR OS, a failure in the network core can be correlated to specific subscribers and services.
- QoS policies, filter policies, and accounting policies are applied to each service instead of correlating parameters and statistics from ports to customers to services.

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

2.3 Service Entities

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

- Customers
- Service Access Points (SAPs)
- Service Distribution Points (for distributed services only)

Figure 5 **Service Entities**

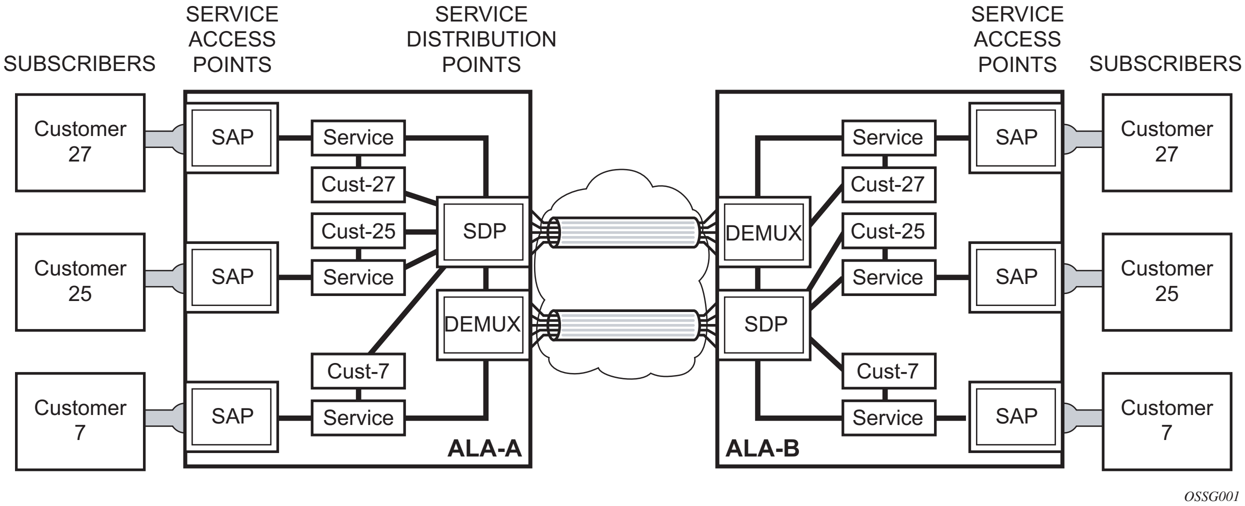

2.3.1 Customers

In this section, 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.

2.3.2 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 a router (for example Figure 6). The SAP configuration requires that slot, XMA or MDA, and port/channel information be specified. The slot, XMA or MDA, and port/channel parameters must be configured prior to provisioning a service (Refer to the XMAs, Cards, MDAs, and Ports sections of the SR OS Interface Configuration Guide).

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

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

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

Figure 6 **7750 SR/7950 XRS Service Access Point (SAP)**

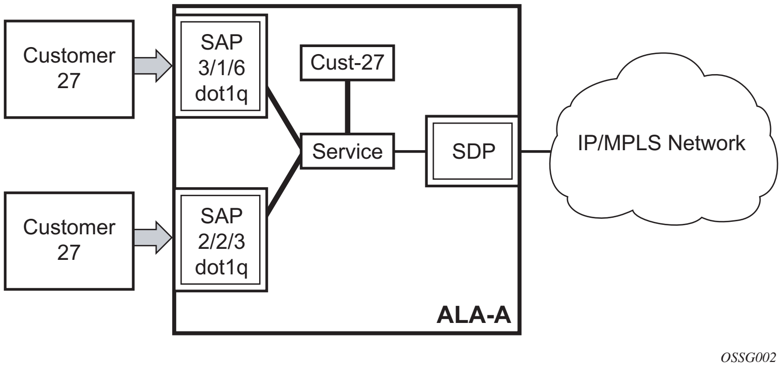

A SAP can also be associated with a pseudowire port rather than an access port. Such SAPs are called pseudowire SAPs. This is only applicable to IES or VPRN services. Pseudowire ports represent pseudowires in enhanced subscriber management (ESM). For a description of pseudowire ports, see the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide*.

2.3.2.1 SAP Encapsulation Types and Identifiers

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

2.3.2.2 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 6 and Figure 7). For example, the port is connected to a multi-tenant unit (MTU) device with multiple downstream customers. The encapsulation ID used to distinguish an individual service is the VLAN ID in the IEEE 802.1Q header.
- QinQ — The QinQ encapsulation type adds a IEEE 802.1Q tag to the 802.1Q tagged packets entering the network to expand the VLAN space by tagging tagged packets, producing a double tagged frame.

There are several 7750 SR encapsulation service options on SONET/SDH channels:

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

- ATM — ATM, ATM-FR, ATM SAP-bridge encapsulation termination Epipe and VPLS.
- Frame Relay — Supports the switched data link layer protocol that handles multiple virtual circuits.

There are several 7450 ESS encapsulation service options on SONET/SDH channels:

- Internet Protocol Control Protocol (IPCP) — Supports a single IP service on a SONET/SDH port. This is typically used for router interconnection using point-to-point protocol (PPP).
- Bridging Control Protocol (BCP-null) — Supports a single service on the SONET/SDH port. This is used for bridging a single service between two devices using PPP over SONET/SDH. The encapsulation ID is always 0 (zero).
- Bridging Control Protocol (BCP-dot1q) — Supports multiple services on the SONET/SDH port. This encapsulation type is used for bridging multiple services between two devices using PPP over SONET/SDH. The encapsulation ID used to distinguish services is the VLAN ID in the IEEE 802.1Q header in the BCP-encapsulated frame.
- Frame Relay — Supports the switched data link layer protocol that handles multiple virtual circuits.

Figure 7 **7750 SR/7950 XRS and 7450 ESS Multiple SAPs on a Single Port/ Channel**

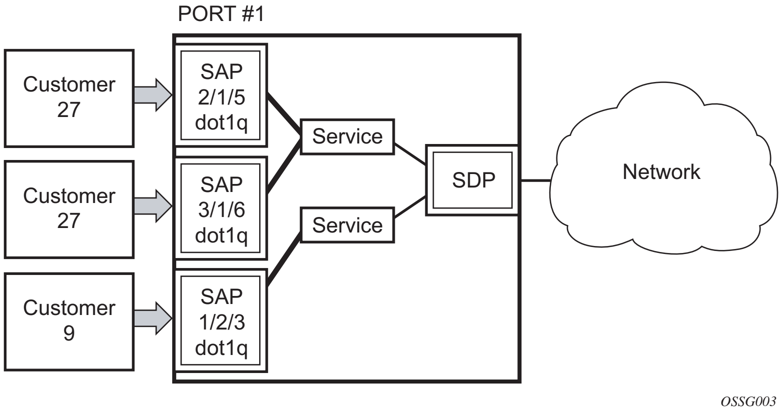

2.3.2.3 Default SAP on a Dot1q Port

This feature introduces default SAP functionality on Dot1q-encapsulated ports. This is similar to the functionality provided by Q1* SAP on QinQ encapsulated ports, meaning that on On dot1q-encapsulated ports where a default SAP is configured, all packets with q-tags not matching any explicitly defined SAPs will be assigned to this SAP. SAPs with default QinQ encapsulation are supported in VPLS, Epipe, IES and VPRN services. Both DHCP snooping and IGMP snooping are supported for QinQ SAPs. In this context, the character "*" indicates default which means allow through. A 0 value means that it should not be there which allows the Qtag to be missing.

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

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

There a few constraints related for the use of default SAP on a Dot1q-encapsulated port:

- This type of SAP is supported only on VPLS and Epipe services and cannot be created in IES and VPRN services as it cannot preserve VLAN tag markings.
- For VPLS SAPs with STP enabled, STP listens to untagged and null-tagged BPDUs only. All other tagged BPDUs are forwarded like other customer packets. This is the same behavior as null-encapsulated ports.
- IGMP snooping is not supported on a default SAP. This would require remembering VLAN tags per hosts. By not allowing IGMP snooping of this SAP, all IGMP packets will be transparently forwarded.
- This type of SAP is mutually exclusive with a SAP defined by explicit null encapsulation (for example, 1/1/1:0). This avoids conflict as to which SAP untagged frames should be associated.

2.3.2.4 QinQ SAPs

A QinQ SAP has the following format:

qinq <port-id | lag-id>:qtag1.qtag2

Where:

- *qtag1* is the outer qtag value - [*, 0 to 4094]
- *qtag2* is the inner qtag value - [*, null, 0 to 4094]

Regular QinQ SAPs have qtag1 and qtag2 values between 1 and 4094. In addition, QinQ Ethernet and LAG ports support the following “default” SAPs that can be enabled by the **new-qinq-untagged-sap** command:

- ‘*.null’ is defined as a default sap for single-tagged frames in a QinQ port. This SAP accepts single tags in the range 0 to 4095 as well as untagged traffic.
- ‘*.*’ is defined as a default sap for double-tagged frames in a QinQ port. This SAP accepts untagged, singly tagged, and doubly tagged frames with tags in the range 0..4095.
- In addition to the above-mentioned SAPs, qtag2 can also be '0' or '*' when qtag1 is an explicit value in the 1 to 4094 range, for instance: 1/1/1:10.0 or 1/1/1:10.*. Assuming qtag1 is the same value, qtag1.* and qtag1.0 are supported in the same QinQ port

A SAP lookup is performed when a new frame arrives to a QinQ port. This 'lookup' is based on the <outer-tag, inner-tag> values of the frame.

Table 3 shows the SAP lookup precedence order for incoming frames with <*qtag1.qtag2*> qtag values.

Table 3 **SAP Lookup Precedence Order for Incoming Frames**

Incoming Frame *qtag1.qtag2*	**System/Port settings [new-qinq-untagged-sap=YES]**					
	SAP Lookup Precedence Order					
	:X.Y	**:X.0**	**:X.***	**:0.***	**:*.null**	**:*.***
x.y	1st		2nd			3rd
x.0		1st	2nd			3rd
0.y				1st		2nd
0.0				1st		2nd
x		1st	2nd		3rd	4th
0				1st	2nd	3rd
<untagged>				1st	2nd	3rd

The following considerations apply to the information described in Table 3:

- All six SAP types (:X.Y, :X.0, :X.*, :0.*, :*.null and :*.*) are supported in the same QinQ port and, in the table, they are ordered from the most specific (left-hand side) to the least specific with the following VID matching ranges:
  - X or Y means <1 to 4094>
  - * means <0 to 4095> or untagged
  - null means 'no tag'
- The user can decide the SAP types that are configured in a specific port. Not all SAP types must be configured in a port.
- Table 3 shows the lookup behavior for ingress frames and priority across SAPs in case more than one can match a given ingress frame. The SAP lookup result for a given frame does not depend on the operational status of the SAP. For instance:
  - In a port with SAPs 1/1/1:0.* and 1/1/1:*.* defined, the SAP lookup for a given frame with VIDs (0, 300) will yield SAP 1/1/1:0.* regardless of its operational status.
  - The frame will only match SAP 1/1/1:*.* when the 0.* SAP is removed from the configuration.
- The following apply to VLAN tag handling:
  - The system will not strip-off any tags for frames entering the default SAPs (:0.*, :*.null or :*.*).
  - No extra tags are added when the system transmits frames on the default SAPs (:0.*, :*.null or :*.*).

The following examples illustrate the SAP classification QinQ ports. The examples assume that the **new-qinq-untagged-sap** command is enabled.

Example - 1

As shown in Figure 8, assuming that the **new-qinq-untagged-sap** command is enabled, the following SAPs are defined on the same port:

- 1/1/1:3000.1001 - business customer - vpls-1
- 1/1/1:2000.1002 - business customer - vpls-2
- 1/1/1:20.0 - BNG traffic - vpls-3
- 1/1/1:20.* - business customer - epipe-4
- 1/1/1:0.* - business customer - epipe-5
- 1/1/1:*.null - wholesaling single tag - epipe-6
- 1/1/1:*.* - wholesaling double tag - epipe-7

Figure 8 **Example 1 SAP Classification QinQ Ports**

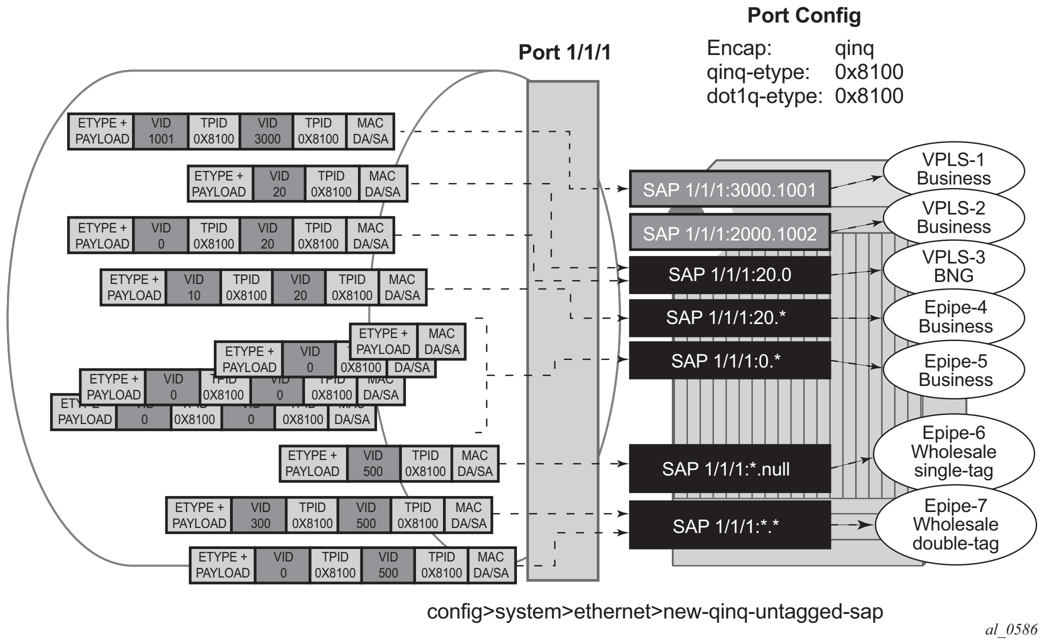

Based on the SAPs configuration described above, the incoming traffic is classified in the following way - notation (outer-VID, inner-VID):

- (3000, 1001) goes to vpls-1
- (20) goes to BNG (vpls-3)
- (20, 0) goes to BNG (vpls-3)
- (20, 10) goes to epipe-4
- untagged, (0), (0, 0), and (0, 10) go to epipe-5
- (500) goes to wholesaling single tag (epipe-6)
- (500, 300) and (500, 0) go to wholesaling double tag (epipe-7)

Example - 2

Figure 9 highlights how untagged, VID=0 tagged frames and 20.X frames are classified in the absence of the 0.* and 20.* SAPs.

Figure 9 Example 2 SAP Classification QinQ Ports

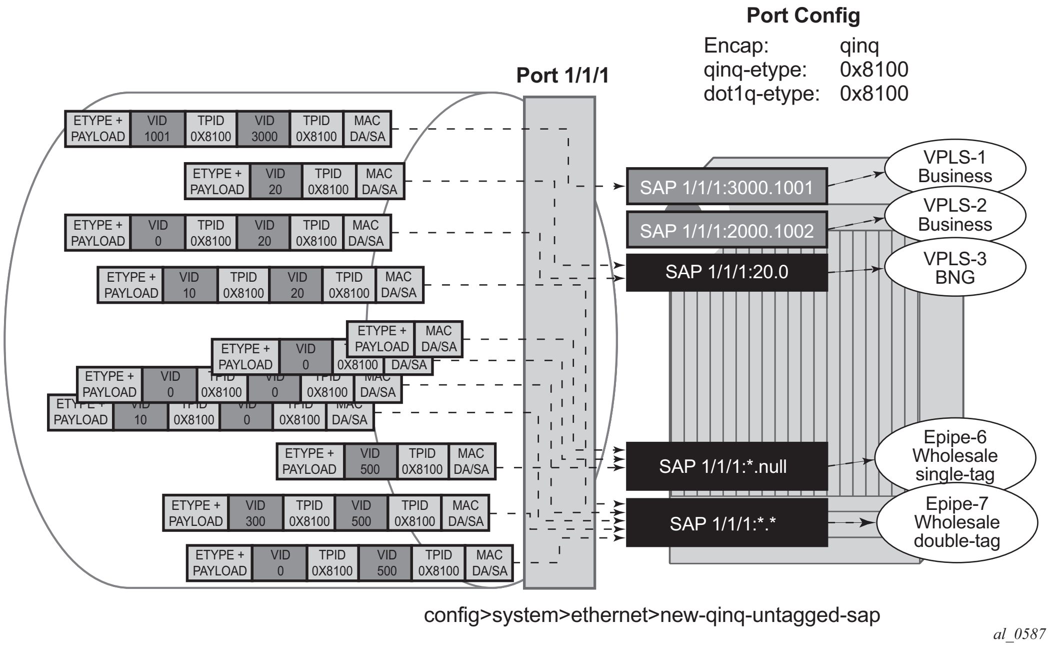
Port Config
Encap: qinq
qinq-etype: 0x8100
dot1q-etype: 0x8100
Port 1/1/1
ETYPE + PAYLOAD | VID 1001 | TPID 0X8100 | VID 3000 | TPID 0X8100 | MAC DA/SA
ETYPE + PAYLOAD | VID 20 | TPID 0X8100 | MAC DA/SA
ETYPE + PAYLOAD | VID 0 | TPID 0X8100 | VID 20 | TPID 0X8100 | MAC DA/SA
ETYPE + PAYLOAD | VID 10 | TPID 0X8100 | VID 20 | TPID 0X8100 | MAC DA/SA
ETYPE + PAYLOAD | MAC DA/SA
ETYPE + PAYLOAD | VID 0 | 0X8100 | DA/SA
ETYPE + PAYLOAD | VID 0 | TPID 0X8100 | VID 0 | TPID 0X8100 | MAC DA/SA
ETYPE + PAYLOAD | VID 10 | TPID 0X8100 | VID 0 | TPID 0X8100 | MAC DA/SA
ETYPE + PAYLOAD | VID 500 | TPID 0X8100 | MAC DA/SA
ETYPE + PAYLOAD | VID 300 | TPID 0X8100 | VID 500 | TPID 0X8100 | MAC DA/SA
ETYPE + PAYLOAD | VID 0 | TPID 0X8100 | VID 500 | TPID 0X8100 | MAC DA/SA
SAP 1/1/1:3000.1001
SAP 1/1/1:2000.1002
SAP 1/1/1:20.0
SAP 1/1/1:*.null
SAP 1/1/1:*.*
VPLS-1 Business
VPLS-2 Business
VPLS-3 BNG
Epipe-6 Wholesale single-tag
Epipe-7 Wholesale double-tag
config>system>ethernet>new-qinq-untagged-sap
al_0587

As outlined in Figure 9, assuming the **new-qinq-untagged-sap** command is enabled, the following SAPs are defined on the same port:

- 1/1/1:3000.1001 - business customer - vpls-1
- 1/1/1:2000.1002 - business customer - vpls-2
- 1/1/1:20.0 - BNG traffic - vpls-3
- 1/1/1:*.null - wholesaling single tag - epipe-6
- 1/1/1:*.* - wholesaling double tag - epipe-7

Incoming traffic - notation (outer-VID, inner-VID)

- (3000, 1001) goes to vpls-1
- (20) goes to BNG (vpls-3)
- (20, 0) goes to BNG (vpls-3)
- (20, 10) goes to wholesaling double tag (epipe-7)
- untagged and (0) go to wholesaling single tag (epipe-6)
- (500) goes to wholesaling single tag (epipe-6)
- (500, 300) and (500, 0) go to wholesaling double tag (epipe-7)
- (0,0), and (0,10) goes to wholesaling double tag (epipe-7)

Note: The system will not add service-delimiting tags with VID=0; however, tags with VID=0 are accepted and classified appropriately.

The following constraints must be considered when configuring default QinQ SAPs (:0.*, :*.null, :*.*):

- Only supported in Ethernet ports or LAG.
- Only supported on Epipe, PBB-Epipe, VPLS and I-VPLS services. They are not supported on VPRN, IES, R-VPLS or B-VPLS services.
- Capture SAPs with encapsulation :*.* cannot coexist with a default :*.* SAP on the same port.
- Inverse-capture SAPs (*.x) are mutually-exclusive with :*.null SAPs.
- *.null SAPs are not supported for Open Flow matching and forwarding.
- The following applies to Eth-CFM:
  - Primary VLAN is not supported.
  - Eth-CFM extractions occur within the service after the packet lookup has determined which service the inbound packet belongs to.
  - All three SAPs (*.null, *.* and 0.*) are treated equally by ETH-CFM. Only untagged CFM PDUs are extracted by a local MEP or MIP. Additional tags in the header may match the service context but are not extracted by ETH-CFM for processing.
  - ETH-CFM PDU transmission encapsulation is based on the SAP configuration. This means that the ETH-CFM PDUs will be transmitted out all three of these SAPs untagged. Care must be taken to ensure that there is no downstream service that may intercept the ETH-CFM PDUs that are not intended for that service. See Table 3 for a description of the SAP lookup precedence order for incoming frames and to understand the potential consequences.
- Default QinQ SAPs do not support the following features:
  - PW-SAPs
  - Eth-tunnel or eth-ring SAPs
  - VLAN-translation *copy-outer*
  - E-Tree root-leaf-tag SAPs
  - Subscriber-management features
  - BPDU-translation
  - Eth-tunnels
  - IGMP-snooping
  - MLD-snooping

2.3.2.5 Services and SAP Encapsulations

The services and SAP encapsulations are listed in Table 4.

Table 4 **Service and SAP Encapsulations**

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

2.3.2.6 SAP Configuration Considerations

When configuring a SAP, consider the following:

- A SAP is a local entity and only locally unique to a given device. The same SAP ID value can be used on another Nokia router.
- There are no default SAPs. All SAPs in subscriber services must be created.
- The default administrative state for a SAP at creation time is administratively enabled.
- When a SAP is deleted, all configuration parameters for the SAP will also be deleted. For Internet Enhanced Service (IES), the IP interface must be shutdown before the SAP on that interface may be removed.
- A SAP is owned by and associated with the service in which it is created in each router.
- A port or channel with a dot1q or BCP-dot1q encapsulation type means the traffic for the SAP is identified based on a specific IEEE 802.1Q VLAN ID value. The VLAN ID is stripped off at SAP ingress and the appropriate VLAN ID placed on at SAP egress. As a result, VLAN IDs only have local significance, so the VLAN IDs for the SAPs for a service need not be the same at each SAP.

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

2.3.2.7 G.8032 Protected Ethernet Rings

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

For further information on Ethernet rings, see G.8032 Ethernet Ring Protection Switching.

2.3.2.8 SAP Bandwidth CAC

This feature provides a bandwidth CAC function per port or per LAG based on an admin bandwidth configured on a SAP and on the associated port or LAG. A booking factor is provided in order to allow over/under booking of the sum of the SAP bandwidth compared to the port/LAG bandwidth.

The admin bandwidth is an abstract value which could represent either, or both, of the ingress or egress bandwidth and is statically configured.

The goal of the CAC function is to ensure that the sum of the admin SAP bandwidth on a port or LAG does not exceed the admin bandwidth configured on that port or LAG.

This is supported on all service Ethernet SAPs, excluding PW SAPs, Ethernet tunnels and subscriber group interface SAPs. It is not supported in a VPLS or Epipe SAP template. It is applicable to both access and hybrid ports or LAGs; in the case of a hybrid port or LAG, the SAP CAC bandwidth only applies to the access operation.

By default a SAP, port or LAG has no admin bandwidth configured in which case it is excluded from the CAC function. Configuring an admin bandwidth on a SAP will cause the CAC function to be enforced.

An admin bandwidth can only be configured on a SAP connected to a port or LAG which itself has an admin bandwidth configured. When a LAG is configured, the admin bandwidth and booking factor on its constituent ports are ignored.

The system tracks the requested and available bandwidth per port or LAG, where the available bandwidth is equal to the admin bandwidth on the port or LAG, with the booking factor applied, minus the sum of admin bandwidth configured on its SAPs. An attempt to increase a SAP's admin bandwidth will fail if there is insufficient available bandwidth on its port or LAG.

The admin bandwidth and booking factor for the port or LAG is configured as follows:

```
configure
    lag <lag-id>
        access
            bandwidth <bw-value>
            booking-factor <percentage>
    port <port-id>
        ethernet
            access
                bandwidth <bw-value>
                booking-factor <percentage>
    service
        [apipe | cpipe | epipe | fpipe | ipipe | vpls] <service-id>
            sap <sap-id>
                bandwidth <bw-value>
        [ies | vprn] <service-id>
            interface <ip-int-name>
                sap <sap-id>
                    bandwidth <bw-value>
```

Changes in admin bandwidth and booking factor are possible dynamically without having to disable the SAP, port or LAG.

Once a SAP has been allocated bandwidth on a port or LAG that bandwidth is allocated to that SAP regardless of whether the SAP and/or port or LAG are up or down (either administratively or operationally). The admin bandwidth must be removed from the SAP configuration in order to free up its bandwidth on the port or LAG. Actions such as clearing the card or MDA, power-cycling the card or removing/inserting a card or MDA do not change the SAP and port or LAG CAC state.

2.3.2.8.1 CAC Enforcement

The CAC is enforced when an admin bandwidth is configured on a SAP (this may be when initially configuring the admin bandwidth or when modifying an existing admin bandwidth value).

The CAC enforcement is achieved by comparing the newly requested SAP admin bandwidth (the incremental admin bandwidth being configured above any currently assigned admin bandwidth) with the available admin bandwidth on its port or LAG.

The operation is as follows:

- If a SAP's admin bandwidth is increased and the incremental requested admin bandwidth is
  - larger than the port or LAG available bandwidth then the command to increase the SAP admin bandwidth fails.
  - smaller or equal to the available port or LAG bandwidth then the incremental bandwidth is subtracted from the available port or LAG bandwidth.
- If a SAP's admin bandwidth is reduced then the available port or LAG bandwidth is increased accordingly.
- If the port or LAG admin bandwidth is increased, the available port or LAG bandwidth is increased accordingly.
- If the port or LAG admin bandwidth is decreased, the available port or LAG bandwidth is decreased accordingly. However, if the resulting available bandwidth would be less than the sum of the currently allocated SAP admin bandwidth on that port or LAG, then the command to decrease the admin bandwidth fails.
- If the port or LAG booking factor is decreased, the available port or LAG bandwidth is decreased accordingly. However, if the resulting available bandwidth would be less than the sum of the currently allocated SAP admin bandwidth on that port or LAG, then the command to decrease the booking factor fails.
- If the SAP admin bandwidth is removed, it is excluded from the SAP bandwidth CAC function. Its admin bandwidth is added to the related port or LAG available bandwidth.
- The port or LAG admin bandwidth can only be removed if all of its SAPs are excluded from the CAC function.

An example is given below. A port is configured with an admin bandwidth of 500 Mb/s, and a SAP on that port with a bandwidth of 10 Mb/s. The show output gives these configured values together with the port's available and booked admin bandwidth. An increase of the SAP admin bandwidth to 600 Mb/s is attempted, which fails as there is insufficient available admin bandwidth on the port.

The port's booking factor is increased to 200% and the increase of the SAP admin bandwidth to 600 Mb/s is then successful as the port's available admin bandwidth becomes 1 Gb/s. The port's booked admin bandwidth is 600 Mb/s and so its available admin bandwidth becomes 400 Mb/s.

```
*A:PE# configure port 1/1/1 ethernet access bandwidth 500000
*A:PE# configure service vpls 1 sap 1/1/1:1 bandwidth 10000
*A:PE# show service id 1 sap 1/1/1:1 detail | match Bandwidth
Bandwidth          : 10000
*A:PE# show port 1/1/1 detail | match expression "Bandwidth | BW"
Access Bandwidth   : 500000                    Booking Factor   : 100
Access Available BW: 490000
Access Booked BW   : 10000
*A:PE# configure service vpls 1 sap 1/1/1:1 bandwidth 600000
MINOR: SVCMGR #2664 Insufficient bandwidth available
*A:PE# show service id 1 sap 1/1/1:1 detail | match Bandwidth
Bandwidth          : 10000
*A:PE# *A:PE# configure port 1/1/1 ethernet access booking-factor 200
*A:PE# configure service vpls 1 sap 1/1/1:1 bandwidth 600000
*A:PE# show service id 1 sap 1/1/1:1 detail | match Bandwidth
Bandwidth          : 600000
*A:PE# show port 1/1/1 detail | match expression "Bandwidth | BW"
Access Bandwidth   : 500000                    Booking Factor   : 200
Access Available BW: 400000
Access Booked BW   : 600000
*A:PE#
```

2.3.3 Connection Profile VLAN SAPs

The **connection-profile-vlan** SAPs (CP SAPs) allow the association of a range of customer VIDs to a given SAP. CP SAPs can be used to build Layer 2 Services that are fully compatible with MEF 10.3 Bundling Service Attributes and RFC 7432 EVPN VLAN Bundle Service interfaces.

The **config>connection-profile-vlan>vlan-range** command defines the range of customer VIDs to be matched when the **connection-profile-vlan** is associated with a dot1q or QinQ SAP. The following CLI output example shows the use of **connection-profile-vlan** in dot1q and QinQ SAPs:

```
A:PE# configure connection-profile-vlan 1 create
        vlan-range 5 to 100
        vlan-range 150 to 300
        vlan-range 350
    exit
A:PE>config>service>vpls# info
-------------------------------------------
<snip>
sap 1/1/1:cp-1 create
    no shutdown
exit
sap 1/1/2:100.cp-1 create
```

```
    no shutdown
exit
sap 1/1/3:cp-1.* create
    no shutdown
exit
<snip>
```

As far as VLAN manipulation is concerned, the CP SAP behavior is equivalent to the default SAP's (when the ingress VID falls into the range configured in the CP), where the range of VIDs included is not service-delimiting and therefore, the VIDs are not pushed/popped. The main differences between the CP SAPs and the default SAPs are:

- Resources — A default SAP consumes one SAP instance, whereas a CP SAP consumes a number of SAP instances equal to the number of VLANs in the range. The amount of SAP instances consumed in the system can be checked by executing the following commands:

```
*A:Dut# tools dump resource-usage system
===============================================================================
Resource Usage Information for System
===============================================================================
                                                  Total   Allocated        Free
-------------------------------------------------------------------------------
<snip>
                                SAP Entries  |   262143           8      262135
===============================================================================

*A:Dut# tools dump resource-usage card 1 fp 1
===============================================================================
Resource Usage Information for Card Slot #1 FP #1
===============================================================================
                                                  Total   Allocated        Free
-------------------------------------------------------------------------------
<snip>
                              SAP Instances  |    63999         254       63745
<snip>
===============================================================================
```

  Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Basic System Configuration Guide* for a complete description of the **tools dump resource-usage** command.

- Unlike the default SAP, a CP SAP cannot co-exist with a vlan SAP that is in the same range. For example: 1/1/1:* and 1/1/1:100 can co-exist; in contrast, 2/1/1:cp-1 (cp-1 = vlan 1 to 200) and 2/1/1:100 cannot co-exist.

Figure 10 shows customer VID processing by SAPs with service-delimiting VIDs, and by CP SAPs. SAP 1/1/1:cp-1 does not strip off or push VID 10, whereas SAP 1/1/1:100 and SAP 1/1/1:200 do strip off and push the corresponding VID.

Figure 10 **VLAN Tag Handling**

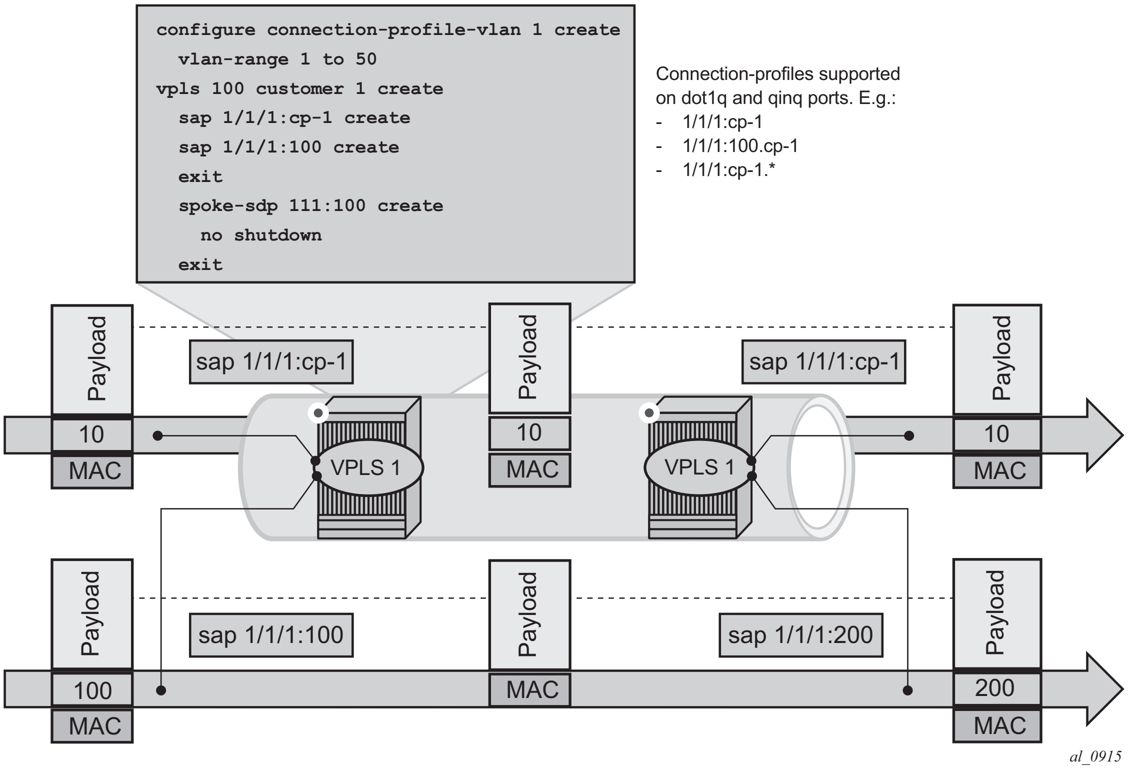

```
configure connection-profile-vlan 1 create
  vlan-range 1 to 50
vpls 100 customer 1 create
  sap 1/1/1:cp-1 create
  sap 1/1/1:100 create
  exit
  spoke-sdp 111:100 create
    no shutdown
  exit
```

Connection-profiles supported on dot1q and qinq ports. E.g.:
- 1/1/1:cp-1
- 1/1/1:100.cp-1
- 1/1/1:cp-1.*

A **connection-profile-vlan** allows the configuration of VLAN ranges with the following characteristics.

- A **vlan-range** can be defined as a single VID (for example, **vlan-range 101**), or two VIDs delimiting the beginning and the end of the range (for example, **vlan-range 105 to 107**).
- Discontinuous ranges are allowed.
- Overlapping ranges are not allowed within the same **connection-profile-vlan**. VLAN range overlapping can exist across different connection-profiles as long as they are not applied to the same port (in the case of dot1q ports), or the same port and service-delimiting tag (in the case of qinq ports). For example:
  - 1/1/1:x.cp-1 and 1/1/1:y.cp-2 can coexist on the same port, where cp-1 includes vids [10-20] and cp-2 includes vids [15-25]
  - If x=y, then the overlapping is not possible in the above case.
- A **connection-profile-vlan** must have at least one range (with a single or multiple VIDs) before it can be associated with a SAP.
- A **connection-profile-vlan** cannot contain an explicitly defined SAP within any of the ranges when the explicit SAP is configured on the same port.
- The configured VLAN ranges cannot contain VIDs 0 or 4095.

- The **connection-profile-vlan** SAPs are supported in Layer 2 services only. No IES or VPRN services can contain CP SAPs.
- CP SAPs are supported on access or hybrid ports but are not on network interfaces.
- CP SAPs are supported in (non-PBB) Epipe and VPLS services.
- CP SAPs support SAP based QoS policies. VID type MAC criteria can be used on CP SAPs to apply specific QoS on a given VLAN within the connection-profile-vlan.
- The legacy OAM commands (**mac-ping**, **mac-trace**, **mac-purge**, and **mac-populate**) do not work with CP SAPs.

2.3.3.1 Using connection-profile-vlan in Dot1q Ports

Table 5 describes the SAP lookup matching order that is applied when **connection-profile-vlan** is used in dot1q ports.

Table 5 **SAP Lookup Matching Order for Dot1q Ports**

Incoming Frame qtag VID value	**SAP lookup precedence order (:0 and :* are mutually-exclusive on the same port)**			
	:X	**:CP**	**:0**	**:***
x (belongs to the CP range)	1st	1st		2nd
0			1st	1st
<untagged>			1st	1st

2.3.3.2 Using connection-profile-vlan in QinQ Ports

Table 6 describes the SAP lookup matching order that is applied when **connection-profile-vlan** is used in QinQ ports.

Table 6 **SAP Lookup Matching Order for QinQ Ports**

Incoming Frame	**System/port settings = new-qinq-untagged-sap**							
qtag1.qtag2	**SAP lookup precedence order** **(assumption: X and Y are defined in CP ranges)**							
	:X.Y	**:X.0**	**:X.CP**	**:CP.***	**:X.***	**:0.***	**:*.null**	**:*.***
x.y	1st		1st	2nd	2nd			3rd
x.0		1st		2nd	2nd			3rd
0.y						1st		2nd
0.0						1st		2nd
x		1st		2nd	2nd		3rd	4th
0						1st	2nd	3rd
<untagged>						1st	2nd	3rd

Consider the following when using connection-profile-vlan (CP) in qinq ports:

- A CP can be defined for inner or outer tags but not both at the same time; for example, :X.CP and :CP.* are possible, but not ':CP.CP'.
- It is important to note that :CP:Y is not allowed; for example, if a CP is defined at the outer VID, the inner VID can only be a '*' or a '0'.
- :0.CP SAPs are not allowed; if the outer VID is 0, the inner VID cannot be a connection-profile-vlan value.
- A CP cannot contain a VID that is associated to an explicitly defined inner or outer tag in a specific port. For example, assuming that X and Y are tags defined in 'CP', a given port can be defined with ":X.CP" or ":Y.CP" but not with ":X.CP" or ":Y.CP".
- The following combinations are allowed:
  - :CP.0 - matches frames with outer tags contained in CP and inner tags 0 or null
  - :CP.* - matches frames with outer tags contained in CP and any inner tags
- In the case where a VLAN tag combination matches different SAPs, the highest priority SAP will be picked, irrespective of its oper-status, as long as the SAP is still created. Therefore, if the SAP is down, the frames will not go to a different SAP. For example, suppose that ingress frames with VIDs 10.25 are classified as part of sap 10.cp-1. Only when sap 10.cp-1 is removed from the configuration will the frames with VIDs 10.25 go to sap cp-1.*.

```
# cp-1 includes vlan ids (10-100).
sap 1/1/4:cp-1.* create
```

```
exit
sap 1/1/4:10.cp-1 create
exit
```

2.3.4 Service Distribution Points

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 Nokia 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 (either GRE, MPLS, or L2tPv3).
- 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 router 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.

2.3.4.1 SDP Binding

To configure a distributed service from ALA-A to ALA-B, the SDP ID (1) (shown in Figure 11) 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 11 **GRE Service Distribution Point (SDP) Pointing From ALA-A to ALA-B**

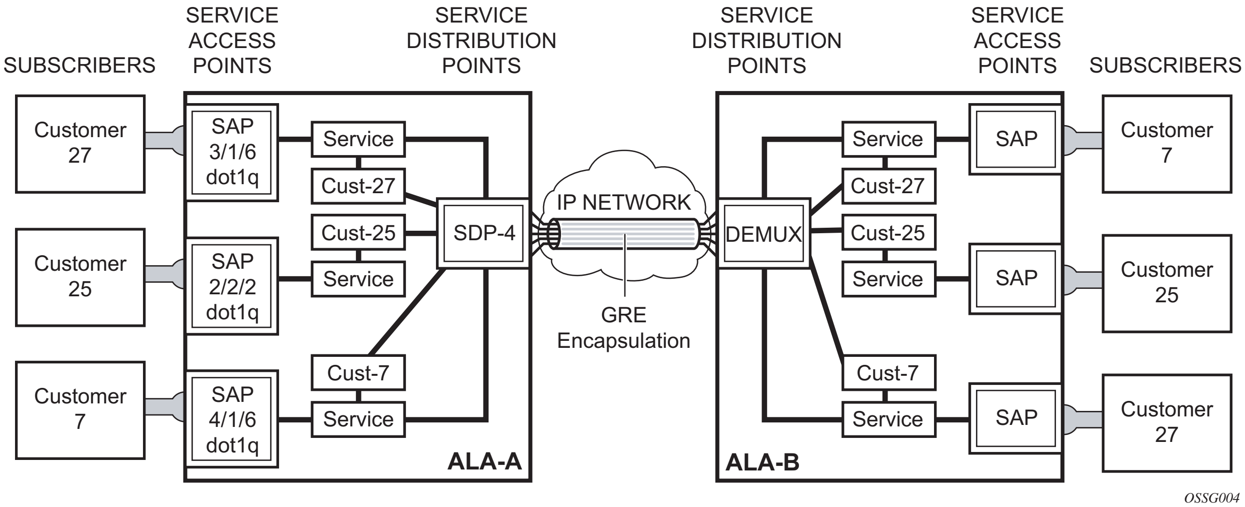

2.3.4.2 Spoke and Mesh SDPs

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

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

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

2.3.4.3 SDP Using BGP Route Tunnel

SDP is enhanced to use BGP route tunnel to extend inter-AS support for routes and services. An SDP can be configured based on service transport method (for example, GRE or MPLS tunnel). 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).

For the 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 and 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 or VPN at far-end PE nodes.

2.3.4.4 SDP Keepalives

SDP keepalives actively monitor the SDP operational state using periodic Nokia SDP ping echo request and echo reply messages. Nokia SDP ping is a part of Nokia’s suite of service diagnostics built on an Nokia 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.

2.3.4.5 SDP Administrative Groups

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

The user first creates the admin groups that are to be used by SDPs on this node:

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

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

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

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

config>service>sdp>sdp-group *group-name*

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

The user then selects which admin groups to include or exclude in a given pseudowire template:

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

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

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

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

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

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

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

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

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

SDP admin groups can be enabled on all router services that make use of the pseudowire template (BGP-AD VPLS service, BGP-VPLS service, BGP-VPWS and FEC129 VLL service). In the latter case, SR OS provides support at the T-PE nodes only.

2.3.4.6 SDP Selection Rules

In the current SDP selection process, all provisioned SDPs with the correct far-end IP address, the correct tunnel-far-end IP address, and the correct service label signaling are considered. The SDP with the lowest admin metric is selected. If more than one SDP with the same lowest metric are found, then the SDP with the highest sdp-id is selected. The type of SDP, GRE or MPLS (BGP/RSVP/LDP) is not a criterion in this selection.

The selection rule with SDP admin groups is modified such that the following admin-group constraints are applied up front to prune SDPs that do not comply:

- If one or more **sdp-include** statement is part of the PW template, then an SDP that is a member of one or more of the included groups will be considered. With the **sdp-include** statement, there is no preference for an SDP that belongs to all included groups versus one that belongs to one or fewer of the included groups. All SDPs satisfying the **admin-group** constraint will be considered and the selection above based on the lowest metric and highest *sdp-id* is applied.
- If one or more **sdp-exclude** statement is part of the PW template, then an SDP that is a member of any of the excluded groups will not be considered.

2.3.4.7 Class-Based Forwarding

2.3.4.7.1 Application of Class-Based Forwarding over RSVP LSPs

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

Figure 12 **Class-Based Forwarding over SDP LSPs**

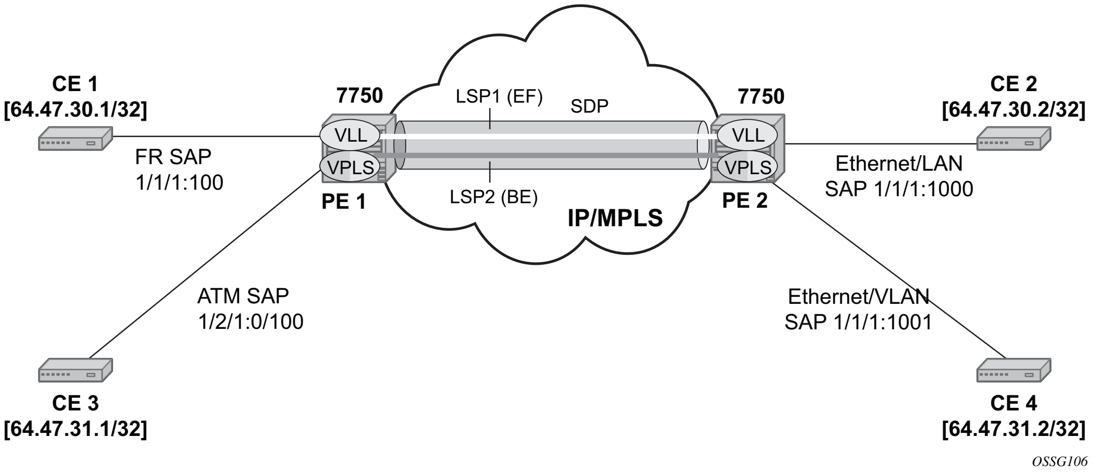

Figure 12 illustrates the use of class-based forwarding to direct packets of a service to specific RSVP or static LSPs that are part of the same SDP based on the packets' forwarding class. The forwarding class of the packet is the one assigned to the packet as a result of applying the ingress QoS policy to the service SAP. The VLL service packets are all classified into the **ef** forwarding class and those that are destined to PE2 are forwarded over LSP1. Multicast and broadcast are classified into the **be** class and are forwarded over LSP2.

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

2.3.4.7.2 Operation of Class-Based Forwarding over RSVP LSPs

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

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

Note: The SDP goes down if the default LSP is down.

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

VLL service packets are forwarded based on their forwarding class only if shared queuing is enabled on the ingress SAP. Shared queuing must be enabled on the VLL ingress SAP if class-forwarding is enabled on the SDP the service is bound to. Otherwise, the VLL packets will be forwarded to the LSP which is the result of hashing the VLL service ID. Since there are eight entries in the ECMP table for an SDP, one LSP ID for each forwarding class, the resulting load balancing of VLL service ID is weighted by the number of times an LSP appears on that table. For instance, if there are eight LSPs, the result of the hashing will be similar to when class based forwarding is disabled on the SDP. If there are fewer LSPs, then the LSPs which were mapped to more than one forwarding class, including the default LSP, will have proportionally more VLL services forwarding to them.

Only user packets are forwarded based on their forwarding class. OAM packets are forwarded in the same way as an SDP with class-based forwarding disabled. In other words, LSP ping and LSP trace messages are queued in the queue corresponding to the forwarding class specified by the user and are forwarded over the LSP being tested. Service and SDP OAM packets, such as service ping, VCCV ping, and SDP ping are queued in the queue corresponding to the forwarding class specified by the user and forwarded over the first available LSP.

Class-based forwarding is not supported for protocol packets tunneled through an SDP. All packets are forwarded over the default LSP.

Class-based forwarding is not supported on a spoke SDP used for termination on an IES or VPRN service. All packets are forwarded over the default LSP.

2.3.4.8 Source IPv4 Address Configuration in GRE SDP and GRE Tunnel

2.3.4.8.1 Introduction and Feature Configuration

When the GRE tunnel is used as part of a provisioned SDP, the following command is relaxed to allow the user to configure a source address for an GRE SDP:

configure>**service**>**sdp**>**local-end** *ip-address*

The default value of the **local-end** parameter is the primary IPv4 address of the system interface. To change the **local-end** address, the SDP must be shut down.

The primary IPv4 address of any local network IP interface, loopback or otherwise, may be used as the source address. The address does not need to match the primary address of an interface which has the MPLS-over-GRE termination subnet configured, unless a GRE SDP or tunnel from the far-end router terminates on this address.

The address of the following interfaces are not supported:

- unnumbered network IP interface
- IES interface
- VPRN interface
- CSC VPRN interface

The following rules apply to the **local-end** command:

- A maximum of 15 distinct address values can be configured for all GRE SDPs under the **configure**>**service**>**sdp**>**local-end** context, and all L2oGRE SDPs under the **config**>**service**>**system**>**gre-eth-bridged**>**tunnel-termination** context.

  The same source address cannot be used in both contexts since an address configured for a L2oGRE SDP matches an internally created interface which is not available to other applications.
- The **local-end** address of a GRE SDP, when different from system, need not match the primary address of an interface which has the MPLS-over-GRE termination subnet configured, unless a GRE SDP or tunnel from the far-end router terminates on this address.

The user must ensure that the local-end address is reachable from the far-end router that terminates the GRE SDP. To help ensure reachability, the interface for this address may be added to IGP or BGP, or a static route may be configured on the far-end router.

The following services can be bound to a GRE SDP when the local-end address is modified:

- VPRN or IES with a spoke-sdp interface
  (**configure**>**service**>**vprn**>**interface**>**spoke-sdp**)
- VPLS with provisioned spoke-sdp
- BGP-AD VPLS and **use-provisioned-sdp** or **prefer-provisioned-sdp** option enabled
- BGP-VPLS and **use-provisioned-sdp** or **prefer-provisioned-sdp** or **prefer-provisioned** option enabled
- Epipe with provisioned spoke-sdp
- Epipe with BGP-VPWS and **use-provisioned-sdp** or **prefer-provisioned-sdp** or **prefer-provisioned** option enabled

For services that support auto-binding to a GRE tunnel, a new CLI command is introduced to configure a single alternate source address per system:

configure>**service**>**system**>**vpn-gre-source-ip** *ip-address*

The default value is the primary IPv4 address of the system interface.

A change to the value of the **vpn-gre-source-ip** parameter can be performed without shutting down the service. Once the new value is configured, the system address will not be used in services that bind to the GRE tunnel.

The primary IPv4 address of any local network IP interface, loopback or otherwise, may be used.

The address of the following interfaces are not supported:

- unnumbered network IP interface
- IES interface
- VPRN interface
- CSC VPRN interface

The following rules apply to the **vpn-gre-source-ip** parameter value:

- This single source address counts towards the maximum of 15 distinct address values per system that are used by all GRE SDPs under the **configure**>**service**>**sdp**>**local-end** context and all L2oGRE SDPs under the **config**>**service**>**system**>**gre-eth-bridged**>**tunnel-termination** context.
- The same source address can be used in both **vpn-gre-source-ip** and **configure**>**service**>**sdp**>**local-end** contexts.
- The same source address cannot be used in both **vpn-gre-source-ip** and **config**>**service**>**system**>**gre-eth-bridged**>**tunnel-termination** contexts because an address configured for a L2oGRE SDP matches an internally created interface that is not available to other applications.
- The **vpn-gre-source-ip** address, when different from system, need not match the primary address of an interface which has the MPLS-over-GRE termination subnet configured, unless a GRE SDP or tunnel from the far-end router terminates on this address.

The following contexts can use a GRE tunnel when the source IP address is modified:

- VPRN service with a SDP (**config**>**service**>**vprn**>**spoke-sdp**)
- VPRN auto-bind-tunnel

The source address cannot be configured for the following services with auto-created GRE-SDP:

- BGP-AD VPLS
- BGP-VPLS
- VGP-VPWS

An alternative solution to bind any one of these services to its own specific GRE SDP with its own source IP address, is to tag a pre-provisioned GRE SDP with a SDP admin-group (**sdp-group** command) and include the admin-group with the PW template binding of this service (**config**>**service**>**pw-template** *policy-id* [**use-provisioned-sdp**]> **sdp-include** *group-name*). The command **prefer-provisioned-sdp** can also be used.

2.3.4.8.2 Feature Operation with T-LDP and BGP Service Label Signaling

The origination function continues to operate as in previous releases. The only change is the ability to insert the user configured address in the source address field of the GRE/IPv4 header as explained in Introduction and Feature Configuration.

Note: The service manager does not explicitly request from the LDP module that an SDP auto-generated T-LDP session for the MPLS-over-GRE SDP uses the source address configured with the **local-end** CLI command. LDP ensures that either a user-configured T-LDP session, or a peer template based auto-created T-LDP session, exists and is connected to the far-end address of the SDP. LDP will use one of these sessions, or will auto-create one using the default local transport address of system.

Consequently, if the source transport address used by the T-LDP control plane session does not match the destination transport address set by the remote PE in the targeted LDP Hello messages, the T-LDP session does not come up.

For example, the setup in Figure 13 will result in both GRE SDP1 and SDP2 to remain down because the targeted Hello adjacency and LDP session will not come up between the two LDP LSRs.

Figure 13 **Mismatched T-LDP Control Plane Parameters**

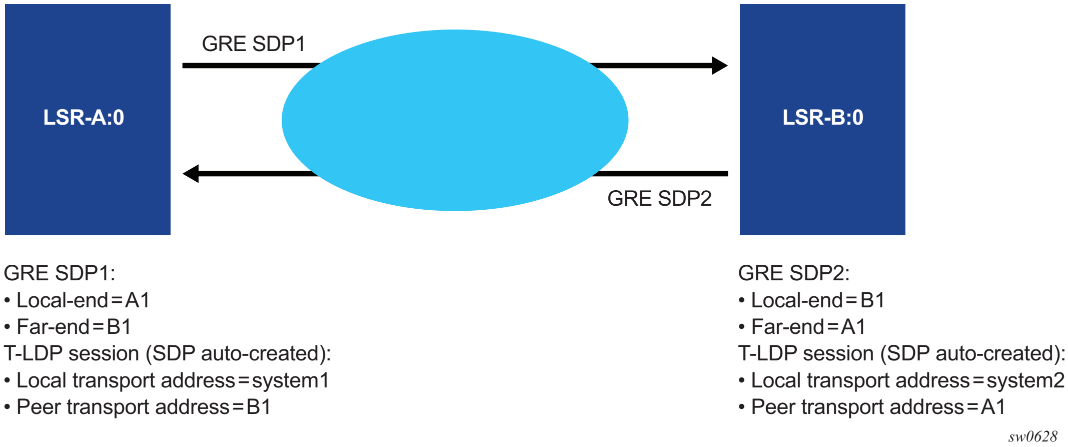
GRE SDP1
LSR-A:0
LSR-B:0
GRE SDP2
GRE SDP1:
• Local-end=A1
• Far-end=B1
T-LDP session (SDP auto-created):
• Local transport address=system1
• Peer transport address=B1
GRE SDP2:
• Local-end=B1
• Far-end=A1
T-LDP session (SDP auto-created):
• Local transport address=system2
• Peer transport address=A1
sw0628

The user must match the local transport address of the T-LDP session to the local-end address of the GRE SDP in both the local and remote PE routers. This can be achieved by manually configuring a T-LDP session to the peer, or by auto-creating a T-LDP session with the targeted peer template feature, and setting the **local-lsr-id** command to the address configured in the **local-end** command of the GRE SDP. In addition, the far-end address must be in a GRE termination subnet at the remote PE and be the primary address of an interface in order for T-LDP to use it as its local LSR ID at the remote PE. Figure 14 shows an example of a correct configuration.

Figure 14 **Proper Setting of T-LDP Control Plane Parameters**

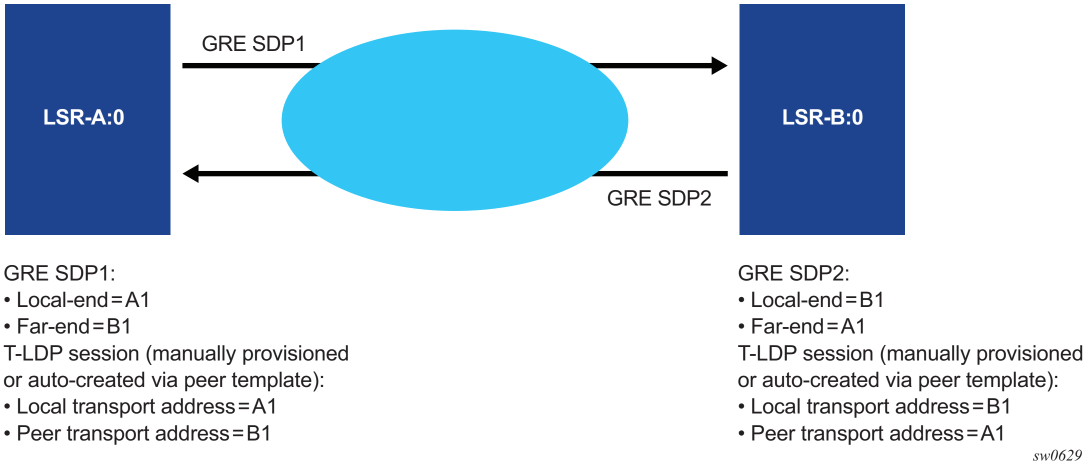
LSR-A:0
GRE SDP1
LSR-B:0
GRE SDP2
GRE SDP1:
• Local-end=A1
• Far-end=B1
T-LDP session (manually provisioned
or auto-created via peer template):
• Local transport address=A1
• Peer transport address=B1
GRE SDP2:
• Local-end=B1
• Far-end=A1
T-LDP session (manually provisioned
or auto-created via peer template):
• Local transport address=B1
• Peer transport address=A1
sw0629

The source address used by the GRE tunnel in the data plane can be different than the local transport address used by T-LDP in the control plane and the GRE SDPs will still come up. For example, the setup in Figure 15 uses at each end the system address for the T-LDP session but uses a loopback interface address as the source address of the GRE SDP.

Figure 15 **Source Address Mismatch between Control and Data Planes**

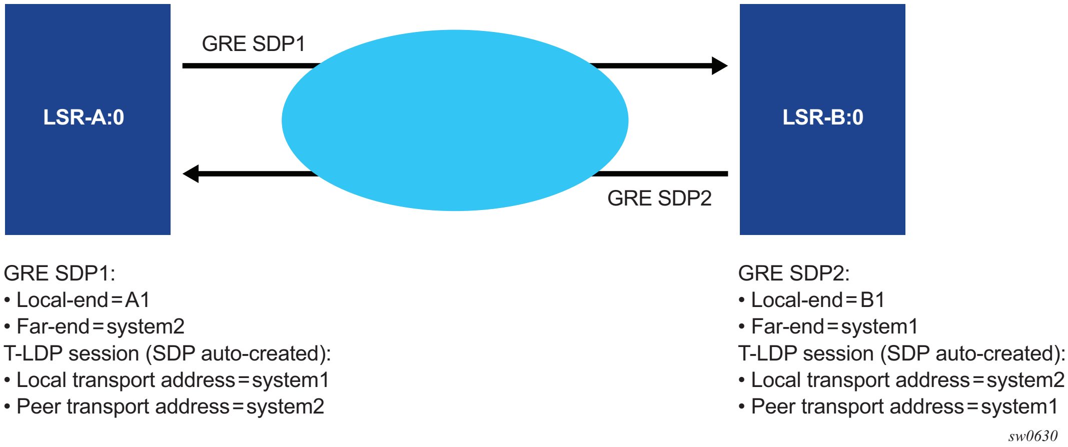
LSR-A:0
GRE SDP1
LSR-B:0
GRE SDP2
GRE SDP1:
• Local-end=A1
• Far-end=system2
T-LDP session (SDP auto-created):
• Local transport address=system1
• Peer transport address=system2
GRE SDP2:
• Local-end=B1
• Far-end=system1
T-LDP session (SDP auto-created):
• Local transport address=system2
• Peer transport address=system1
sw0630

Note: The LDP uses a priority mechanism to select which parameters to use to instantiate a T-LDP session to the same far-end transport address. A manually provisioned T-LDP session overrides one that is signaled using the targeted peer template which overrides one that is auto-created by an SDP. This is done automatically by LDP by issuing, an ad-hoc update to the Hello message to the far-end with the new parameters. As long as the corresponding change is performed at the far-end router to match the local-end parameter change (for example, changing the local transport address requires a change of the far-end transport address in the remote LSR to the same value) the T-LDP session remains up while the Hello adjacency is being synchronized by both LSRs.

The same recommendation applies when the SDP uses BGP for signaling the VC labels of the services. The user must configure the BGP session to the peer and set the **local-address** CLI command under the BGP group context or under the neighbor context to the address configured in the **local-end** command of the GRE SDP.

Replies to OAM messages such as an SDP keep-alive and sdp-ping are sent by the far-end PE using the MPLS-over-GRE encapsulation to the source address of the received OAM message. This means, the source transport address of the T-LDP control plane session or the BGP control plane session is used for the signaling of the VC-label in the local PE. Replies to OAM messages when the VC label is static are sent to the source address of the local PE. In all cases however, the system can properly extract them to the CPM as long as the subnet of that local interface is reachable.

2.3.5 SAP and MPLS Binding Loopback with MAC Swap

SAPs and MPLS SDP bindings within Ethernet services, Epipe and VPLS, may be placed into a loopback mode that allows all packets that arrive on the looped entity to be reflected back into the service. The function is specific to the entity on which the loopback is configured and is non-disruptive to other SAPs and SDP bindings on the same port or LAG.

Epipe and PBB Epipe service constructs support both ingress and egress loopbacks on Ethernet SAPs or MPLS SDP bindings.

VPLS and I-VPLS service constructs support both in ingress and egress loopback on Ethernet SAPs or MPLS SDP bindings.

Do not enable this functionality in the core PBB context because there is no ISID awareness. If this feature is enabled within the core PBB context all traffic that arrives on the B-SAP or B-MPLS binding will be looped back into the PBB context without regard for ISID or customer specific MAC headers.

An ingress loopback configured on the entity will have the following effects on forwarding for the entity:

- Traffic arriving on the entity will be looped back to the same entity, via the fabric.
- Traffic that is attempting to egress that entity from another SAP or SDP binding within the service will be blocked.

Essentially an ingress loopback function will isolate the SAP or MPLS SDP binding from the rest of the service. Figure 16 uses a simple Epipe service to illustrate the various touch points and processing that occurs on a packet that is processed by an ingress loopback as it moves through the network element.

Figure 16 **Ingress Loopback**

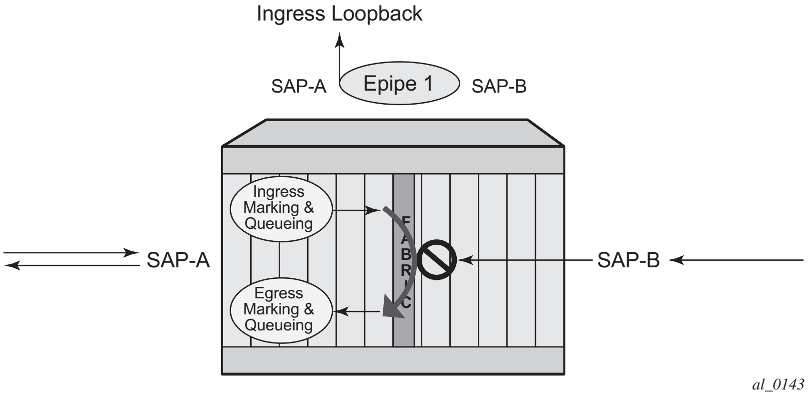

An egress loopback configured on the entity will have the following effects on the forwarding for the entity.

- Traffic that arrives on any service SAP or SDP binding that is forwarded to an egress that is in loopback will be looped back into the service.
- Any traffic that is attempting to gain access to the service from that entity (ingress the network element from the entity) will be dropped.

In the case of the egress loopback, the SAP or MPLS SDP binding is not isolated from the rest of the service it remains part of the service and reflects traffic back into the service. Extreme care must be used when considering the application of an egress loopback in a VPLS or I-VPLS service. Since a VPLS service rely on MAC based forwarding any packet that arrives at an egress loopback will be reflected back into the service and use MAC based forwarding to apply the proper forwarding decision. If this is a live multipoint service with active endpoints this could have very negative effects on the service and the clients connected to this service. Even if the forwarding database is primed any broadcast, unknown or multicast that arrives in the service will arrive on the egress loopback and will be reflected back into the service causing at the very least duplication of all of this type of traffic.

Figure 17 uses a simple Epipe service to illustrate the various touch points and processing that occurs on a packet that is processed by an egress loopback as it moves through the network element. Egress processing will not perform queuing functions on the egress it will only perform the functions of the forwarding plane like remarking.

Figure 17 **Egress Loopback**

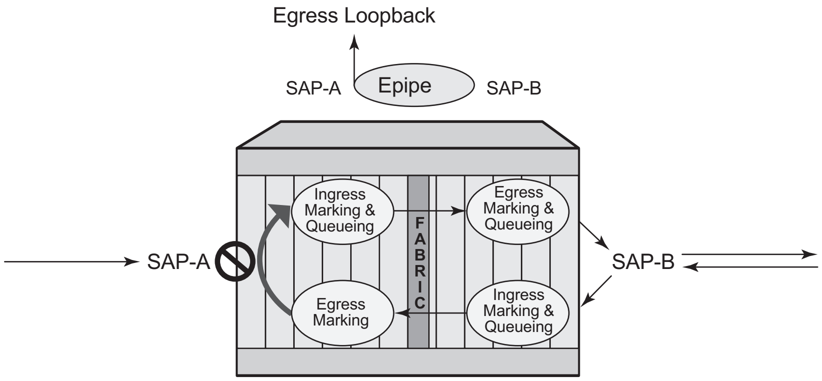
Egress Loopback
SAP-A
Epipe
SAP-B
Ingress Marking & Queueing
Egress Marking & Queueing
FABRIC
SAP-A
SAP-B
Egress Marking
Ingress Marking & Queueing

The operational state of the SAP or MPLS SDP binding will not change as a result of the loopback function. This means a SAP or MPLS SDP binding that is operationally up will not change state strictly because of the loopback be started or stopped. Of course control protocols that are attempting to gain access via the entity that is not allowing packets to enter the service will eventually time out.

Care must be taken when considering the use of control protocols in a service with enabled loopbacks. The operator must be very aware of the impact that interrupting control protocols can have on the state of the SAP. When SAPs are dynamically created using a protocol or a protocol is required to maintain the operational state of the SAP, interruption of this control protocol will cause the SAP to fail. Other SAPs linking their state to a failed SAP will react to that failure as well. This loopback function is per Ethernet SAP or MPLS SDP binding. This means that all traffic that is extracted and sent to the CPM prior to the loopback process will all be looped back to in the direction it was received, or in the case of VPLS, back into the service. All service based control protocols that are included with this service should be removed to ensure the loopback process is handling the packets and not some other function on the node that can extract the control protocol but never respond because the service is block. However, there may be instances where an operator would want to continue to run control protocols for the service during a loopback. For example, Down MEPs on an Ethernet SAP could continue to process ETH-CFM packets if the loopback is on the mate Ethernet SAP and was configured as an egress loopback.

By default no MAC swap functions are performed. Options are available to allow for various MAC swap functions. Table 7 lists the various options and functions based on the configured **mac-swap** and associated options.

Table 7 **MAC-SWAP Configuration and Options**

Configuration		**Reflection with Inbound DA**			
Action	**Options**	**Unicast (Learned)**	**Unicast (Unknown)**	**Broadcast**	**Multicast**
mac-swap	no options	Swap SA to DA Swap DA to SA	Swap SA to DA Swap DA to SA	Drop	Drop
mac-swap	mac	Swap SA to DA Swap DA to SA	Swap SA to DA Swap DA to SA	Swap SA to DA Static MAC= SA	Swap SA to DA Static MAC= SA
mac-swap	mac + all	Swap SA to DA Static MAC= SA	Swap SA to DA Static MAC= SA	Swap SA to DA Static MAC= SA	Swap SA to DA Static MAC= SA
none	none	No swapping	No swapping	No swapping	No swapping

Only the outer Layer 2 header can be manipulated.

In order for the loopback function to operate, the service must be operationally up, and the SAP, port, or LAG must be administratively up. In the case of a LAG, the LAG must have members port that are administratively up. If any of these conditions are not met, the loopback function will fail.

A SAP that is configured for egress loopback is not required to be operationally up, and the cabling does not need to be connected to the port. However, all necessary hardware must be installed in the network element for the ingress packets to be routed to the egress. Ghost ports do not support loopback operations.

An Epipe service will enter an operationally Down state when one of the SAPs is non-operational. The service state will remain or be returned to an operational state if the **ignore-oper-down** command is configured under the non-operational SAP. A VPLS service will remain operational as long as one SAP in the service is operational. However, if the SAP is a VPLS is configured over a LAG, the SAP is removed from the forwarding table if it has a non-operational state, and, consequently, packets will never reach the egress. The **process-cpm-traffic-on-sap-down** command can be configured under the VPLS SAP over a LAG to allow the LAG SAP to be reached even with a non-operational SAP.

If the service state is not operational or the egress SAP is not reachable via the forwarding plane, the traffic will never arrive on the SAP to be looped.

MPLS SDP bindings must be operationally up or the loopback function will fail.

In order to configure this functionality the operator is required to us use the *tools* hierarchy. In this specific case, the loopback tools supporting this functionality may be configured through CLI or through SNMP. However, these commands are never resident in the configuration. This means the loopback will survive high availability events that cause one CPM to change from standby to active, as well as ISSU function or IOM resets (hard or soft). However the function will not survive a complete node reboot.

In the case on SNMP, it is possible to configure a static mac address for the mac swap function without actually invoking the mac-swap. This is not possible through the CLI.

This function requires a minimum of IOM/IMM.

This feature is mutually exclusive with functions that use mirroring.

Figure 18 shows an example for placing sap 1/1/10:2.2 in service id 2 (an Epipe) in an active loopback mode with a mac-swap for all broadcast and multicast destined packets.

Figure 18 **Active Loopback Mode**

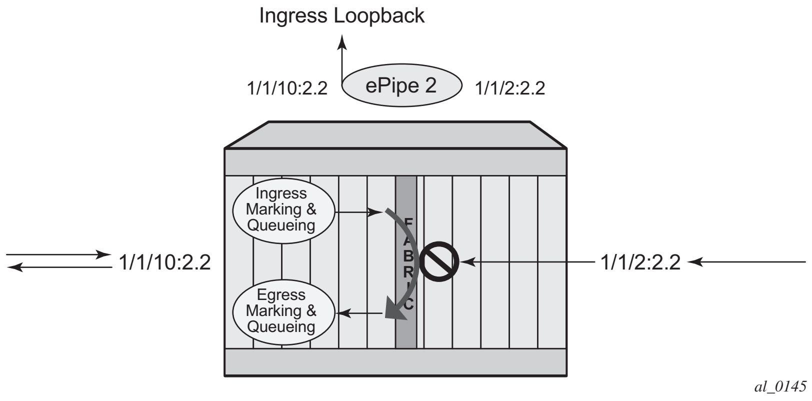

```
show service id 2 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 2                   Vpn Id            : 0
Service Type      : Epipe
Name              : (Not Specified)
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 07/08/2013 09:57:02
Last Mgmt Change  : 07/08/2013 09:56:49
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 2                   SDP Bind Count    : 0
Per Svc Hashing   : Disabled
Force QTag Fwd    : Disabled
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/2:2.2                            qinq         1522    1522    Up   Up
sap:1/1/10:2.2                           qinq         1522    1522    Up   Up
===============================================================================
tools perform service id 2 loopback eth sap 1/1/10:2.2 start ingress mac-swap mac
00:00:00:00:00:88 00:00:00:00:00:88


tools dump service loopback
===============================================================================
Service Ethernet Loopback Points
===============================================================================
Identifier                               Svc ID      Type  Swap    Swap    Oper
                                                           Unicast Mlt/Br
-------------------------------------------------------------------------------
SAP 1/1/10:2.2 qinq                      2           ingr  SA<->DA static  up
-------------------------------------------------------------------------------
No. of Service ethernet loopback points: 1
===============================================================================
tools dump service id 2 loopback sap 1/1/10:2.2
===============================================================================
Service ID 2 SAP 1/1/10:2.2 Loopback
===============================================================================
Identifier (SAP)      : 1/1/10:2.2 qinq
Service ID            : 2
Type                  : Ingress
MAC Swap
  Unicast             : SA<->DA
  Multicast/Broadcast : Static
  Static MAC          : 00:00:00:00:00:88
SAP Oper State        : Up
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
CPM Ingress           : 491790                  46721290

Forwarding Engine Stats
Dropped               : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0
Off. Managed          : 0                       0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
-------------------------------------------------------------------------------
===============================================================================
```

To stop the loopback, a simple **stop** command is required.

```
tools perform service id 2 loopback eth sap 1/1/10:2.2 stop
```

2.3.6 Promiscuous ETH-LBM Mode of Operation

ETH-CFM MEPs support the processing of ETH-CFM PDUs without interrupting the flow of service data. ETH-CFM processing identifies, processes, and responds to the appropriate ETH-CFM PDUs directed at the target MEP using domain-level logic comparison, equal to and lower. This behavior lends itself well to the testing of connectivity, performance monitoring, and path information. The pinpoint ETH-CFM processing logic also lends itself well to service activation testing streams encapsulated in ETH-LBM frames.

The **lbm-svc-act-responder** configuration option allocates additional resources to the associated MEP to process high-speed service activation streams encapsulated in ETH-LBM frames. When a MEP is created with this option, it streamlines the processing of the inbound ETH-LBM frame. This is accomplished by performing basic ETH-CFM header parsing, replacing the inbound ETH-LBM operational code (03) with the outbound ETH-LBR operational code (02), swapping source and destination MAC addresses, and reflecting any Data TLVs and other data contained in the PDU without validation. A MEP configured with the **lbm-svc-act-responder** configuration option will operate in promiscuous ETH-LBM mode.

Promiscuous ETH-LBM mode bypasses some checks and extended functions typically performed by a MEP. In this mode, the MEP will not validate the Layer 2 destination MAC address of the arriving ETH-LBM frame to ensure that it matches the MEP. ETH-LB system statistics and per-MEP statistics, as well as ETH-LB specific counters, will not be incremented. CFM debugging will not be available for these ETH-LB packets.

Only ETH-LBM PDUs at the same domain level as the MEP that is configured with the **lbm-svc-act-responder** function will access the additional resources required to accommodate high-speed service activation processing. Normal processing of ETH-CFM packets will occur for all other ETH-CFM PDUs that arrive on the MEP with the same domain level. The MEP will also process and terminate the lower levels as per normal processing. To ensure proper handling of the service activation stream encapsulated in the ETH-LBM PDU, the level of all ETH-LBM packets in the stream must equal that of the target MEP with the **lbm-svc-act-responder** command.

This mode of operation is supported for Up and Down MEPs in Epipe and VPLS services as well as for base router interfaces. This functionality requires a minimum of FP3 hardware.

2.4 Multi-Service Sites

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

Multi-service sites are anchor points to create an ingress and egress virtual scheduler hierarchy. When a site is created, it must be assigned to a chassis slot or port. When a site is created, it must be assigned to a chassis slot or port with the exception of the 7450 ESS-1 in which the slot is set to 1. When scheduler policies are defined for ingress and egress, the scheduler names contained in each policy are created according to the parameters defined in the policy. Multi-service customer sites exist for the sole purpose of creating a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs).

The scheduler policy association with the customer site normally prevents the scheduler policy from being deleted until after the scheduler policy is removed from the customer site.

When the multi-service customer site is created, an ingress and egress scheduler policy association does not exist. This does not prevent the site from being assigned to a chassis slot or prevent service SAP assignment. After the site has been created, the ingress and egress scheduler policy associations can be assigned or removed at any time.

Each customer site must have a unique name within the context of the customer. Modifications made to an existing site immediately affect all SAPs associated with the site. Changing a scheduler policy association can cause new schedulers to be created and existing policers and queues on the SAPs to no longer be orphaned. Existing schedulers on the site may cease to exist, causing queues relying on that scheduler to be orphaned.

2.5 G.8031 Protected Ethernet Tunnels

G.8031 Protected Ethernet Tunnels is supported only on the 7450 ESS and 7750 SR.

The Nokia implementation of Ethernet Tunnels offers ITU-T G.8031 specification compliance to achieve 50 ms resiliency for failures in a native Ethernet backbone for native Layer 2 networks.

Ethernet Automatic Protection Switching (APS) as defined in ITU-T recommends G.8031 provides a linear 1:1 or 1+1 protection switching mechanism for VLAN-based Ethernet networks. The OS implementation of G.8031 supports 1:1 linear protection through implementation of point-to-point Ethernet Tunnels providing a working and protecting Ethernet circuit, where the path providing the protection is always available through health-monitoring. The 1:1 model is common practice for packet based services since it makes best use of available bandwidth.

Within each path, Y.1731 Maintenance Entity Group (MEG) Endpoints (MEPs) are used to exchange 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 working path by one of the mechanisms triggers to move from working to protecting circuits. 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 OS supports message timers as low as 10 milliseconds so the restoration times are comparable to SONET/SDH. Alternatively, 802.3ah (Ethernet in the First Mile) or simple Loss of Signal can act as a trigger for a protection switch where appropriate.

Revertive or non-revertive behavior can be configured based on service provider environment. Revertive behavior is commonly deployed since it restores the traffic to a predictable state.

Ethernet APS can be configured on any port configured for access mode using dot1q or Q-in-Q encapsulation enabling support for Ethernet APS protected services on the service edge towards the customer site, or within the Ethernet backbone. ELINE, E-LAN, and E-Tree services can be afforded Ethernet APS protection and, although the Ethernet Tunnel providing the protection has a working/protecting path that is presented to the service as a single logical entity to the service layer. The intention of this is to cause minimum disruption to the service during Ethernet APS failure detection and recovery.

Figure 19 **PBB G.8031 Protected Ethernet Tunnel Example**

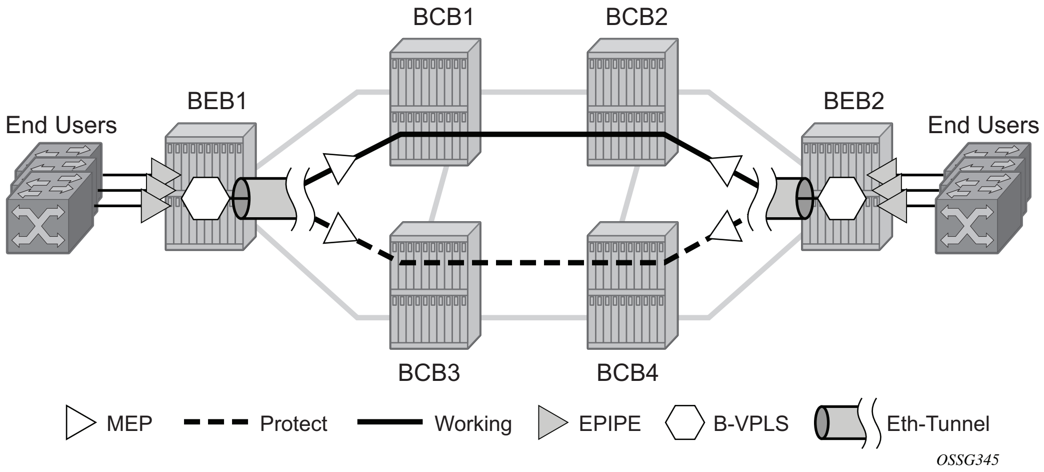

In the implementation, the Ethernet tunnel is a logical interface for a SAP defined Layer 2 service similar to a LAG. The implementation offers ITU G.8031 1:1 compliance as well as some added capabilities such as fate sharing and emulated LAG support.

- Synchronization between services such that both send and receive on the same Ethernet path in stable state.
- Revertive/non-revertive choices.
- Emulated-LAG co-existence.

It is important that the configuration for the various services does not change when a new Ethernet tunneling type is introduced on the backbone side. This is achieved by using a SAP to map the eth-tunnel object into service instance.

The member port and control tag defined under each eth-tunnel path are then used for encapsulating and forwarding the CCMs and the G.8031 PDUs used for protection function, the latter frames being sent only on the secondary path. The configuration of the active path is also used to instantiate the SAP object in the forwarding plane.

If a failure of a link or node affects the primary eth-tunnel path, the services will fail to receive the CC messages exchanged on that path or will receive a fault indication from the link layer OAM module.

For fault detection using CCMs, a number of 3.5 CC intervals 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 received the fault indication, the protection module will declare the associated path down, then sends an indication to the remote protection module to switch the transmit direction to the backup path.

In order to address unidirectional failures, the RDI bit will be set in CC messages transmitted in the reverse direction upon detection of failure at the receiving service. The same applies for link layer OAM. Until the protection switch indication arrives from the remote node, the local node will continue to receive frames from both primary and backup paths to avoid the loss of in-flight packets.

In case of direct connectivity between the nodes, there is no need to use Ethernet CCM messaging for liveliness detection. Link level detection mechanisms like LoS (Loss of Signal) or IEEE 802.3ah link layer OAM can be used to detect link or nodal failure. This can be achieved by not provisioning a MEP on the primary path.

Using the Ethernet Tunnel as a building block for Ethernet APS protection it is possible to provide different protection schemes with different fate-dependency; or indeed to mix protected and non-protected services on the same physical port.

The simplest model is the fate-independent model where each Ethernet Tunnel supports its own protection using Y.1731 CCMs for example. In this case a single VLAN Tag may be used for control and data traffic. In cases where Ethernet Tunnels can be guaranteed to share a common physical path, it is possible to implement a fate-sharing model. This approach provides the advantage of reducing the amount of Ethernet OAM signaling because only one control tag determines the fate of many user tags.

Epipe using BGP-MH site support for Ethernet tunnels (see the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information) offers an enhancement to Ethernet tunnels enabling an Ethernet edge device using G.8031 to support multi-chassis redundancy for Epipe services. The G.8031 device configuration is standard on the Ethernet edge device, but the active link is controlled by BGP-multihoming just as with VPLS services. This Epipe feature offers a standards-based alternative for multihomed access.

Figure 20 **PBB Fate-Independent Ethernet Tunnels**

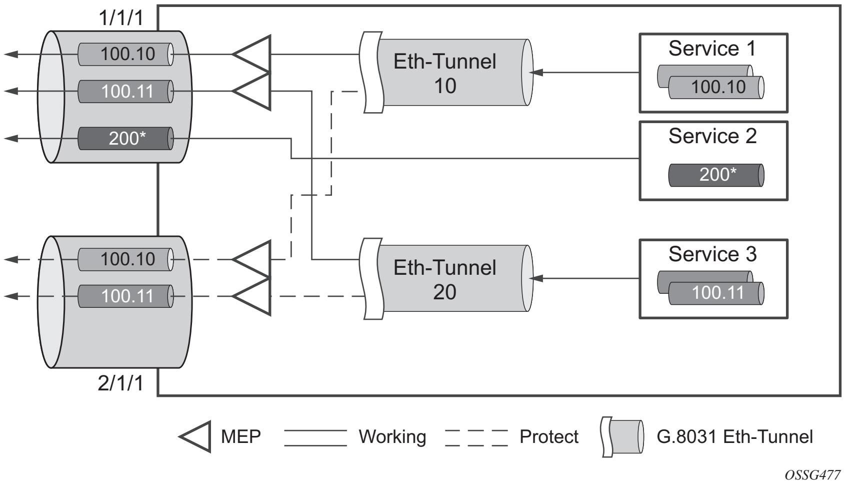

Figure 21 **PBB Fate Sharing Ethernet Tunnels**

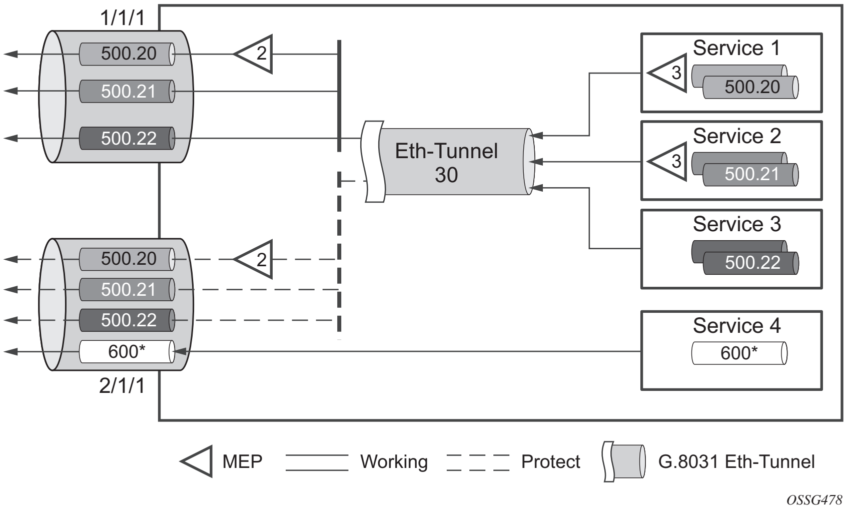

One of the advantages of access redundancy using Ethernet APS is that because it operates at the VLAN level protection mechanisms can be varied between services supported on the physical port. For example, it is possible to provide a protected service for “Premium” customers and also provide non-protected services for “Standard” users on the same physical port.

2.5.1 OAM Considerations

Ethernet CFM can be enabled on each individual path under an Ethernet tunnel. 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 as low as 10 msec. Different CCM intervals can be supported on the primary and secondary paths in an Ethernet tunnel.

MEPs can still be configured under the services independent of the Ethernet Tunnels.

The following rules control the interaction between the MEP defined under the eth-tunnel path and the MEP defined in the service:

- The down MEPs configured on the eth-tunnel paths must be lower level than any down.
- MEPs configured on the associated SAP. The same applies for Virtual MEPs associated with services such as BVPLS. Checks are provided to prevent the user from configuring anything that violates the above rule. An error message is generated to indicate the mismatch.
- Other service MEPs (up direction, down higher levels) are allowed with no restriction.
- Any down MEP on the associated SAP will transmit only over the active path entity.

2.5.2 QoS Considerations

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

The protection mode '8031-1to1' (default) activates only the primary path at any point in time, guaranteeing the use of the desired QoS resources.

Ethernet tunnel 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 bouncing of the Ethernet tunnel, 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.

2.5.3 Mirroring and Lawful Intercept Considerations

Mirroring and Lawful Intercept (LI) cannot use the eth-tunnel as a source. Also, a SAP configured on an eth-tunnel cannot be used as mirror destination. The CLI blocks the above options. The SAP configured on the eth-tunnel, a filter associated with it and the member ports in the **eth-tunnel> path** context can be used as mirror and LI source.

2.5.4 Support Service and Solution Combinations

The Ethernet tunnels are supported Layer 2 service VLL, VPLS and B-VPLS instances. The following considerations apply:

- Only ports in access or hybrid mode can be configured as eth-tunnel path members. The member ports can be located on the same or different IOMs or MDAs.
- Dot1q and QinQ ports are supported as eth-tunnel path members.
- The same port cannot be used as member in both a LAG and an Ethernet Tunnel but LAG emulation is supported.
- A mix of regular and multiple eth-tunnel SAPs and pseudowires can be configured in the same services.
- Split horizon groups in VPLS and BVPLS are supported on eth-tunnel SAPs. The use of split horizon groups allows the emulation of a VPLS model over the native Ethernet core, eliminating the need for P-MSTP.
- LAG Emulation offers another method offering MSTP or P-MSTP over Ethernet Tunnels.
- MC-LAG access multi-homing into services is supported in combination with Ethernet tunnels.

2.5.5 LAG Emulation using Ethernet Tunnels

Ethernet Tunnels can provide G.8031 Ethernet APS protection as described in G.8031 Protected Ethernet Tunnels, or they can operate in a load-sharing manner providing an emulated LAG function. Moreover, as multiple Ethernet Tunnels can be provisioned on the same physical link(s), it is possible that two physical links could support one or more Ethernet Tunnels supporting APS protection for protected services whilst concurrently supporting one or more Ethernet Tunnels in load-sharing mode for non-protected services.

When Ethernet Tunnels have the protection type set to load-sharing, the precedence is configured to secondary, making the tunnels equal in order to implement load-sharing capability. A path threshold parameter allows the load-sharing group to be declared down if the number of paths drops equal to or lower than the threshold value. The ‘lag-emulation’ context provides access to conventional LAG parameters such as the adapt-qos mode (link, port-fair or distributed bandwidth distribution) and per-fp-ing-queuing to ensure that only one ingress queue is instantiated for every physical link supported on the same FP complex.

A typical use case for LAG emulation is to allow unprotected Ethernet services to capitalize on the LAG capability. RSTP and MSTP can also be used to network VPLS or B-VPLS over the Ethernet tunnels. LAG Emulation is also recommended when you use BGP-MH site support for Ethernet tunnels.

2.6 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 (Ethernet-ring) is built on Ethernet OAM and often referred to as Ring Automatic Protection Switching (R-APS).

Ethernet-rings are supported on VPLS SAPs (VPLS, I-VPLS, B-VPLS). VPLS services supporting Rings SAPs can connect to other rings and Ethernet service using VPLS and R-VPLS SAPs. Ethernet-rings enables rings for core network or access network resiliency. A single point of interconnection to other services is supported. The Ethernet-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 Ethernet-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 SR OS implementation can be used for interconnecting access rings and to provide traffic engineered backbone rings.

Ethernet-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 SR OS implementation supports DOT1q, QinQ and PBB encapsulation for data ring instances. The control channel supports dot1q and QinQ encapsulation. The control channel can support DOT1Q while the data channels use queuing if the global **config>system**>**ethernet**>**new-qinq-untagged-sap** command is enabled.

2.6.1 Overview of G.8032 Operation

R-APS messages that carry the G.8032 protocol are sent on dedicated protocol VLAN called the Ethernet Ring Protection (ERP) 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 the ring 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. 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 than 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 22 depicts this operational state:

Figure 22 **0-1 G.8032 Ring in the Initial State**

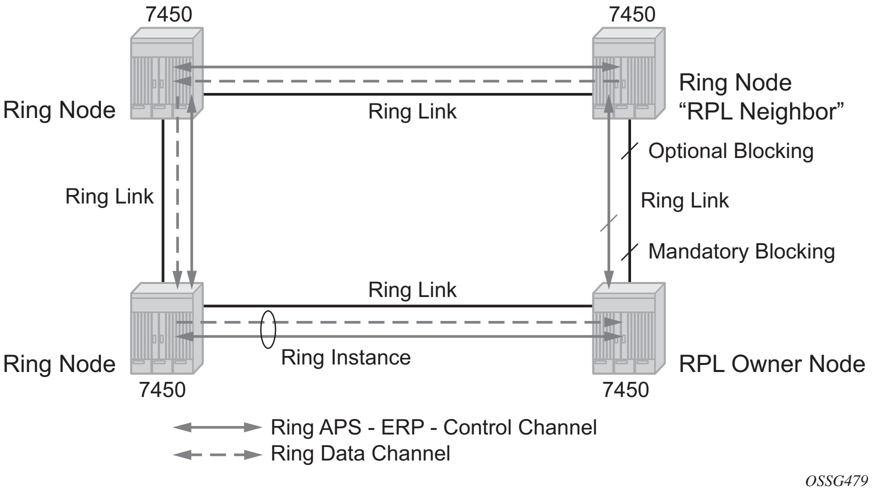

When a ring failure occurs, a node or nodes detecting the failure (enabled by Y.1731 OAM CC monitoring) send R-APS message in both directions. This allows the nodes at both ends of the failed link to block forwarding to the failed link preventing it from becoming active. In revertive mode, the RPL Owner 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. Figure 23 depicts the failed link scenario.

Figure 23 **0-1 G.8032 Ring in the Protecting State**

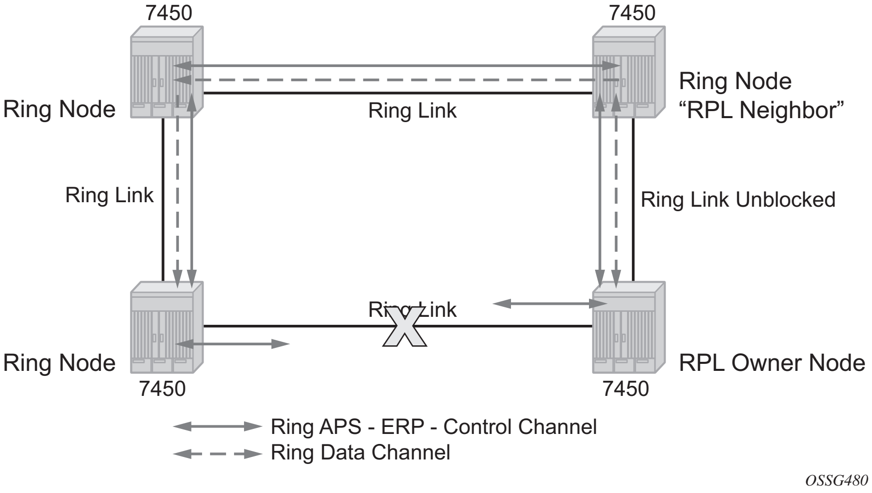

Once the failed link recovers, the nodes that blocked the link again send the R-APS messages indicating no failure this time. This in turn triggers RPL owner to block the RPL link and indicate the blocked RPL link the ring in R-APS message, which when received by the nodes at the recovered link cause 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 router supports message timers as low as 10 milliseconds (also 100 ms) so the restoration times are comparable to SONET/SDH. 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 Ethernet-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 services (using PBB Epipes with the B-VPLS configured with Ethernet rings), E-LAN services, and E-Tree data 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 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 Ethernet-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 CCMs exchanged on that segment or will receive a fault indication from the Link Layer OAM module. CCMs are optional but MEPs are always configured to provide G.8032 control. Note that the forwarding of CCMs and G.8032 R-APS messages continues in the control VPLS even if the service or its SAPs are administratively shut down. The Ethernet ring instance can be shut down if it is desired to stop the operation of the ring on a node.

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 router implementation allows flooding of traffic during the flushing interval to expedite traffic recovery.

Figure 24 illustrates a resilient Ring Service. In the 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 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 closed loop. A ring can restore connectivity when any one node or link fails to all remaining nodes within the 50 ms transfer time (signaling time after detection).

Figure 24 **0-3 Ring Example**

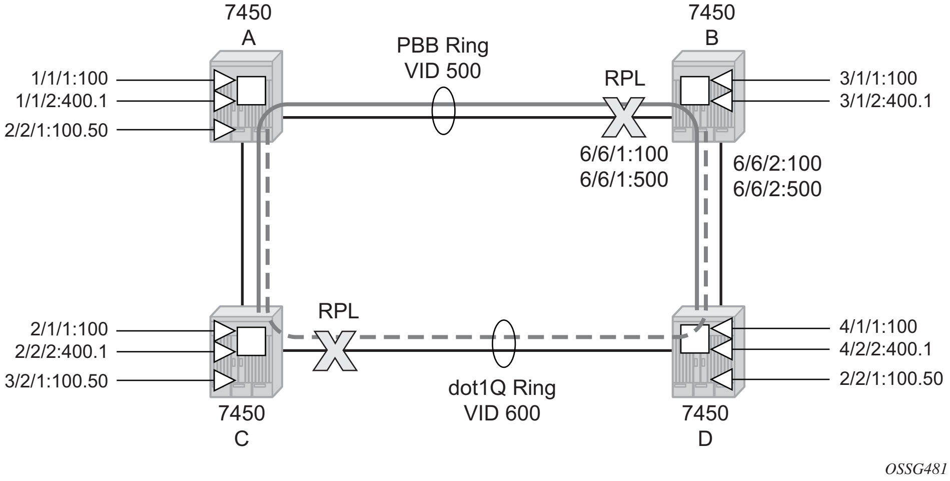

Sample Configuration:

```
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  6/6/1 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
    no shutdown
    exit
    service
        vpls 10 customer 1 create // Ring APS SAPs
            description "Ring Control VID 100"
                sap 6/6/1:100 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 6/6/1:500 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 vpls 1000 i-vpls // CPE traffic
    sap 3/1/1:100 create // CPE SAP
        pbb
            backbone-vpls 40 isid 1000
              exit
        exit
    no shutdown
    exit
    service vpls 1001 i-vpls // CPE traffic
    sap 3/1/2:400.1 create   // CPE SAP
        pbb
            backbone-vpls 40 isid 1001
              exit
        exit
    no shutdown
    exit
```

2.6.2 Ethernet Ring Sub-Rings

Ethernet Sub-Rings offer a dual redundant way to interconnect rings. The router 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 25 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 25 **0-4 G.8032 Sub-Ring**

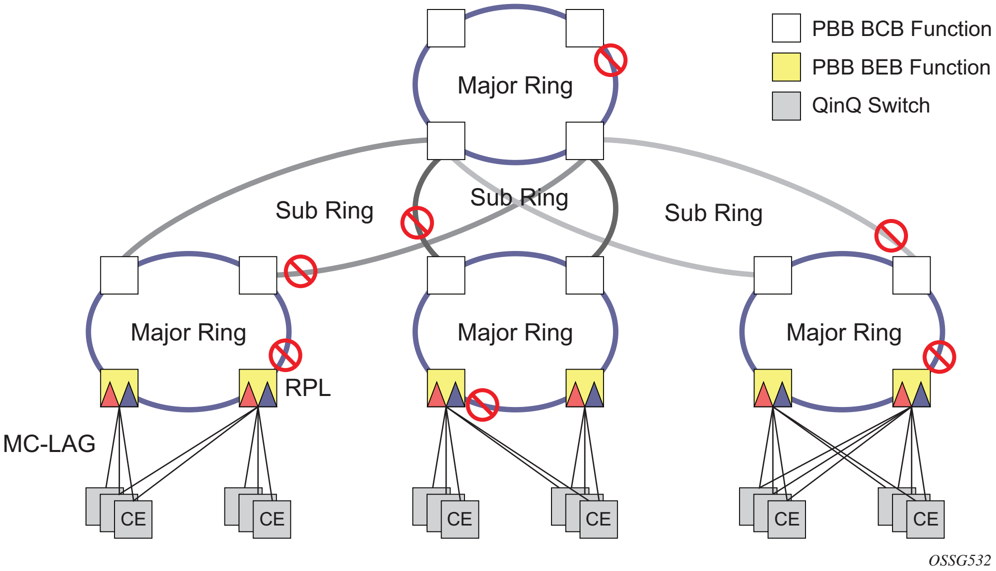

Sub-Rings and Major Rings run similar state machines for the ring logic, however there are some differences. When Sub-Rings protect a link, the flush messages are propagated to the major ring. (A special configuration allows control of this option on the router.) When major rings change topology, the 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.

2.6.2.1 Virtual and Non-Virtual Channel

The 7450 ESS, 7750 SR, and 7950 XRS support both the virtual channel and non-virtual channel for sub-ring control communication. In the virtual channel mode, a dedicated VID, other than the major ring RAPs control channel is configured as a data instance on the major ring. This allows the sub-ring control messages and state machine logic to behave similar to a major ring. In the non-virtual channel mode, the sub-ring is only connected by the RAPs control channels on the sub-ring itself. This mode offers slightly less redundancy in the RAPs messaging than the virtual channel mode since sub-ring RAPs messages are not propagated across the major ring. When non-virtual link is configured, the protocol allows RPL messages over the sub-ring blocked link.

Figure 26 **0-5 Sub-Ring Configuration Example**

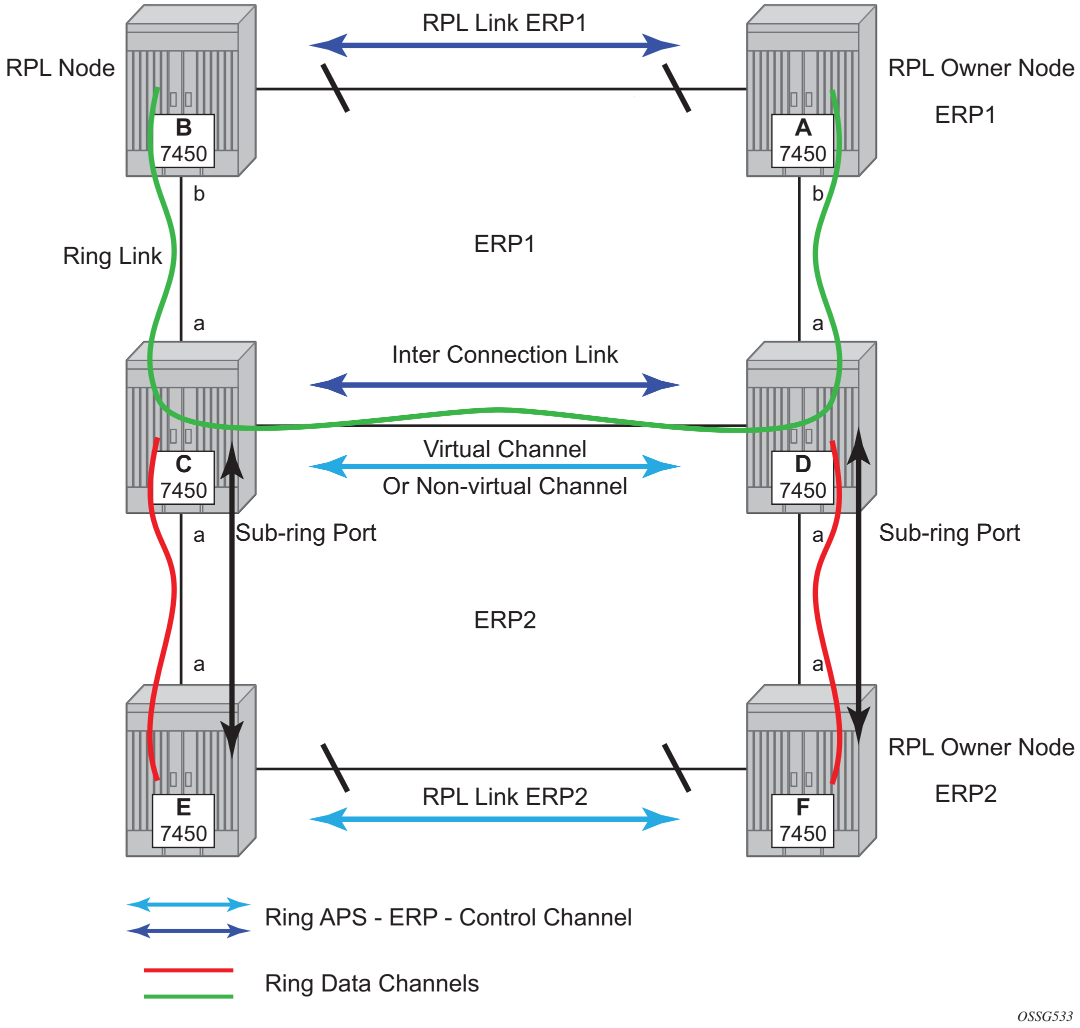

Sub-ring configuration is similar to major ring configuration and consists of three parts: Ethernet-ring instance configuration, control VPLS configuration, and data VPLS configuration (data instance or data channel). The Ethernet-ring configuration of a sub-ring is tied to a major ring and only one path is allowed.

Note: A split horizon group is mandatory to ensure that Sub-Ring control messages from the major ring are only passed to the sub-ring control.

As with a major ring, the forwarding of CCMs and G.8032 R-APS messages continues in the control VPLS even if the service or its SAPs are administratively shut down. The Ethernet ring instance can be shut down if it is desired to stop the operation of the ring on a node.

The data VPLS can be configured on the major ring, and in the example, shares the same VID (SAP encapsulation) on both the major ring and the sub-ring to keep data on the same VLAN ID everywhere.

Note: Like other services in the router, the encapsulation VID is controlled by SAP configuration and the association to the controlling ring is by the Ethernet-ring, ring-id.

The following illustrates a sample sub-ring configuration on Node C:

```
eth-ring 2
        description "Ethernet Sub Ring on Ring 1"
        sub-ring virtual-link // Using a virtual link
            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
```

If the sub-ring had been configured as a non-virtual-link, the sub-ring configuration above and on all the other sub-ring nodes for this sub-ring would become:

```
        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
```

If the sub-ring had been configured as a non-virtual-link, the configuration above would then become:

```
vpls 100 customer 1 create
    description "Control VID 100 for Ring 2 Interconnection"
    sap 1/1/3:100 eth-ring 2 create
        stp shutdown
    exit
    no shutdown
exit
```

The 7450 ESS, 7750 SR, and 7950 XRS allow for a special configuration of the non-virtual link sub-ring that can be homed to a VPLS service illustrated in Figure 27. This is an economical way to have a redundant ring connection to a VPLS service. This is currently supported only for dot1Q and QinQ sub-rings and only on LDP based VPLS. The primary application for this is access rings that require resiliency. This configuration shows the configuration for a sub-ring at an interconnection node without a virtual channel and interconnected to a VPLS. A VPLS service 1 is used to terminate the ring control. The Ethernet ring data SAP appears in the associated LDP based VPLS service 5.

Figure 27 **0-6 Sub-Ring Homed to VPLS**

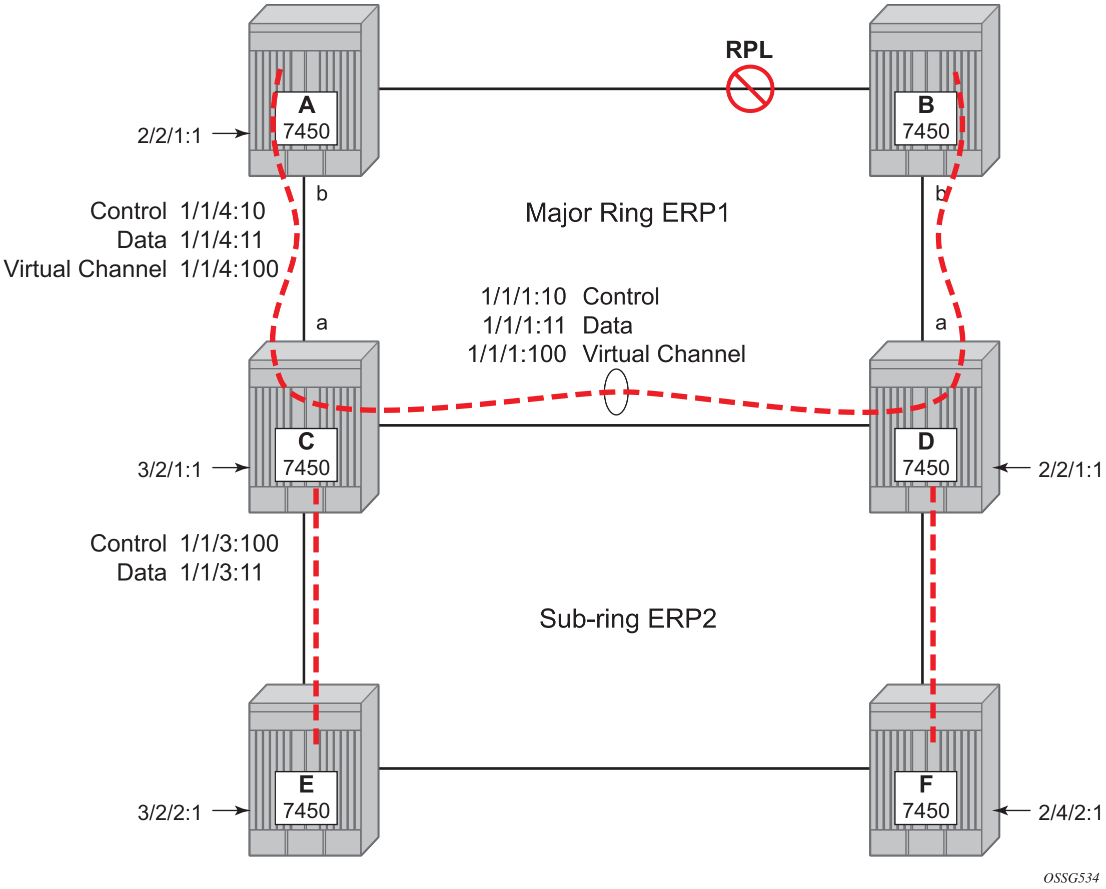

The following is a sample sub-ring configuration for VPLS (at PE1):

```
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

# Configuration for the ring control interconnection termination:
  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 LDP based 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
```

Ethernet-rings and sub-rings offer a way to build a scalable resilient Ethernet transport network. Figure 28 illustrates a hierarchical ring network using PBB where dual homed services are connected to a PBB based Ethernet ring network.

Figure 28 **0-7 Multi Ring Hierarchy**

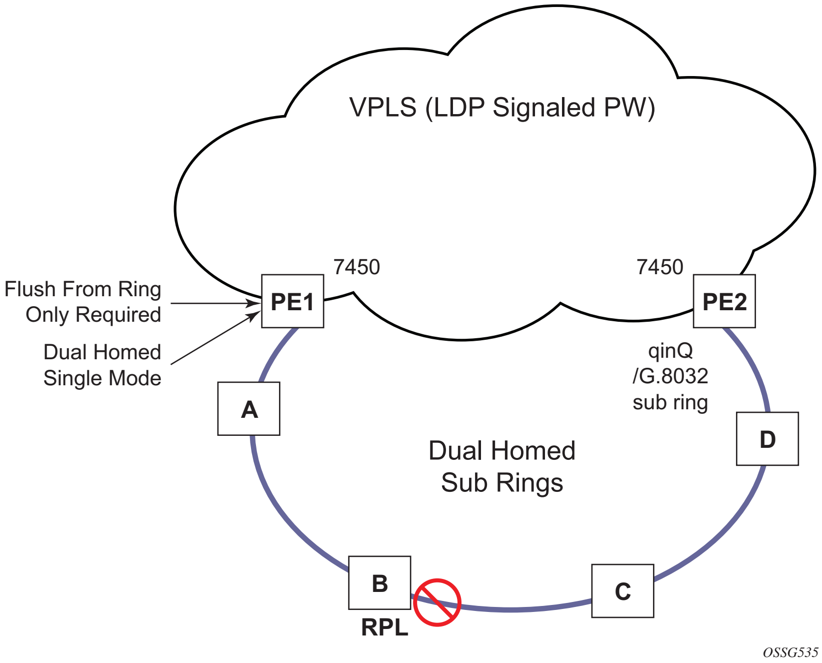

The major rings are connected by sub-rings to the top level major ring. These sub-rings require virtual channel and will not work with non-virtual channel. Ring flushing is contained to major rings, or in the case of a sub-ring link or node failure, to the sub-ring and the directly attached major rings.

2.6.2.2 LAG Support

Ethernet-rings support LAG on Ethernet rings SAPs. However, the use of LAG impact the response time for resiliency. In many cases, the use of multiple ring instances each on a single link may be more suitable from a resiliency and QoS standpoint than using LAG on Ethernet rings in a given topology. If sub 100ms response is not required, LAG is an option for Ethernet-rings.

2.6.3 OAM Considerations

Ethernet CFM is enabled by configuring MEPs on each individual path under an Ethernet ring. Only down MEPs can be configured on each of them and optionally, CCM sessions can be enabled to monitor the liveliness of the path using interval of 10 or 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 (LOS) for example, which is controlled by configuring **no-ccm-enable**.

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

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

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 sections of the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN*.

2.6.4 Support Service and Solution Combinations

The Ethernet rings are supported Layer 2 service, VPLS, I-VPLS, R-VPLS and B-VPLS instances. The following considerations apply:

- Only ports in access mode can be configured as Ethernet-ring paths. The ring ports can be located on the same or different media adapter cards.
- Dot1q and QinQ ports are supported as Ethernet-ring path members.
- A mix of regular and multiple Ethernet-ring SAPs and pseudowires can be configured in the same services.

2.7 Internal Objects Created for L2TP and NAT

Some services such as L2TP LNS (L2TP Network Server) and NAT (Network Address Translation) automatically create service objects for internal use. In particular, an IES service with ID 2147483648 is created. In that service, or in configured VPRN services, service objects such as interfaces, SAPs and related objects can be automatically created for internal use.

Named objects reserved for internal use have a name that starts with "_tmnx_". Objects with a numeric identifier created for internal use have an identifier from a reserved range.

The general rules for objects reserved for internal use:

- Will appear in CLI show commands and MIB walks output;
- Will appear in the output of **info detail** commands but will never be in the output of **admin save** [**detail**].

It may be possible to enter the CLI node of such an object, but it is not possible to change anything. It may also be possible to set the value of one of its objects to the current value with SNMP, but it will never be possible to change any value.

2.8 Ethernet Unnumbered Interfaces

The ability to configure Ethernet Unnumbered interfaces has been added to support some service types for IPv4. The unnumbered interface capability has been available for other interface types on SR OS. Unnumbered Ethernet allows point-to-point interfaces to borrow the address from other interfaces such as system or loopback interfaces.

This feature enables unnumbered interfaces for some routing protocols (IS-IS and OSPF). Support for routing is dependent on the respective routing protocol and service. This feature also adds support for both dynamic and static ARP for unnumbered Ethernet interfaces to allow interworking with unnumbered interfaces that may not support dynamic ARP.

The use of unnumbered interface has no effect on IPv6 routes but the unnumbered command must only be used in cases where IPv4 is active (IPv4 only and mixed IPv4/IPv6 environments). When using an unnumbered interface for IPv4, the loopback address used for the unnumbered interface must have IPv4 address. Also, interface type for the unnumbered interface will automatically be point-to-point.

2.9 ECMP and Weighted ECMP for Services Using RSVP and SR-TE LSPs

ECMP over MPLS LSPs refers to spraying packets across multiple named RSVP and SR-TE LSPs within the same ECMP set. The ECMP-like spraying consists of hashing the relevant fields in the header of a labeled packet and selecting the next-hop tunnel based on the modulo operation of the output of the hash and the number of ECMP tunnels. Only LSPs with the same lowest LSP metric can be part of the ECMP set.

In weighted ECMP, the load-balancing weight of the LSP is normalized by the system and then used to bias the amount of traffic forwarded over each LSP. The weight of the LSP is configured using the **config**>**router**>**mpls**>**lsp**>**load-balancing-weight** *weight* and **config**>**router**>**mpls**>**lsp-template**>**load-balancing-weight** *weight* commands.

If one or more LSPs in the ECMP set do not have **load-balancing-weight** configured, and the ECMP is set to a specific next hop, regular ECMP spraying is used.

Weighted ECMP is supported for VPRN Layer 3 services. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for more information.

Weighted ECMP is supported for the following Layer 2 services over RSVP-TE tunnels:

- Epipe VLLs
- Ipipe VLLs
- LDP VPLS
- BGP-AD VPLS with provisioned SDPs

Class Based Forwarding (CBF) and weighted ECMP are mutually exclusive for VLL and VPLS services.

For services that use an explicitly configured SDP, weighted ECMP is configured under the SDP used by the service with the **config**>**service**>**sdp**>**weighted-ecmp** command. By default, weighted ECMP is disabled.

For VLL and VPLS services, when a service uses a provisioned SDP on which weighted ECMP is configured, a path is selected based on the configured hash. Paths are then load-balanced across LSPs within an SDP according to the normalized LSP weights. Additional fields may be taken into account for VPLS services based on the commands in the **config**>**service**>**load-balancing** context.

2.10 Network Group Encryption (NGE)

This section provides information to configure network group encryption (NGE) on the VSR.

2.10.1 NGE Overview

The network group encryption (NGE) feature enables end-to-end encryption of MPLS services, Layer 3 user traffic, and IP/MPLS control traffic. NGE is an encryption method that uses a group-based keying security architecture, which removes the need to configure individual encryption tunnels to achieve network-wide encryption.

NGE relies on the NSP NFM-P to manage the network and apply encryption to specific MPLS services, Layer 3 user traffic, or control plane traffic depending on the security requirements of the network. Operators designate traffic types that require added security and then apply NGE to those traffic types using the NSP NFM-P. The NSP NFM-P also acts as the network-wide NGE key manager, downloading encryption and authentication keys to nodes and performing hitless rekeying of the network at operator-defined intervals. For more information about managing NGE within a network, see the *NSP NFM-P User Guide*.

Figure 29 shows an NGE network with NSP NFM-P services, control plane configuration, and key management.

Figure 29 **NGE Network with NSP NFM-P Management**

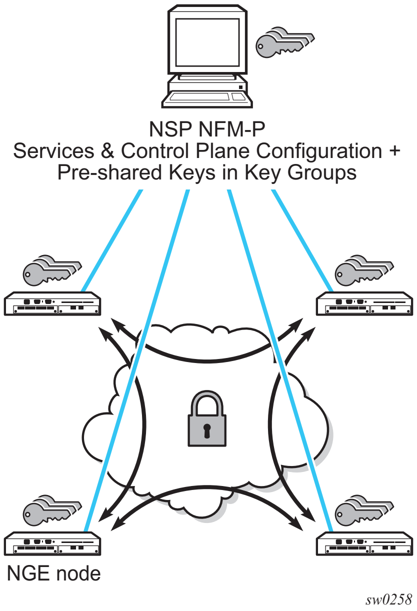
NSP NFM-P
Services & Control Plane Configuration +
Pre-shared Keys in Key Groups
NGE node
sw0258

NGE provides three main types of encryption to secure an IP/MPLS network:

- SDP encryption — MPLS user plane encryption enabled on MPLS tunnels (SDPs) supporting VPRN or IES services using spoke SDPs, VPLS using spoke or mesh SDPs, routed VPLS into VPRN, Epipes, and Cpipes
- VPRN encryption — MP-BGP-based VPRN-level encryption using auto-bind of LDP, GRE, MPLSoUDP, RSVP-TE, and MPLS (LDP or RSVP-TE) tunnels
- router interface — Layer 3 user plane and control plane encryption

NGE is supported on the following adapter cards and platforms:

- VSR-I
- VSR-a CN8
- VSR-a NGE

2.10.1.1 NGE Key Groups and Encryption Partitions

NGE allows a tiered approach to managing encryption keys in a network using key groups by configuring services, or router interfaces to use specific key groups depending on security policies for the service and network topology.

Figure 30 shows a typical application of NGE key-group partitioning in which there are several critical levels (tiers) of security that need to be considered. In this example, the protection of Distribution Automation and Field Area Network (DA/FAN) equipment are less critical than the Transmission or Distribution Substation network equipment. Ensure that nodes more at risk of a security breach do not contain more critical information than is necessary. Therefore, encryption keys for the sensitive portions of the network (such as control center traffic) should not be available on nodes that are at risk. The NGE feature enables operators to partition and distribute encryption keys among different services, NGE domains, or nodal groups in a network. NGE partitions are enabled by configuring different key groups per security partition and applying those key groups as needed.

Figure 30 **Key-Group Partitioning**

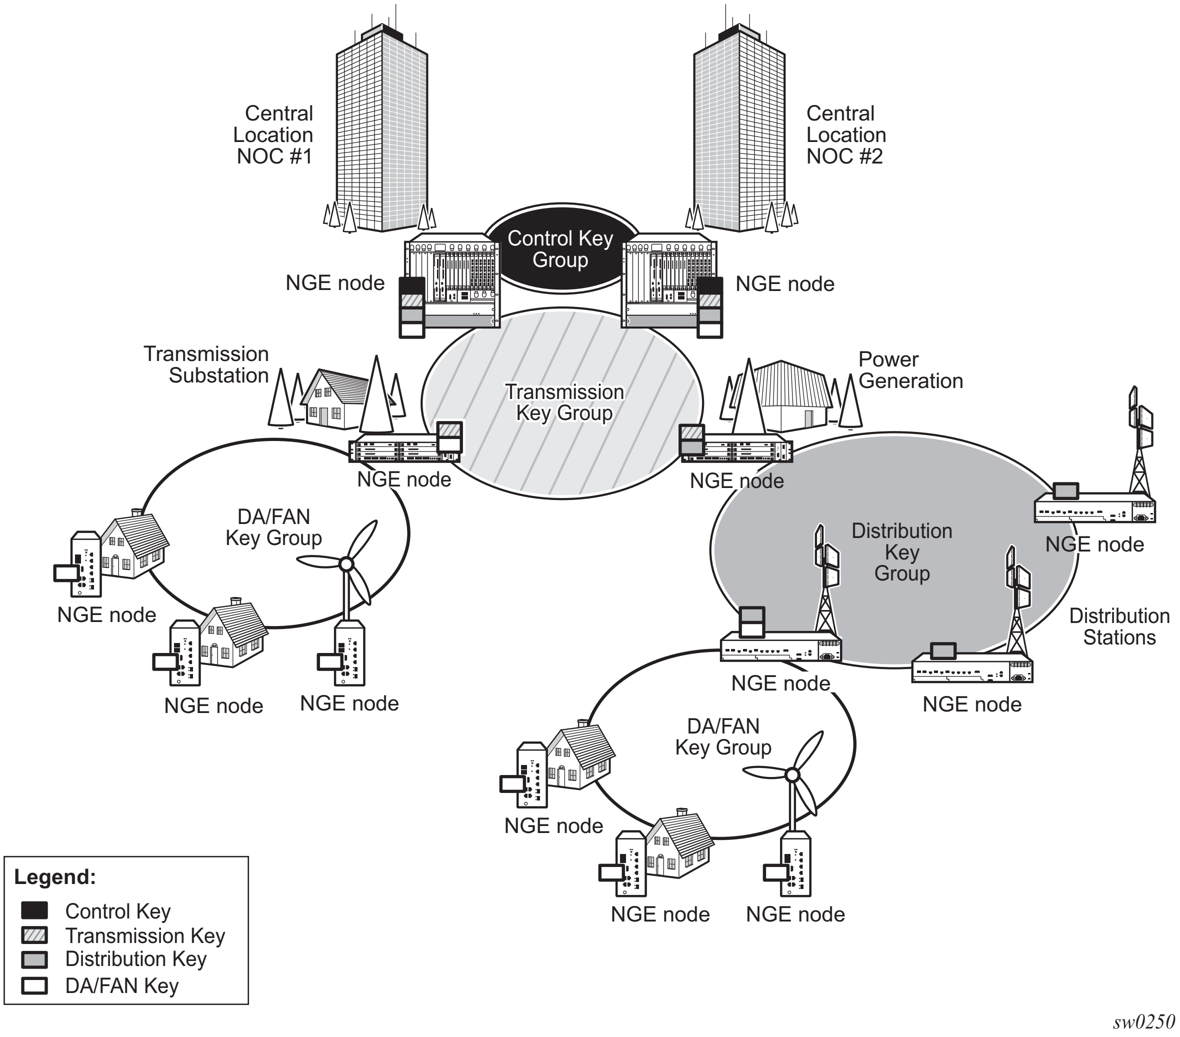

Another application of key-group partitioning allows different parts of an organization to have their own method of end-to-end communication without the need to share encryption keys between each organization. If two partitions need to communicate between themselves, gateway nodes configured with both key groups allow inter-organization traffic flows between the key group partitions, as needed.

2.10.1.2 Network Services Platform Management

The NGE feature is tightly integrated with the NSP NFM-P. The following functions are provided by the NSP NFM-P:

- managing and synchronizing encryption and authentication keys within key groups on a network-wide basis
- configuring NGE on MPLS services and managing associated key groups
- configuring NGE on router interfaces and managing associated key groups
- coordinating network-wide rekeying of key groups

The NSP NFM-P acts as the key manager for NGE-enabled nodes and allocates the keys in key groups that are used to perform encryption and authentication. The NSP NFM-P ensures that all nodes in a key group are kept in synchronization and that only the key groups that are relevant to the associated nodes are downloaded with key information.

The NSP NFM-P performs network-wide hitless rekeying for each key group at the rekeying interval specified by the operator. Different key groups can be rekeyed at different times if desired, or all key groups can be rekeyed network-wide at the same time.

For more information on NSP NFM-P management, refer to the “Service Management” section in the *NSP NFM-P User Guide*.

2.10.2 Key Groups

Key groups are used to organize encryption keys into distinct groups that allow a user to partition the network based on security requirements. A key group contains the following elements:

- an encryption algorithm—see Key Group Algorithms
- an authentication algorithm—see Key Group Algorithms
- a list of security associations (SAs)—see Security Associations

- an active outbound SA—see Active Outbound SA

Figure 31 illustrates the use of key groups (KGs), security associations (SAs), and security parameter indices (SPIs). The VSR-1 and VSR-2 both have the same set of key groups configured. One path uses key group 1 (KG1) and the other uses key group 2 (KG2). Each key group contains the elements listed above. Key group 1 has four live keys, SPI_1 through SPI_4, and SPI_3 is the active outbound SA. The active outbound SA is identified by its SPI, and this SPI is embedded in the NGE packet.

Each SA listed in a key group, indexed by an SPI, specifies a single key for encryption and a single key for authentication. Packets transmitted or received that reference a particular SPI use the keys in the SA for that SPI when performing encryption and authentication.

Before enabling encryption, key groups must be configured on the node. Only after a key group is configured can it be assigned to an SDP or VPRN services.

Figure 31 **Key Groups and a Typical NGE Packet**

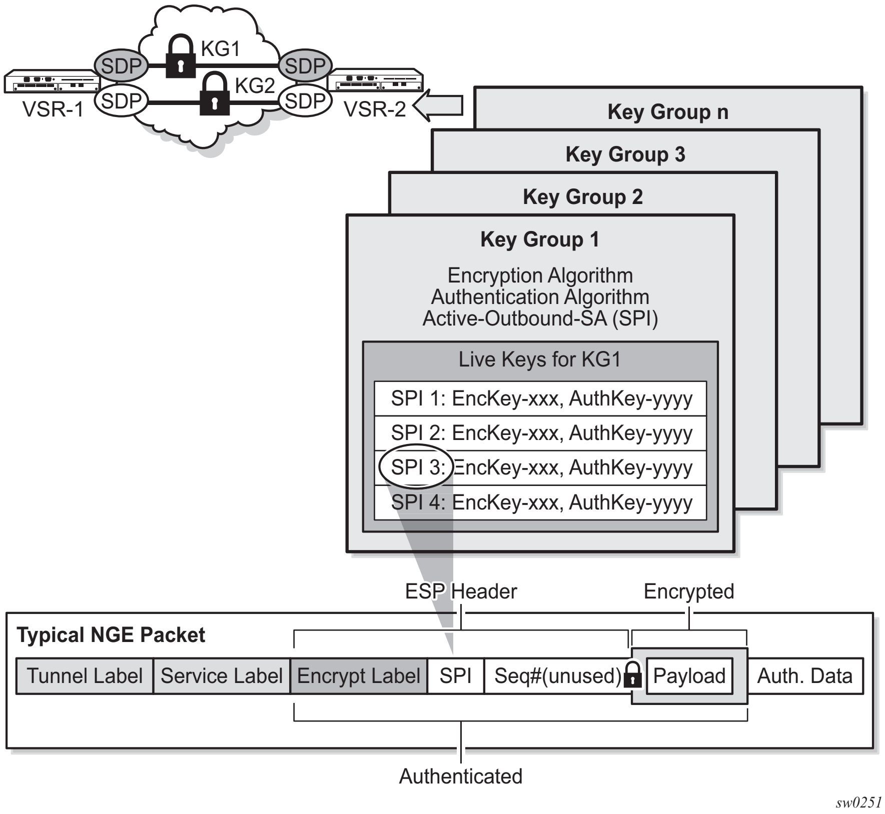

2.10.2.1 Key Group Algorithms

All SAs configured in a key group share the same encryption algorithm and the same authentication algorithm. The size and values required by a particular key depend on the requirements of the algorithms selected (see lists below). One encryption algorithm and one authentication algorithm must be selected per key group.

Encryption algorithms available per key group include:

- AES128 (a 128-bit key, requiring a 32-digit ASCII hexadecimal string)
- AES256 (a 256-bit key, requiring a 64-digit ASCII hexadecimal string)

Authentication algorithms available per key group include:

- HMAC-SHA-256 (a 256-bit key, requiring a 64-digit ASCII hexadecimal string)
- HMAC-SHA-512 (a 512-bit key, requiring a 128-digit ASCII hexadecimal string)

Encryption and authentication strengths can be mixed depending on the requirements of the application. For example, 256-bit strength encryption can be used with 512-bit strength authentication.

The configured algorithms cannot be changed when there is an existing SA configured for the key group. All SAs in a key group must be deleted before a key group algorithm can be modified.

Key values are not visible in CLI or retrievable using SNMP. Each node calculates a 32-bit CRC checksum for the keys configured against the SPI. The CRC can be displayed in the CLI or read by SNMP. The purpose of the CRC is to provide a tool to check consistency between nodes, thereby verifying that each node is set with the same key values while keeping the actual key values hidden.

2.10.2.1.1 Encapsulating Security Payload

The NGE feature uses the Encapsulating Security Payload (ESP) protocol according to IETF RFC 4303. ESP maintains data integrity, ensuring privacy and confidentiality for encrypted traffic.

The ESP protocol used by NGE relies on symmetric ciphers, meaning that the same key is used for encryption and decryption. The NGE node supports Cipher Block Chaining (CBC) encryption mode. Block ciphers used by NGE include:

- AES128 with a 128-bit key using 128-bit blocks
- AES256 with a 256-bit key using 128-bit blocks

For authentication, the integrity check value (ICV) size is as follows:

- HMAC-SHA-256 (16 bytes or 128 bits)
- HMAC-SHA-512 (32 bytes or 256 bits)

2.10.2.2 Security Associations

Each key group has a list of up to four security associations (SAs). An SA is a reference to a pair of encryption and authentication keys that are used to decrypt and authenticate packets received by the node and to encrypt packets leaving the node.

For encrypted ingress traffic, any of the four SAs in the key group can be used for decryption if there is a match between the SPI in the traffic and the SPI in the SA. For egress traffic, only one of the SAs can be used for encryption and is designated as the active outbound SA. Figure 31 illustrates these relationships.

As shown in Figure 31, each SA is referenced by an SPI value, which is included in packets during encryption and authentication. SPI values must be numerically unique throughout all SAs in all key groups. If an SPI value is configured in one key group and an attempt is made to configure the same SPI value in another key group, the configuration is blocked.

Note: Keys are entered in clear text using the **security-association** command. After configuration, they are never displayed in their original, clear text form. Keys are displayed in an encrypted form, which is indicated by the system-appended **crypto** keyword when an **info** command is run. The NGE node also includes the **crypto** keyword with an **admin**>**save** operation so that the NGE node can decrypt the keys when reloading a configuration database. For security reasons, keys encrypted on one node are not usable on other nodes (that is, keys are not exchangeable between nodes).

2.10.2.2.1 Active Outbound SA

The active outbound SA is specified by the SPI referencing the specific SA used to encrypt and authenticate packets egressing the node for the SDP or service using the key group. The SPI value for the active outbound SA is included in the ESP header of packets being encrypted and authenticated.

2.10.3 Services Encryption

NGE provides the ability to encrypt MPLS services using key groups that are configured against these services. These services include:

- VLL service (Epipe and Cpipe)
- VPRN service using Layer 3 spoke-SDP termination
- IES service using Layer 3 spoke-SDP termination
- VPLS service using spoke and mesh SDPs
- routed VPLS service into a VPRN or IES
- MP-BGP-based VPRNs

For services that use SDPs, all tunnels may be either MPLS LSPs (RSVP-TE, LDP, or static LSP), or GRE or MPLSoUDP tunnels.

For MP-BGP services, **auto-bind-tunnel** is supported using LDP, GRE, MPLSoUDP, RSVP-TE, or MPLS (LDP or RSVP-TE).

2.10.3.1 Services Encryption Overview

NGE adds a global encryption label to the label stack for encrypting MPLS services. The global encryption label must be a unique network-wide label; in other words, the same label must be used on all nodes in the network that require NGE services. The label must be configured on individual nodes before NGE can become operational on those nodes.

The global encryption label is used to identify packets that have an NGE-encrypted payload and is added to the bottom of the stack. This allows network elements such as LSRs, ABRs, ASBRs, and RRs to forward NGE packets without needing to understand NGE or to know that the contents of these MPLS packets are encrypted. Only when a destination PE receives a packet that needs to be understood at the service layer does the PE check for an encryption label, and then decrypt the packet.

After the global encryption label is set, it should not be changed. If the label must be changed without impacting traffic, all key groups in the system should first be deleted. Next, the label should be changed, and then all key groups should be reconfigured.

The NSP NFM-P helps to coordinate the distribution of the global encryption label and ensures that all nodes in the network are using the same global encryption label.

Figure 32 illustrates the NGE MPLS, GRE, or MPLSoUDP label stack.

Figure 32 **NGE MPLS/GRE/MPLSoUDP Label Stack**

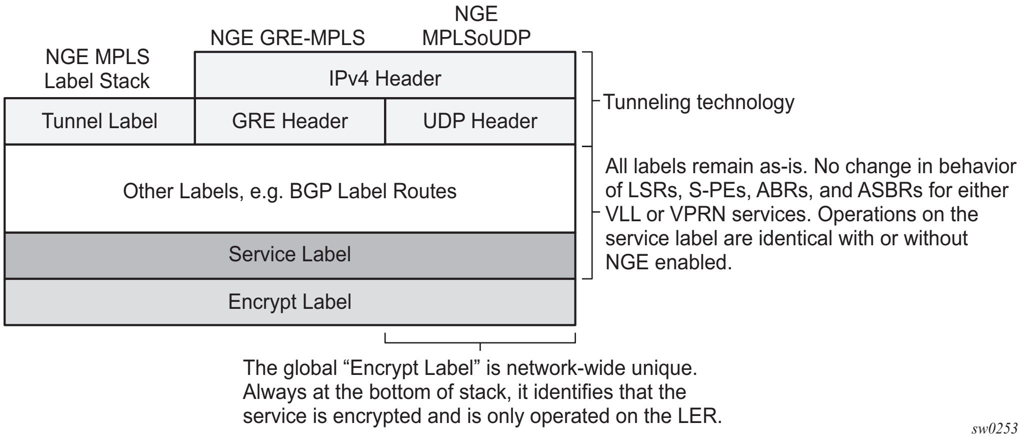
NGE MPLS Label Stack
NGE GRE-MPLS
NGE MPLSoUDP
IPv4 Header
Tunnel Label
GRE Header
UDP Header
Tunneling technology
Other Labels, e.g. BGP Label Routes
Service Label
Encrypt Label
All labels remain as-is. No change in behavior of LSRs, S-PEs, ABRs, and ASBRs for either VLL or VPRN services. Operations on the service label are identical with or without NGE enabled.
The global “Encrypt Label” is network-wide unique. Always at the bottom of stack, it identifies that the service is encrypted and is only operated on the LER.
sw0253

Figure 33 illustrates VPRN and PW (with control word) packet formats using NGE encryption.

Figure 33 **NGE Encryption and Packet Formats**

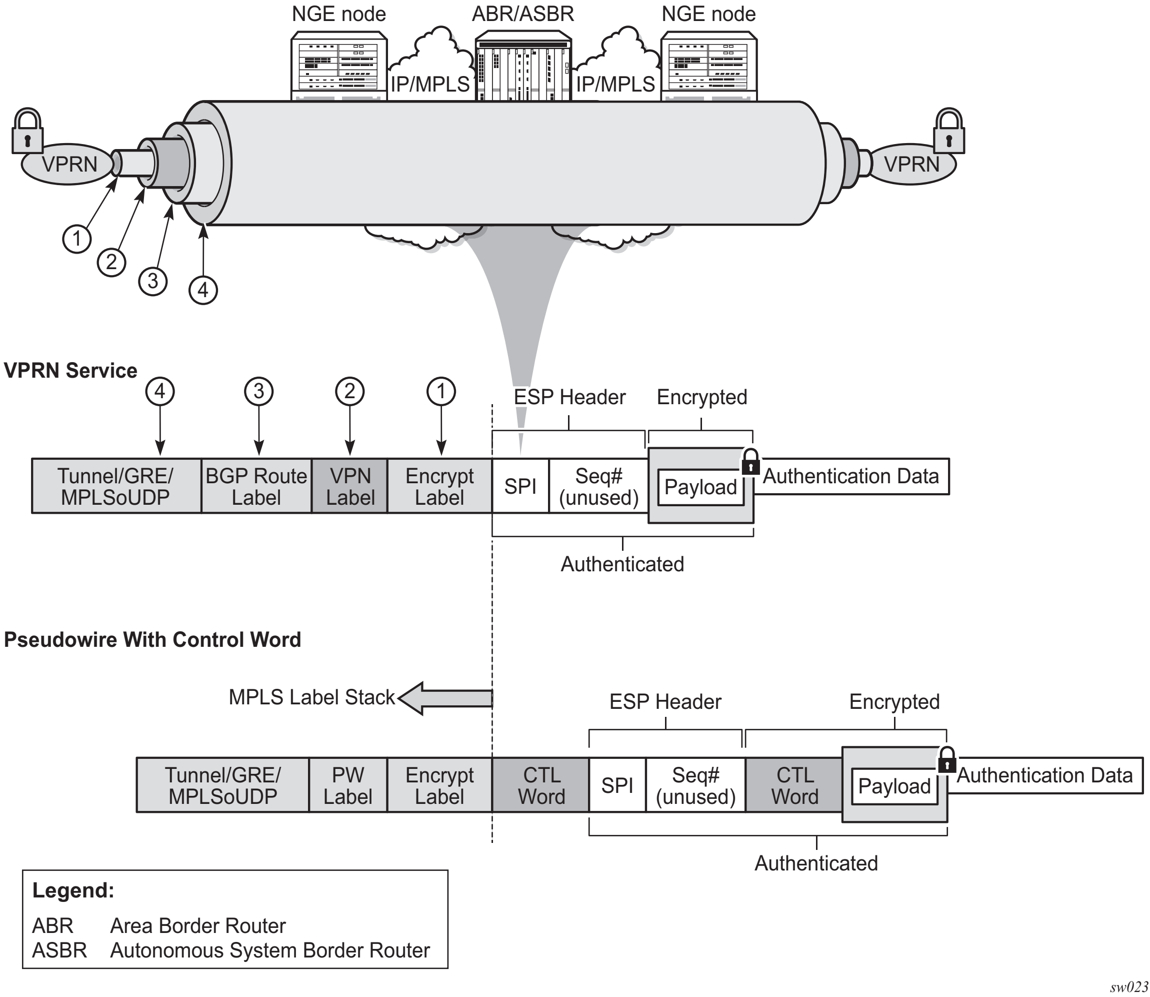

2.10.3.2 Assigning Key Groups to Services

Assigning key groups to services requires configuring an inbound and outbound key group for directional processing on a per-service basis (see Figure 34).

Figure 34 **Inbound and Outbound Key-Group Assignments**

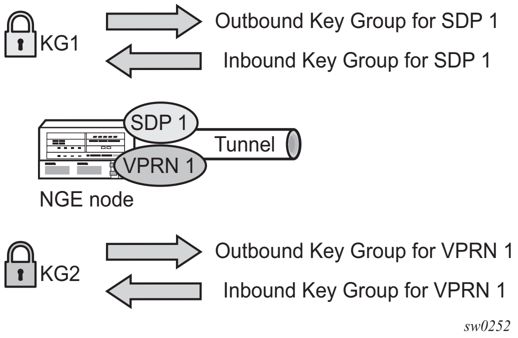
KG1
Outbound Key Group for SDP 1
Inbound Key Group for SDP 1
SDP 1
Tunnel
VPRN 1
NGE node
KG2
Outbound Key Group for VPRN 1
Inbound Key Group for VPRN 1
sw0252

The outbound key group identifies which key group to use for traffic that egresses the node for the service. The inbound key group ensures that ingress traffic is using the correct key group for the service.

If the inbound key group is not set, the node ensures that packets are either unencrypted or are using one of the valid key groups configured in the system.

In most deployment scenarios, the inbound and outbound key groups are identical; however, it is possible to configure different key groups as the outbound and the inbound key groups, as this is not checked by the node.

Including an inbound and outbound direction when assigning key groups to services allows users to:

- gracefully enable and disable NGE for services
- move services from one key-group domain to another domain without halting encryption

The NGE feature makes use of the NSP NFM-P to help manage the assignment of key groups to services on a network-wide basis. Refer to the *NSP NFM-P User Guide* for more information.

2.10.3.3 Pseudowire Switching for NGE Traffic

For VLL services, the NGE node supports PW switching of encrypted traffic from one PW to another. There are three scenarios that are supported with regard to PW switching of traffic:

- PW switch using the same key group

When a PW is using an encrypted SDP, the PW may be switched to another PW that is also using an encrypted SDP, where both SDPs are in the same key group. In this case, to perform the PW switch, the NGE node leaves the encrypted payload unchanged and swaps the labels as needed for passing traffic between PWs.

- PW switch using different key groups

  When a PW is using an encrypted SDP, the PW may be switched to another PW that is also using an encrypted SDP, where both SDPs are in different key groups. In this case, the NGE node decrypts the traffic from the first SDP by using the configured key group for that SDP, and then re-encrypts the traffic by using the egress SDP's key group egress SPI ID.

- PW switch between an encrypted and unencrypted PW

  When traffic is switched from an encrypted PW to an unencrypted PW, the traffic is decrypted before it is sent. The converse occurs in the reverse direction (that is, traffic from an unencrypted PW to an encrypted PW gets encrypted before it is sent).

Refer to “Pseudowire Switching” in the *7450 ESS, 7750 SR, 7950 XRS, and VSR OAM and Diagnostics Guide* for more information.

2.10.3.4 Pseudowire Control Word for NGE Traffic

The control word is a configurable option for PWs and is included in PW packets when it is configured. When **control-word** is enabled and NGE is used, the datapath creates two copies of the CW. One CW is both encrypted and authenticated, and is inserted after the ESP header. The other CW is not encrypted (clear form) and is inserted before the ESP header.

For cases where PW switching is configured, the NGE node ensures—in the CLI and with SNMP—that both segments of the PW have consistent configuration of the control word when encryption is being used.

2.10.3.5 VPRN Layer 3 Spoke-SDP Encryption and MP-BGP-based VPRN Encryption Interaction

The encryption configured on an SDP used to terminate the Layer 3 spoke SDP of a VPRN always overrides any VPRN-level configuration for encryption.

- When VPRN encryption is enabled, all routes resolved using MP-BGP are encrypted or decrypted using the VPRN key group.

- When Layer 3 spoke-SDP encryption is enabled, all routes resolved using the Layer 3 interface are encrypted or decrypted using the SDP's key group.

Some examples are as follows.

- If a VPRN is enabled for encryption while a Layer 3 spoke SDP for the same VPRN is using an SDP that is not enabled for encryption, then traffic egressing the spoke SDP is not encrypted.
- If a VPRN is disabled for encryption while a Layer 3 spoke SDP for the same VPRN is using an SDP that is enabled for encryption, then traffic egressing the spoke SDP is encrypted.
- If a VPRN is enabled for encryption using key group X, while a Layer 3 spoke SDP for the same VPRN is using key group Y, then traffic egressing the spoke SDP is encrypted using key group Y.

The commands used for these scenarios are **config**>**service**>**sdp**>**encryption-keygroup** and **config**>**service**>**vprn**>**encryption-keygroup**.

2.10.3.6 NGE and RFC 3107

When RFC 3107 is enabled on the node and NGE traffic is crossing the Area Border Router (ABR) between two VPRN domains, the same key group must be used between the two domains.

Note: It is the responsibility of the network operator to ensure key group consistency across the (ABR).

2.10.4 NGE Packet Overhead and MTU Considerations

NGE adds overhead packets to services. Table 8 shows the additional overhead for the worst-case scenario of MPLS services encryption. Table 9 shows the additional overhead for the worst-case scenario of router interface. Additional overhead depends on which encryption and authentication algorithms are chosen.

Table 8 **NGE Overhead for MPLS**

Item	Number of Bytes
Encryption label	4

Table 8 **NGE Overhead for MPLS (Continued)**

Item	Number of Bytes
ESP	24
ICV	32
Padding	17
Control word copy	4
Total	**81**

For MP-BGP-based VPRNs, the total is 77 bytes because the control word copy is not required.

Table 9 **NGE Overhead for Router Interface**

Item	Number of Bytes
ESP	24
ICV	32
Padding	17
Total	**73**

For Layer 3 packets for router interface encryption, the total is 73 bytes because the encryption label and control word copy are not required.

The overhead values in Table 8 must be considered for services that are supported by NGE.

➡ **Note:** Currently, the port MTU has a default value of 1572 bytes. This value is too low for outbound traffic when NGE is enabled. Users must configure new MTU values to adjust for the overhead associated with NGE, as outlined in Table 10 for MPLS-based and GRE-based services. For details on configuring MTU, refer to the “MTU Configuration Guidelines” section in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide*.

The calculations in Table 10 show how NGE overhead affects SDP MTU and service MTU values for MPLS-based, GRE-based, and VPRN-based services. The calculations are with and without NGE enabled.

Table 10 **Accounting for NGE Overhead SDP and Service MTU — Calculation Examples**

Service Type	MTU Values With and Without NGE Enabled
MPLS-based services	**SDP MTU (MPLS)** = 1572 (network port MTU) – 14 (Ethernet header) – 4 (outer label) – 4 (inner label) = 1550 bytes (without NGE enabled) => 1469 bytes (with NGE enabled)
	Service MTU (MPLS) considerations with NGE enabled • Layer 3 spoke IP MTU (MPLS) = 1469 – 14 (inner Ethernet header) = 1455 bytes • PW spoke SDP MTU (MPLS) = SDP MTU = 1469 bytes
GRE-based services	**SDP MTU (GRE)** = 1572 (network port MTU) – 14 (Ethernet header) – 20 (IP header) – 4 (GRE header) – 4 (inner label) = 1530 bytes (without NGE enabled) => 1449 bytes (with NGE enabled)
	Service MTU (GRE) considerations with NGE enabled • Layer 3 Spoke IP MTU (GRE) = 1449 – 14 (inner Ethernet header) = 1435 bytes • PW Spoke MTU (GRE) = SDP MTU = 1449 bytes
VPRN-based services	Each interface inherits its MTU from the SAP or spoke SDP to which it is bound and the MTU value can be manually changed using the **ip-mtu** command.
	MP-BGP-based VPRN services The MTU of the egress IP interface is used. When NGE is enabled on a VPRN service, customers must account for the additional 77 bytes of overhead needed by NGE for any egress IP interface used by the VPRN service.

When an unencrypted Layer 3 packet ingresses the node and routing determines that the egress interface is a router interface NGE-enabled interface, the node calculates whether the packet size will be greater than the MTU of the egress interface after the router interface NGE overhead is added. If the packet cannot be forwarded out from the network interface, an ICMP message is sent back to the sender and the packet is dropped. Users must configure new MTU values to adjust for the overhead associated with NGE.

If an IP exception ACL that matches the ingressing packet exists on the egress interface, the MTU check applied to the ingress packet includes the router interface NGE overhead. This is because the ingress interface cannot determine which IP exceptions are configured on the egress interface, and therefore the worst-case MTU check that includes the router interface NGE overhead is performed.

2.10.5 Statistics

Statistics specific to NGE are counted for the following main areas:

- key group
- SPI
- MDA
- service

2.10.6 Remote Network Monitoring Support

Remote network Monitoring (RMON) can be used in conjunction with NGE statistics to provide event and alarm reporting. This can be used by customers to detect security breaches of NGE traffic flows and provide real-time reporting of such events.

Threshold crossing alerts and alarms using RMON are supported for SNMP MIB objects, including NGE.

2.10.7 Configuration Notes

This section describes NGE configuration guidelines and caveats. For more information about configuring NGE using the NSP NFM-P, see the *NSP NFM-P User Guide*.

Consider the following caveats when performing NGE configuration tasks.

- The authentication and encapsulation keys must contain the exact number of hexadecimal characters required by the algorithm used. For example, using sha256 requires 64 hexadecimal characters.
- The key group bound to an SDP or service must be unbound from that SDP or service before the active outgoing SA for the key group can be removed.

- The active outgoing SA must be removed (deconfigured) before the SPI can be deleted from the SA list in the key group.
- The encryption or authentication algorithm for a key group cannot be changed if there are any SAs in the key group.
- The encryption configured on an SDP used to terminate the Layer 3 spoke SDP of a VPRN (enabled or disabled) always overrides any VPRN-level configuration for encryption. See section “VPRN Layer 3 Spoke-SDP Encryption and MP-BGP-based VPRN Encryption Interaction” in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for more information.
- The NSP NFM-P provides configuration parameters that are not configurable using the CLI. See Network Services Platform Management for more information.

To enable NGE for an SDP or VPRN service:

Step 1. Install the outbound direction key group on each node for the service.

Step 2. Install the inbound direction key group on each node for the service.

To enable NGE for a router interface:

Step 1. Enable **group-encryption** on the interface.

Step 2. Configure the outbound key group.

Step 3. Configure the inbound key group.

To change NGE from one key group to another key group for an SDP or VPRN service:

Step 1. Remove the inbound direction key group from each node for the service.

Step 2. Change the outbound direction key group on each node for the service.

Step 3. Install the new inbound direction key group on each node for the service.

To change NGE from one key group to another key group for a router interface:

Step 1. Remove the inbound key group.

Step 2. Configure the new outbound key group.

Step 3. Configure the new inbound key group.

To disable NGE for an SDP or VPRN service:

Step 1. Remove the inbound direction key group from each node providing the service.

Step 2. Remove the outbound direction key group from each node for the service.

To disable NGE for a router interface:

Step 1. Remove the inbound key group.

Step 2. Remove the outbound key group.

Step 3. Disable group encryption on the interface.

2.11 Service Creation Process Overview

Figure 35 displays the overall process to provision core and subscriber services.

Figure 35 **Service Creation and Implementation Flow**

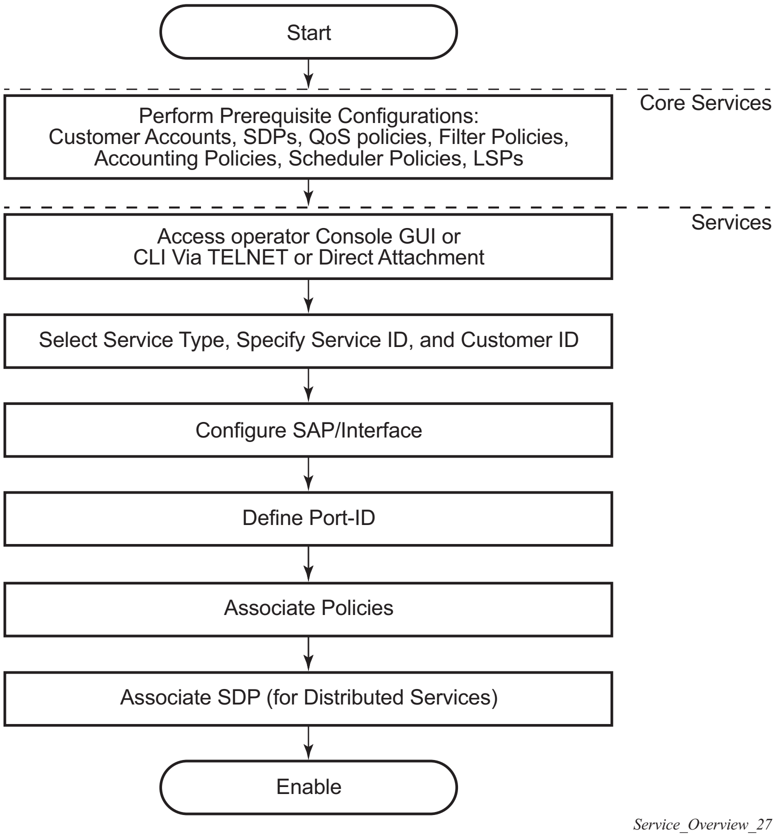
Start
Core Services
Perform Prerequisite Configurations:
Customer Accounts, SDPs, QoS policies, Filter Policies,
Accounting Policies, Scheduler Policies, LSPs
Services
Access operator Console GUI or
CLI Via TELNET or Direct Attachment
Select Service Type, Specify Service ID, and Customer ID
Configure SAP/Interface
Define Port-ID
Associate Policies
Associate SDP (for Distributed Services)
Enable
Service_Overview_27

2.12 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.

2.12.1 Phase 1: Core Network Construction

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

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

2.12.2 Phase 2: Service Administration

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

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

2.12.3 Phase 3: Service Provisioning

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

2.13 Configuration Notes

This section describes service configuration caveats.

2.13.1 General

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

Core tasks include the following:

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

Subscriber services tasks include the following:

- Create Apipe, Cpipe, Epipe, Fpipe, IES, Ipipe, VPLS or VPRN services on the 7750 SR.
- Create Epipe, IES, Ipipe, VPLS or VPRN services on the 7450 ESS or 7950 XRS.
- Bind SDPs
- Configure interfaces (where required) and SAPs
- Create exclusive QoS and filter policies

2.14 Configuring Global Service Entities with CLI

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

2.14.1 Service Model Entities

The Nokia service model uses logical entities to construct a service. The service model contains four main entities to configure a service.

- Subscribers
- SDPs
- Services:
  - ATM VLL (Apipe) services — Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information
  - Circuit Emulation services (Cpipe) — Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information
  - Ethernet Pipe (Epipe) services — Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information
  - Frame Relay VLL (Fpipe) services — Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information
  - IP Interworking VLL (Ipipe) services — Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information
  - Virtual Private LAN Service (VPLS) — Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information
  - Internet Enhanced Service (IES) — Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for more information
  - Virtual Private Routed Network (VPRN) service — Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for more information
- Service Access Points (SAPs)

- Ethernet Pipe (Epipe) Services — Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information
- ATM VLL (Apipe) SAP — Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information
- Frame Relay VLL (Fpipe) SAP — Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information
- VPLS SAP — Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information
- IES SAP — Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for more information
- VPRN Interface SAP — Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for more information

2.14.2 Basic Configurations

The most basic service configuration must have the following:

- A customer ID
- A service type
- A service ID — An optional service name can also be configured in addition to the service ID. Service names are optional. 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.
- A SAP identifying a port and encapsulation value
- An interface (where required) identifying an IP address, IP subnet, and broadcast address
- For distributed services: an associated SDP

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

```
A:ALA-B>config>service# info detail
#------------------------------------------
...
```

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

2.14.3 Common Configuration Tasks

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

2.14.3.1 Configuring Customers

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

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

2.14.3.1.1 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
multi-service-site customer-site-name [create]
     assignment {port port-id | card slot}
     description description-string
     egress
     egress
          agg-rate
               burst-limit size [bytes | kilobytes]
               limit-unused-bandwidth
               queue-frame-based-accounting
               rate kilobits-per-second
          policer-control-policy name
               scheduler-override
               scheduler scheduler-name [create]
                    parent {[weight weight] [cir-weight
                         cir-weight]}
                    rate pir-rate [cir cir-rate]
               scheduler-policy scheduler-policy-name
     ingress
          scheduler-override
               scheduler scheduler-name [create]
                    parent {[weight weight] [cir-weight
                         cir-weight]}
                    rate pir-rate [cir cir-rate]
          scheduler-policy scheduler-policy-name
phone phone-number
```

The following displays a basic customer account configuration.

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

2.14.3.1.2 Configuring Multi-Service-Sites

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

After a multi-service site is created, it must be assigned to a chassis slot or port.

Note: The 7450 ESS-1 model multi-service site assignment configuration defaults to slot 1.

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

CLI Syntax:

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

The following displays a customer’s multi-service-site configuration.

```
A:ALA-12>config>service# info
-------------------------------------------
..
       customer 5 create
           multi-service-site "EastCoast" create
               assignment card 4
               ingress
```

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

The following shows an example of a customer’s 7450 ESS multi-service-site configuration.

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

2.14.3.2 Configuring an SDP

The most basic SDP must have the following:

- a locally unique SDP identification (ID) number
- the system IP address of the originating and far-end routers
- an SDP encapsulation type, either GRE or MPLS

2.14.3.2.1 SDP Configuration Tasks

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

Consider the following SDP characteristics:

- SDPs can be created as either GRE or MPLS.
- Each distributed service must have an SDP defined for every remote 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 an Nokia router 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 Nokia nodes.

Note: 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:

Step 1. Specify an originating node.

Step 2. Create an SDP ID.

Step 3. Specify an encapsulation type.

Step 4. Specify a far-end node.

2.14.3.2.2 Configuring an SDP

Use the following CLI syntax to create an SDP and select an encapsulation type. If **gre** or **mpls** is not specified, the default encapsulation type is **gre**.

Note: When specifying the far-end IP address, a tunnel is created. A the path from Point A to Point B is created. When configuring a distributed service, an SDP ID must be specified. Use the **show service sdp** command to display the qualifying SDPs.

When specifying MPLS SDP parameters, specify an LSP or enable LDP. There cannot be two methods of transport in a single SDP except if the **mixed-lsp** command is specified. 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. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR MPLS Guide MPLS Configuration Guide* for configuration and command information.

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

CLI Syntax:

```
config>service>sdp sdp-id [gre | 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
               ldp (only for MPLS SDPs)
               lsp lsp-name [lsp-name](only for MPLS
                 SDPs)
path-mtu octets
signaling {off | tldp}
no shutdown
```

The following displays an example of a GRE SDP, an LSP-signaled MPLS SDP, and an LDP-signaled MPLS SDP configuration.

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

2.14.3.2.3 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 must revert the SDP to a higher priority LSP type when one becomes available by using the following command:

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

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>sdp-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.

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 line card that will activate it to forward service packets according to the following priority order:

Step 1. RSVP LSP type. Up to 16 RSVP LSPs can be entered by the user and programmed by the service manager in ingress line card to load balance service packets. This is the highest priority LSP type.

Step 2. LDP LSP type. One LDP FEC programmed by service manager but ingress line card can use up to 16 LDP ECMP paths for the FEC to load balance service packets when ECMP is enabled on the node.

Step 3. BGP LSP type. One RFC 3107-labeled BGP prefix programmed by the service manager. The ingress line card can use more than one next-hop for the prefix.

In the case of the RSVP or 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: 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 **config>router>ldp>tunnel-down-damp-time** command.

If the value of the **sdp-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.

If class based forwarding is enabled for this SDP, the forwarding of the packets over the RSVP LSPs will be based on the FC of the packet as in current implementation. When the SDP activates the LDP LSP type, then packets are forwarded over the LDP ECMP paths using the regular hash routine.

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.

2.15 Ethernet Connectivity Fault Management (ETH-CFM)

The IEEE and the ITU-T have cooperated to define the protocols, procedures and managed objects to support service based fault management. Both IEEE 802.1ag standard and the ITU-T Y.1731 recommendation support a common set of tools that allow operators to deploy the necessary administrative constructs, management entities and functionality, Ethernet Connectivity Fault Management (ETH-CFM). The ITU-T has also implemented a set of advanced ETH-CFM and performance management functions and features that build on the proactive and on demand troubleshooting tools.

CFM uses Ethernet frames and is distinguishable by ether-type 0x8902. In certain cases the different functions will use a reserved multicast Layer 2 MAC address that could also be used to identify specific functions at the MAC layer. The multicast MAC addressing is not used for every function or in every case. The Operational Code (OpCode) in the common CFM header is used to identify the PDU type carried in the CFM packet. CFM frames are only processed by IEEE MAC bridges.

IEEE 802.1ag and ITU-T Y.1731 functions that are implemented are available on the SR and ESS platforms.

This section of the guide will provide configuration example for each of the functions. It will also provide the various OAM command line options and show commands to operate the network. The individual service guides will provide the complete CLI configuration and description of the commands in order to build the necessary constructs and management points.

Table 11 lists and expands the acronyms used in this section.

Table 11 **ETH-CFM Acronym Expansions**

Acronym	Expansion	Supported Platform
1DM	One way Delay Measurement (Y.1731)	All
AIS	Alarm Indication Signal	All
CCM	Continuity check message	All
CFM	Connectivity fault management	All
CSF	Client Signal Fail (Receive)	All
DMM	Delay Measurement Message (Y.1731)	All
DMR	Delay Measurement Reply (Y.1731)	All

Table 11 **ETH-CFM Acronym Expansions (Continued)**

Acronym	Expansion	Supported Platform
ED	Ethernet Defect (Y.1731 sub OpCode of MCC)	All
LBM	Loopback message	All
LBR	Loopback reply	All
LMM	(Frame) Loss Measurement Message	Platform specific
LMR	(Frame) Loss Measurement Response	Platform specific
LTM	Linktrace message	All
LTR	Linktrace reply	All
MCC	Maintenance Communication Channel (Y.1731)	All
ME	Maintenance entity	All
MA	Maintenance association	All
MD	Maintenance domain	All
MEP	Maintenance association end point	All
MEP-ID	Maintenance association end point identifier	All
MHF	MIP half function	All
MIP	Maintenance domain intermediate point	All
OpCode	Operational Code	All
RDI	Remote Defect Indication	All
TST	Ethernet Test (Y.1731)	All
SLM	Synthetic Loss Message	All
SLR	Synthetic Loss Reply (Y.1731)	All
VSM	Vendor Specific Message (Y.1731)	All
VSR	Vendor Specific Reply (Y.1731)	All

2.15.1 Facility MEPs

Facility MEPs have been introduced to improve scalability, reduce operational overhead, and provide fate sharing without requiring service MEPs. This allows for fault notification for Epipe services that share a common transport. Facility MEPs recognize failure based solely on ETH-CFM detection mechanisms.

There are a total of four facility MEPs, as described below:

- Port (physical) — Detects port failure where LoS may be hidden by some intervening network
- LAG (logical) — Validates the connectivity of the LAG entity
- Tunnel (logical) — Enables fate sharing of a MEP configured on a QinQ encapsulated access LAG and outer VLAN-ID.
- Router IP Interface (logical) — Validates the Layer 2 connectivity between IP endpoints (troubleshooting only – no CCM functions)

In general, a Facility MEP detects failure conditions using ETH-CFM at the Ethernet Transport layer. The detection is based solely on the MEP entering a fault state as a result of ETH-CC. Conditions outside the scope of ETH-CFM do not directly influence the state of the MEP. However, these outside influences have indirect influence. For example, upon a failure of a port, CCM messages cannot reach the destination. This condition causes the MEP to enter a fault state after the 3.5*interval expires, with the only exception being the acceptance of AIS on a Tunnel MEP. AIS received on all other facilities MEPs are discarded silently when normal level matching targets the local facility MEP.

Facility MEPs are supported as part of a down MEP only. Facility MEPs validate the point to point Ethernet transport between two end points. Facility MEPs do not validate switching functions that are not part of the point to point Ethernet transport. Instead, service MEPs validate switching functions that are not part of the point to point Ethernet transport.

A facility MEP allows for the scaling improvements using fate sharing and leveraging OAM mapping. The OAM mapping functions are part of the fault propagation functions and allow ETH-CFM to move from alarms only to network actions. Service based MEPs are not required to generate AIS in reaction to a facility MEP fault. OAM mapping and generation of fault via fault-propagation means or the AIS function are only available for Epipe services. There is no equivalent AIS generation as part of the facility fault for VPLS, IES, and VPRN. There is no service MEP required to have the SAP transition in the VPLS, IES, and VPRN service context. Normal SAP transition functions does not occur when these services are configured to accept the tunnel fault, or in reaction to a facility fault, where the underlying port or LAG transitions the SAP.

Note: Do not exceed the platform-specific scaling limits. A single facility fault may trigger the generation of many service level faults, ensure that the specific ETH-CFM processing power of the network element and any configured rate controlling features for the service are not exceeded. Exceeding the network element scaling properties may lead to OAM packet loss during processing and result in undesirable behavior.

The implementation of facility MEPs must adhere to all platform-specific specifications. For example, sub-second enabled CCM MEPs are supported on port based MEPs. However, any platform restrictions preventing the sub-second enabled MEPs override this capability and require the operator to configure CCM intervals that are supported for that specific platform.

Facility MEPs are created in the same manner as service MEPs, both related to the ETH-CFM domain and association. However, the association used to build the facility MEP does not include a bridge-identifier. The CLI ensures that a bridge id is not configured when the association is applied to a facility MEP.

Service MEPs and Facility MEPs may communicate with each other, as long as all the matching criteria are met. Since facility MEPs use the standard ETH-CFM packets, there is nothing contained in the packet that would identify an ETH-CFM packet as a facility MEP or Service MEP.

Facility MEPs are not supported on ports that are configured with Eth-Tunnels (G.8031) and only facility MEPs of 1 second and above are supported on the ports that are involved in an Eth-Ring (G.8032).

2.15.1.1 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. By default, an fault in the CCM MEP state machine generates AIS when it is configured. Table 12 illustrates the ETH-CC defect condition groups, configured low-priority-defect setting, priority and defect as it applies to fault propagation. AIS maintains its own low-priority-defect option which can be used to exclude the CCM defect RDI from triggering the generation of AIS.

Table 12 **Defect Conditions and Priority Settings**

Defect	Low Priority Defect	Description	Causes	Priority
DefNone	n/a	No faults in the association.	Normal operations.	n/a

Table 12 **Defect Conditions and Priority Settings (Continued)**

Defect	Low Priority Defect	Description	Causes	Priority
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 MEP-ID.	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

A facility MEP may trigger two distinct actions as a result of fault. Epipe services generate AIS that have been configured to do so as a result of a failure. The level of the AIS is derived from the facility MEP. Multiple **client-meg-levels** can be configured under the facility MEP to allow for operational efficiency in the event a change is required. However, only the lowest AIS level is generated for all the linked and applicable services. VPLS, IES and VPRN SAPs transition the SAP state that are configured to react to the facility MEP state. In addition, Epipe services may also take advantages of OAM and mapping functions.

Before implementing facility MEPs, it is important to understand the behavior of AIS and Fault propagation. Nokia advises considers the following recommendations listed below before enabling or altering the configuration of any facility MEP. These steps must be tested on each individual network prior to building a maintenance operational procedure (MOP).

- Do not configure AIS on the facility MEP until the ETH-CCM has been verified. For instance, when a local MEP is configured with AIS prior to the completion of the remote MEP, the AIS is immediately generated when the MEP enters a fault state for all services linked to that facility MEP.
- Disable the **client-meg-level** configuration parameter when changes are being made to existing functional facility MEPs for AIS. Doing this stops the transmit function but maintains the ability to receive and understand AIS conditions from the network.
- Set the **low-priority-defect** parameter to **noXconn** in order to prevent the MEP from entering a defect state, triggering SAP transitions and OAM mapping reactions.

It is important to consider and select what types of fault conditions causes the MEP to enter a faulty state when using fault propagation functions.

The **ccm-hold-timers** supported on port-based MEPs configured with a sub-second interval. The **ccm-hold-timers** prevents the MEP from entering a failed state for 3.5 times the CCM interval plus the additional hold timer.

2.15.1.2 General Detection, Processing and Reaction

All Facility MEPs that support CCM functions must only have one remote MEP peer. Facilities MEPs validate point-to-point logical or physical Ethernet transports. Configure service MEPs if multipoint-service validation is required.

There are three distinct functions for a Facility MEP:

- General Detection: Determines that a fault has occurred. In this case, the MEP performs its normal functions such as: recognizing the fault condition, maintaining the local errors and reporting based on low-priority-setting, and taking no further action. This is the default.

- Fault Processing: By default, there is no action taken as a result of a MEP state machine transition beyond alarming. In order to take action which may include a SAP operational state change, generation of AIS, or fault propagation and mapping, the appropriate facility fault configuration parameter must be configured and enabled. The general reaction to a fault is described below. More details are including the section describing the functions of the individual facility MEPs.
  - Port—Affects link operational status of the port. Facility failure changes the operational state to Link Up. This indicates that the port has been brought down as a result of OAM MEP Fault. This operational state has the equivalent function to port down condition.
  - LAG—Affects link operational status of the LAG. Facility failure changes the operational state of the LAG to DOWN. This indicates that the LAG has be brought down as a result of OAM MEP Fault.
  - Tunnel MEP—Enters faulty state and will further impact the operational state of the SAPs linked to the tunnel MEP state.
    - Epipe SAP remains operationally up, SAP's flag set to **OamTunnelMEPFault**
    - Ipipe SAP remains operationally up, SAP's flag set to **OamTunnelMEPFault**
    - VPLS, IES and VPLS SAPs transition to operationally **down**, the SAP's flag is set to **OamTunnelMEPFault.** SAP operational states and flags are affect only by the **tunnel-fault** configuration option.
  - Router IP Interface— Affects operational status of the IP Interface.
- Propagation: Services appropriately linked to the Facility MEP take the following service-specific actions:
  - Epipe generates AIS or use Fault Propagation and OAM mappings.
  - VPLS does not propagate fault using AIS unless service-based MEPs are configured and contain MEP-specific AIS configuration. SAP transitions will occur when the facility MEP failure is recognized by the service.
  - IES and VPRN, as Layer 3 functions, act as boundaries for Layer 2 fault processing. No propagation functions occur beyond what is currently available as part of fault propagation: SAP down.
- AIS-enable configuration options: Epipe services support the **ais-enable** configuration option under the SAP hierarchy level. This structure, outside of the MEP context, creates a special link between the Epipe service SAP and the facility MEP. If a facility MEP enters a fault state, all Epipe service SAPs with this configuration generate lowest-level AIS at the level configured under the facility MEP. As with fault propagation, AIS generation is restricted to Epipe services only. The actions taken by the other services are described in more detail in the relevant facility MEP sections.

Note: Facility MEPs do not support the generation of AIS to an explicitly configured endpoint. An explicitly configured endpoint abstracts multiple endpoints within its context; for example, pseudowire (PW) redundancy. Although the linkage of a facility MEP to an Epipe, and AIS generation triggered as a result of the facility MEP failure can be configured, AIS generation is not supported and will be unpredictable. When an explicit endpoint is configured, service-based MEPs are required when AIS generation is the desired behavior.

2.15.1.3 Port-Based MEP

There is an increase in services that share the same facilities, and that service-based ETH-CFM, although very granular, comes at an operational and scalability cost. Configuring a MEP on a physical port allows ETH-CFM to detect Ethernet transport failures, raise a facility alarm, and perform local fault processing. A facility event is coordinated to the services or functions using the affected port.

The port-based MEP is intended to validate physical connectivity to the peer MEP, and provide on-demand and scheduled troubleshooting, and performance management functions.

Port facility MEPs are advantageous in cases where port-to-port connectivity issues are obscured, similar to the deployment use cases for *IEEE 802.3 Clause 57 – Operation, Administration and Maintenance* (formerly 802.3ah). *Clause 57* specification limits the transmit rate to 10 packets/s, or a send duration of 100 ms. To more quickly detect port failure conditions between two peers, a port-based facility MEP may be configured to use the supported sub-second CCM intervals. One-second and above timers are also available for configuration in cases where aggressive timers are not necessary. All platform-specific requirements must be met for the desired interval. ETH-CFM and IEEE 802.3ah Clause 57 can influence the operational state of the port over which they are configured.

The 802.3ah and ETH-CFM protocols cannot simultaneously control the individual port operational state. Both protocols can be decoupled from the port operational state. The 802.3ah protocol defaults to influencing the port operational state. This can be modified by using the **config**>**port**>**ethernet**>**efm**>**ignore-efm-state** command. The ETH-CFM protocol ETH-CC defaults to alarm-only without influencing the port operational state. This can be modified by using the **config**>**port**>**ethernet**>**eth-cfm**>**mep**>**facility-fault**. The 802.3ah and ETH-CFM protocol combinations that conflict with the single-port operational control rule are rejected with a configuration error.

Port-level ETH-CFM PDUs are sent untagged because they are not specific to any service or VLAN. The ETH-CFM packets generated from a port-based facility MEP must use an ETH-CFM level of 0 or 1. Any ETH-CFM PDU that arrives untagged on a port matching the level for the port-based facility MEP will be terminated and processed by the port-based MEP.

Do not use MEPs configured with level 0 to validate logical transport or services. Consideration should be given to blocking all non-customer (5-7) levels at the entry point of the network.

It is not expected that faults from other parts of the network will be propagated and terminated on a port-based facility MEP. This type of facility MEP provides a one-to-one validation with a single remote MEP across on a physical port, allowing locally detected faults to be propagated to the endpoints of the network.

A physical port may only have a single port-based facility MEP. Since the purpose of the MEP is to control the port state, more than one is not required per port.

When a port enters the link up operational state due to ETH-CFM, the MEP continues to transmit and received in order to properly clear the condition. However, when the port fails for reasons that are not specific to ETH-CFM, it stops transmit and receive functions until the condition is cleared. This is different than the behavior of a service MEP, because facility MEPs only supports Down MEPs, while some service-based MEPs support UP and Down MEPs. In the case of UP MEPs, a single port failure may not prevent all the CCMs from egressing the node. So the operational method for service-based MEPs remains the same: continuing to increase the counter for CCM transmit in the event of port failure, regardless of the reason. The transmit ETH-CCM counters do not apply to sub-second CCM-enabled MEPs.

There are two types of port in the context of port-based facility MEPs. The first type are ports that are not part of a LAG, referred to as non-member ports. The second type of ports are ports that are part of a LAG, referred to member ports, and have slightly different reactions to fault. MEPs configured directly on either type of port will act the same. However, a MEP configured on a non-member port and a MEP configured on a member port handle fault propagation differently.

When a port-based facility MEP causes the port to enter the operational state Link Up, normal processing occurs for all higher level functions. If the port is a member port, unless the entire LAG enters a non-operational state, the SAP configured on the LAG remains operational. A facility MEP on a member port has no direct influence on the SAP. The purpose of a facility MEP on a member port is to provide feedback to the LAG. The LAG performs the normal computations in response to a port down condition. A facility MEP configured on a non-member port does have direct control over the SAPs configured on the port. Therefore, when a port fails, all the SAPs transitions to the operation state down. When this occurs, fault may be propagated using AIS for those Epipe services that are AIS-enabled under the SAP. For the services that have MEPs configured on the SAP or the binding, fault propagation occurs. For VPLS, IES and VPRN services, normal reaction to a SAP entering a down state occurs.

When a LAG is administratively shutdown, the member ports are shutdown automatically. As a result, packet reception is interrupted, causing ETH-CFM functions running on physical member ports to lose connectivity. Therefore, the CFM functions on member ports are somewhat tied to the LAG admin status in this case.

Note: LAG convergence time is not affected by a facility MEP on a member port once the port has entered the link up operational state. The ETH-CFM failure of a port-based MEP acts as the trigger to transition the port.

Figure 36 provides an example of how an ETH-CFM failure reacts with the various services that share that port. The green Epipe service generates AIS as a result of the port failure using the **client-meg-level** command configured on the port facility MEP. The multipoint service takes location configured action when the SAP transitions to the down operational state. The blue Epipe service is not affected by the port link up state as a result of ETH-CFM fault.

Figure 36 **Fault Handling Non-Member Port**

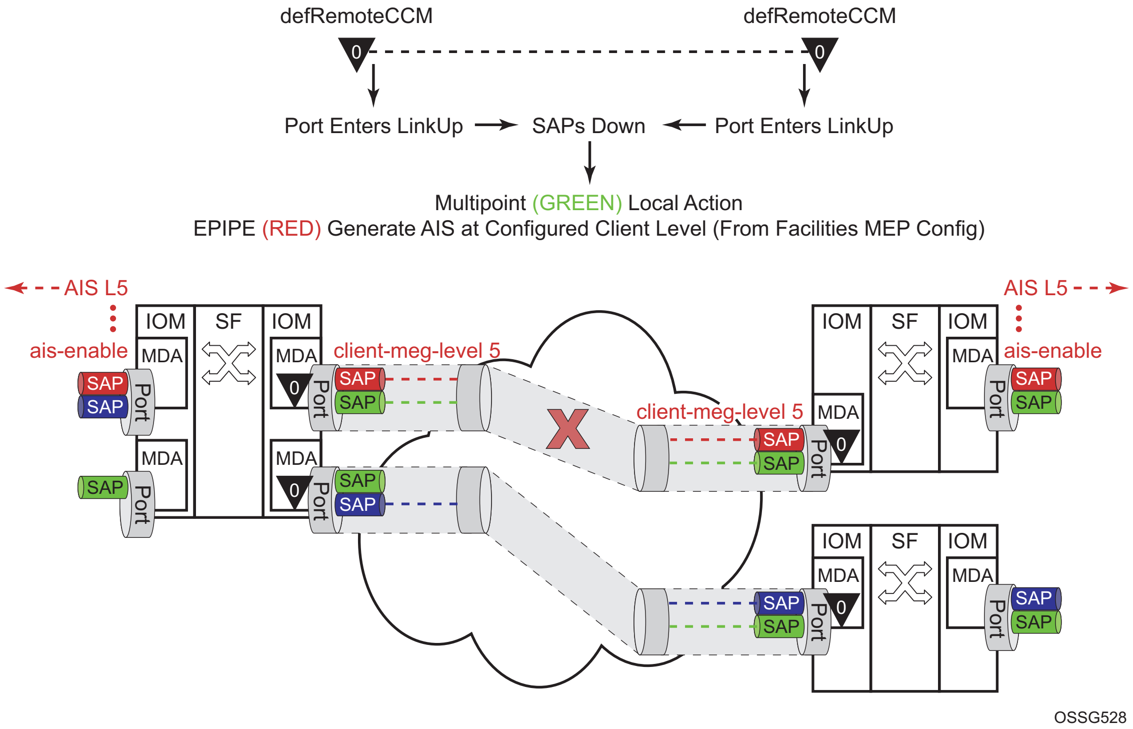

A debounce function has been implemented to prevent notifying every port state change if a port bounces multiple times within a window. Up to four notifications will be accepted in a three second window. If the third port state is a down state change the fourth will be ignored. If the fourth port state change is a down state change it will be processed. After that no further state changes will be accepted for the duration of the three second timer. This helps ensure that the port is not artificially held in the UP state when it is not operation. Following the processing of that last port state change, the third or fourth, the latest state change will be held and processed at the expiration of the three second hold timer.

Port based facility MEPs are not allowed on a port that is configured with G.8031 Ethernet Tunnels.

Example: Port-Based MEP Configuration

Figure 37 displays an example of how port-based MEPs and defect conditions translate into service awareness without service-based MEPs. From the two nodes perspective, they are aware they are directly connected at the port. The two nodes are unaware of any of the cross connections that allow this to occur.

Figure 37 **Port-Based MEP Example**

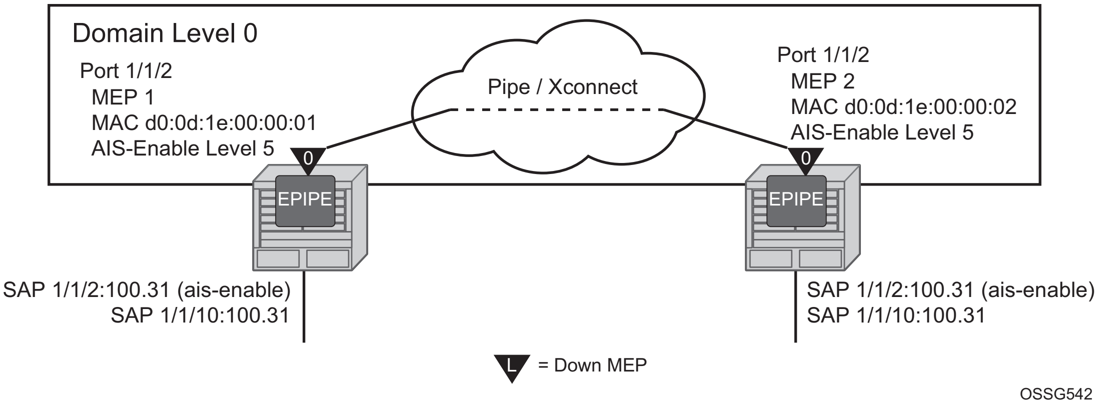

Configure port-based MEPs with the **facility-fault** option and **ais-enable client-meg-level** command. When the MEP enters any defect state, an AIS is generated to any Epipe service that has the ais-enable configured under the **sap>eth-cfm** hierarchy.

NODE1

```
config>eth-cfm# info
----------------------------------------------
        domain 10 format none level 0
            association 1 format icc-based name "FacilityPort0"
                ccm-interval 1
                remote-mepid 2
            exit
        exit
----------------------------------------------

config>port# info
----------------------------------------------
        ethernet
            mode access
            encap-type qinq
            eth-cfm
                mep 1 domain 10 association 1
                    ais-enable
                        client-meg-level 5
                    exit
                    facility-fault
     ccm-enable
                    mac-address d0:0d:1e:00:00:01
                    no shutdown
                exit
            exit
        exit
        no shutdown
----------------------------------------------

config>service>epipe# info
----------------------------------------------
            sap 1/1/2:100.31 create
                eth-cfm
                    ais-enable
                exit
            exit
            sap 1/1/10:100.31 create
            exit
            no shutdown
----------------------------------------------
```

NODE2

```
config>eth-cfm# info
----------------------------------------------
        domain 10 format none level 0
            association 1 format icc-based name "FacilityPort0"
                ccm-interval 1
                remote-mepid 1
            exit
        exit
----------------------------------------------

config>port# info
----------------------------------------------
        ethernet
            mode access
            encap-type qinq
            eth-cfm
                mep 2 domain 10 association 1
                    ais-enable
                        client-meg-level 5
                    exit
                    facility-fault
                    ccm-enable
                    mac-address d0:0d:1e:00:00:02
                    no shutdown
                exit
            exit
        exit
        no shutdown
----------------------------------------------

config>service>epipe# info
----------------------------------------------
            sap 1/1/2:100.31 create
                eth-cfm
                    ais-enable
                exit
            exit
            sap 1/1/10:100.31 create
            exit
            no shutdown
----------------------------------------------
```

There are two different levels of fault to consider: Port State/Operational State driven by the low-priority-defect setting and the generation of AIS driven by the defect state for the MEP.

If the low-priority-defect is left at the default macRemErrXcon setting, then port state may not match on both nodes. If an unidirectional failure is introduced for port-based MEPs, then RDI is received on one of the nodes and the other node would report and react to RemoteCCM (timeout). The RDI defect is below the default low-priority-defect in priority, and the port would remain operationally UP and the port state would remain UP. The MEP that has timed out the peer MEP takes port level action because this defect is higher in priority than the default low-priority-defect. The port state is recorded as Link Up and the Port is operationally down with a Reason Down: ethCfmFault. To avoid this inconsistency, set the **low-priority-defect** setting to detection unidirectional failures using the allDef option.

The following **show** commands reveal the condition mentioned above within the network. Node 1 is receiving RDI and Node 2 has timed out its peer MEP.

NODE1

```
#show port
===============================================================================
Ports on Slot 1
===============================================================================
Port        Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id          State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
…snip..
1/1/2       Up    Yes  Up      1522 1522    - accs qinq xcme
…snip..
```

```
#show port 1/1/2
===============================================================================
Ethernet Interface
===============================================================================
Description        : 10/100/Gig Ethernet SFP
Interface          : 1/1/2                      Oper Speed       : 1 Gbps
Link-level         : Ethernet                   Config Speed     : 1 Gbps
Admin State        : up                         Oper Duplex      : full
Oper State         : up                         Config Duplex    : full
Physical Link      : Yes                        MTU              : 1522
…snip…

#show eth-cfm mep 1 domain 10 association 1
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index           : 10                       Direction         : Down
Ma-index           : 1                        Admin             : Enabled
MepId              : 1                        CCM-Enable        : Disabled
Port               : 1/1/2                    VLAN              : 0
Description        : (Not Specified)
FngState           : fngReset                 ControlMep        : False
LowestDefectPri    : macRemErrXcon            HighestDefect     : none
Defect Flags       : bDefRDICCM
Mac Address        : d0:0d:1e:00:00:01        ControlMep        : False
CcmLtmPriority     : 7
CcmTx              : 1481                     CcmSequenceErr    : 0
Fault Propagation  : disabled                 FacilityFault     : Notify
MA-CcmInterval     : 1                        MA-CcmHoldTime    : 0ms
Eth-1Dm Threshold  : 3(sec)                   MD-Level          : 0
Eth-Ais:           : Enabled                  Eth-Ais Rx Ais:   : No
Eth-Ais Tx Priorit*: 7                        Eth-Ais Rx Interv*: 1
Eth-Ais Tx Interva*: 1                        Eth-Ais Tx Counte*: 3019
Eth-Ais Tx Levels  : 5
Eth-Tst:           : Disabled
…snip…

# show service sap-using eth-cfm facility
===============================================================================
Service ETH-CFM Facility Information
===============================================================================
SapId             SvcId                            SAP AIS  SAP Tunnel SVC Tunnel
                                                            Fault      Fault
-------------------------------------------------------------------------------
1/1/2:100.31      100                              Enabled  Accept     Ignore
-------------------------------------------------------------------------------
No. of Facility SAPs: 1
===============================================================================

NODE2
# show port
===============================================================================
Ports on Slot 1
===============================================================================
Port        Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id          State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
```

```
…snip..
1/1/2       Up    Yes  Link Up 1522 1522    - accs qinq xcme
…snip..

# show port 1/1/2
===============================================================================
Ethernet Interface
===============================================================================
Description        : 10/100/Gig Ethernet SFP
Interface          : 1/1/2                      Oper Speed       : N/A
Link-level         : Ethernet                   Config Speed     : 1 Gbps
Admin State        : up                         Oper Duplex      : N/A
Oper State         : down                       Config Duplex    : full
Reason Down        : ethCfmFault
Physical Link      : Yes                        MTU              : 1522
…snip…

# show eth-cfm mep 2 domain 10 association 1
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index           : 10                       Direction         : Down
Ma-index           : 1                        Admin             : Enabled
MepId              : 2                        CCM-Enable        : Enabled
Port               : 1/1/2                    VLAN              : 0
Description        : (Not Specified)
FngState           : fngDefectReported        ControlMep        : False
LowestDefectPri    : macRemErrXcon            HighestDefect     : defRemoteCCM
Defect Flags       : bDefRemoteCCM
Mac Address        : d0:0d:1e:00:00:02        ControlMep        : False
CcmLtmPriority     : 7
CcmTx              : 5336                     CcmSequenceErr    : 0
Fault Propagation  : disabled                 FacilityFault     : Notify
MA-CcmInterval     : 1                        MA-CcmHoldTime    : 0ms
Eth-1Dm Threshold  : 3(sec)                   MD-Level          : 0
Eth-Ais:           : Enabled                  Eth-Ais Rx Ais:   : No
Eth-Ais Tx Priorit*: 7                        Eth-Ais Rx Interv*: 1
Eth-Ais Tx Interva*: 1                        Eth-Ais Tx Counte*: 3515
Eth-Ais Tx Levels  : 5
Eth-Tst:           : Disabled
…snip…

# show service sap-using eth-cfm facility
===============================================================================
Service ETH-CFM Facility Information
===============================================================================
SapId             SvcId                          SAP AIS  SAP Tunnel SVC Tunnel
                                                          Fault      Fault
-------------------------------------------------------------------------------
1/1/2:100.31      100                            Enabled  Accept     Ignore
-------------------------------------------------------------------------------
No. of Facility SAPs: 1
===============================================================================
```

2.15.1.4 LAG Based MEP

LAG bundled ports provide both protection and scalability. Down MEPs configured on a LAG validates the connectivity of the LAG. Failure of this MEP causes the LAG to enter an operational down state. SAPs connected to the operationally down LAG transitions to operationally down. This triggers the configured reaction and processing similar to that of the port-based facility MEP. AIS is generated for those Epipe services with AIS enabled under the SAP. Local processing occurs for VPLS, IES and VPRN services that have experienced the SAP failure as a result of the LAG based SAP. Furthermore, fault propagation is invoked for any SAP with fault propagation operations enabled as a result of the failed LAG based SAP. LAG-based MEPs must be configured with a direction down.

LAG ETH-CFM PDUs are sent untagged because they are not specific to any service or VLAN. When running the combination of LAG-based MEPs and port-based MEPs, domain-level nesting rules must be adhered to for proper implementation, and is enforced by the CLI on the local node. As stated earlier, do not configure logical non-port-based MEPs, including service-based MEPs, to use level 0 for the ETH-CFM packets.

Since the recognition of fault is determined entirely by the ETH-CFM function, timeout conditions for the MEP occurs in 3.5 times the CCM interval. The LAG admin state or other failures that causes the LAG to completely fail, does not directly influence the MEP. The state of the MEP can only be influenced by the ETH-CFM function, specifically ETH-CC.

Since the LAG-based MEP selects a single member port to forward ETH-CFM packets, port-based facilities MEPs must be deployed to validate the individual member ports. Functional tests that require the ability to test individual member ports need to be performed from the port-based MEPs. The LAG-based MEPs validate only the LAG entity.

Figure 38, provides an example how an ETH-CFM failure reacts with the various services that share that LAG. There is only one way the LAG state can trigger the propagation of failure, and that is using ETH-AIS. The carrier must enable CCM at the LAG level and a ETH-CCM defect condition exists. The red Epipe service generates AIS as a result of the LAG failure using the **client-meg-level** parameter configured on the LAG facility MEP. The green multipoint service takes location-configured action when the SAP transitions to the down operational state.

Figure 38 **Fault Handling LAG MEP**

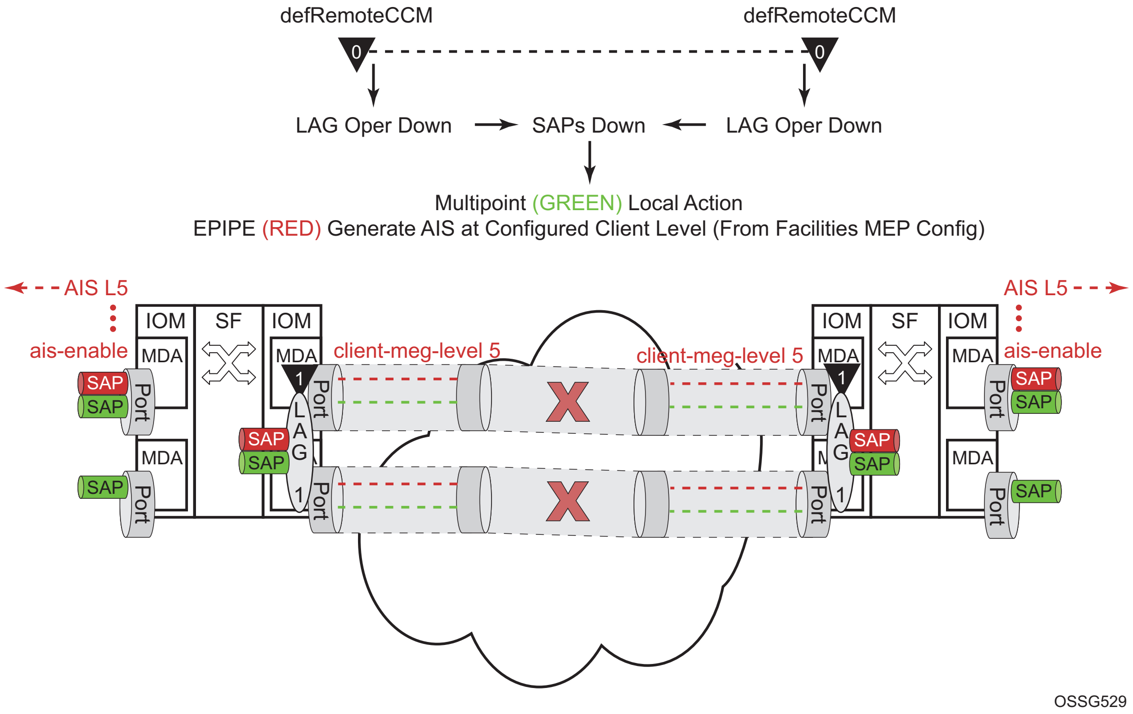

LAG-based MEP are supported for MultiChassis LAG (MC-LAG) configurations.

A LAG facility MEP must not be configured with **facility-fault** when it is applied to an MC-LAG. Traffic will black hole when the LAG Facility MEP enters a defect state. The LAG enters an operational down state but the MC-LAG does not switch over to the peer node. This restriction does not include Tunnel Facility MEPs which are applied to a LAG with an outer VLAN. Tunnel facility MEPs do not control the operational state of the LAG because they are outer VLAN specific.

Example: LAG MEP Configuration

Figure 39 uses a port-based MEP to validate port-to-port connectivity.

Figure 39 LAG MEP Example

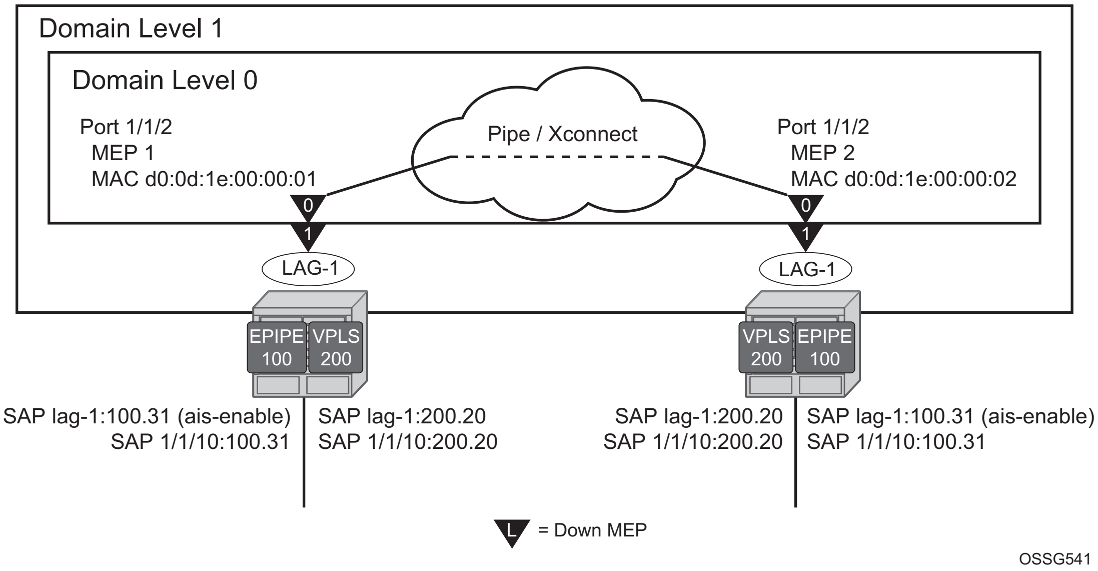

With the introduction of the LAG, the port no longer has direct control over the services SAPs. The ais-enable command has been disabled from the port for this reason. The low-priority-defect condition has been modified to react to all defect conditions “allDef”, avoiding the unidirectional issue demonstrated in the previous port-based MEP example. A LAG MEP is built on top the LAG with the **facility-fault** option and **ais-enable** command with the associated client-meg-level. This allows the Epipe services to generate AIS when the LAG MEP enters any defect condition. This example introduce the use of a VPLS service. VPLS, IES and VPRN services do not support the generation of AIS as a result of a facility MEP failure. However, all service SAPs which correspond to the failed facility will transition to a down state. Epipe service also generates AIS in this example.

NODE1

```
config>eth-cfm# info
----------------------------------------------
        domain 1 format none level 1
            association 1 format icc-based name "FacilityLag01"
                ccm-interval 1
                remote-mepid 22
            exit
        exit
        domain 10 format none level 0
            association 1 format icc-based name "FacilityPort0"
                ccm-interval 1
                remote-mepid 2
            exit
        exit
----------------------------------------------

config>port# info
```

```
----------------------------------------------
        ethernet
            mode access
            encap-type qinq
            eth-cfm
                mep 1 domain 10 association 1
                    facility-fault
                    ccm-enable
                    low-priority-defect allDef
                    mac-address d0:0d:1e:00:00:01
                    no shutdown
                exit
            exit
            autonegotiate limited
        exit
        no shutdown
----------------------------------------------

config>lag# info
----------------------------------------------
        mode access
        encap-type qinq
        eth-cfm
            mep 11 domain 1 association 1
                ais-enable
                    client-meg-level 5
                exit
 ccm-enable
                facility-fault
                low-priority-defect allDef
                no shutdown
            exit
        exit
        port 1/1/2
        no shutdown
----------------------------------------------

config>service# info
----------------------------------------------
        customer 1 create
            description "Default customer"
        exit
        epipe 100 customer 1 create
            sap 1/1/10:100.31 create
            exit
            sap lag-1:100.31 create
                eth-cfm
                    ais-enable
                exit
            exit
            no shutdown
        exit
        vpls 200 customer 1 create
            stp
                shutdown
            exit
            sap 1/1/10:200.20 create
            exit
            sap lag-1:200.20 create
            exit
            no shutdown
        exit
----------------------------------------------
```

NODE2

```
config>eth-cfm# info
----------------------------------------------
        domain 1 format none level 1
            association 1 format icc-based name "FacilityLag01"
                ccm-interval 1
                remote-mepid 11
            exit
        exit
        domain 10 format none level 0
            association 1 format icc-based name "FacilityPort0"
                ccm-interval 1
                remote-mepid 1
            exit
        exit
----------------------------------------------

config>port# info
----------------------------------------------
        ethernet
            mode access
            encap-type qinq
            eth-cfm
                mep 2 domain 10 association 1
                    facility-fault
                    ccm-enable
                    low-priority-defect allDef
                    mac-address d0:0d:1e:00:00:02
                    no shutdown
                exit
            exit
            autonegotiate limited
        exit
        no shutdown
----------------------------------------------

config>lag# info
----------------------------------------------
        mode access
        encap-type qinq
        eth-cfm
            mep 22 domain 1 association 1
                ais-enable
                    client-meg-level 5
                exit
                facility-fault
                ccm-enable
                low-priority-defect allDef
                no shutdown
            exit
        exit
        port 1/1/2
        no shutdown
----------------------------------------------

config>service# info
----------------------------------------------
        customer 1 create
            description "Default customer"
        exit
        epipe 100 customer 1 create
            sap 1/1/10:100.31 create
            exit
            sap lag-1:100.31 create
                eth-cfm
                    ais-enable
                exit
            exit
            no shutdown
        exit
        vpls 200 customer 1 create
            stp
                shutdown
            exit
            sap 1/1/10:200.20 create
            exit
            sap lag-1:200.20 create
            exit
            no shutdown
        exit
----------------------------------------------
```

A fault is introduced that only affects the LAG MEP. The port MEP continues to validate the port, meaning that the port remains operationally up and the lag transitions to operation down. The LAG transition causes all the SAPs tied to the LAG to transition to down. The VPLS service reacts normally with the configured behavior as a result of a SAP down condition. The Epipe SAP also transitions to down, causing the operational state of the Epipe service to transition to down. In this case, AIS is enabled under the SAP in the service those AIS packets will still be generated out the mate SAP.

Output from one of the nodes is included below. Since both react in the same manner, output from both nodes is not shown.

NODE1

```
#show port
===============================================================================
Ports on Slot 1
===============================================================================
Port        Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id          State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
…snip..
1/1/2       Up    Yes  Up      1522 1522    - accs qinq xcme
…snip..
```

```
show eth-cfm mep 11 domain 1 association 1
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index           : 1                        Direction        : Down
Ma-index           : 1                        Admin            : Enabled
MepId              : 11                       CCM-Enable       : Disabled
Port               : lag-1                    VLAN             : 0
Description        : (Not Specified)
FngState           : fngDefectReported        ControlMep       : False
LowestDefectPri    : allDef                   HighestDefect    : defRDICCM
Defect Flags       : bDefRDICCM
Mac Address        : 90:f3:ff:00:01:41        ControlMep       : False
CcmLtmPriority     : 7
CcmTx              : 4428                     CcmSequenceErr   : 0
Fault Propagation  : disabled                 FacilityFault    : Notify
MA-CcmInterval     : 1                        MA-CcmHoldTime   : 0ms
Eth-1Dm Threshold  : 3(sec)                   MD-Level         : 1
Eth-Ais:           : Enabled                  Eth-Ais Rx Ais:  : No
Eth-Ais Tx Priorit*: 7                        Eth-Ais Rx Interv*: 1
Eth-Ais Tx Interva*: 1                        Eth-Ais Tx Counte*: 1085
Eth-Ais Tx Levels  : 5
Eth-Tst:           : Disabled
…snip…

# show service sap-using eth-cfm facility
===============================================================================
Service ETH-CFM Facility Information
===============================================================================
SapId             SvcId                           SAP AIS  SAP Tunnel SVC Tunnel
                                                           Fault      Fault
-------------------------------------------------------------------------------
lag-1:100.31      100                             Enabled  Accept     Ignore
lag-1:200.20      200                             Disabled Accept     Ignore
-------------------------------------------------------------------------------
No. of Facility SAPs: 2
===============================================================================

# show eth-cfm cfm-stack-table facility
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM, A = AisRx
===============================================================================
CFM Facility Port Stack Table
===============================================================================
Port     Tunnel    Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
1/1/2     0          0 Down         10          1    1 d0:0d:1e:00:00:01 ------
===============================================================================
===============================================================================
CFM Facility LAG Stack Table
===============================================================================
Lag      Tunnel    Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
lag-1     0          1 Down          1          1   11 90:f3:ff:00:01:41 R-----
===============================================================================
```

2.15.1.5 Tunnel Based MEP

The concept of a logical tunnel carrying many unique and individual services has been deployed in many networks on QinQ encapsulated access ports where the outer VLAN represents the common transports and the inner VLAN represents the specific service. Typically, the tunnel transparently passes frames from multiple services through some common network. Tunnel MEPs are logically configured on the Port or LAG and outer VLAN for access ports use QinQ Ethernet encapsulation. Service processing is done after the tunnel MEP. This means that any service-based MEPs are required to be a higher level than that of the tunnel MEP. Tunnel MEPs are only supported on LAGs that are configured with QinQ encapsulation and must specify the outer VLAN.

The Tunnel MEP must validate connectivity between the tunnel end points. As with all facility MEPs, this is a point-to-point relationship between the local MEP and one remote MEP. By default, the MEP configured at the tunnel level performs only alarming functions. Actionable functions such as AIS, SAP transition, and fault propagation requires the operator to enable these functions.

The tunnel MEP must first be configured to take action when the MEP enters a fault state, similar to all other facilities MEPs. In order for the individual services to share the fate of the tunnel, each service must accept the facility MEP state. This is service-dependent and depends on the desired goals. Services share the tunnel fate based on the lag-id and the outer VLAN.

Epipe services support the **ais-enable** configuration option on the SAP. Enabling this option generates AIS in the event the tunnel MEP has entered a fault state as a result of ETH-CC failure, similar to other facility MEPs. However, since the individual SAPs configured within the different services are not directly affected by the tunnel MEP, an additional configuration is necessary to perform local SAP transitions, in the case of VPLS, IES and VPRN services and OAM mapping functions for Epipe services.

The **tunnel-fault** service-level command configured on an Epipe allows SAP flags to be set and fault propagation and OAM mapping functions between technology. The operational state of the SAP remains up. The operator needs to determine if the AIS generation of fault propagation is the best approach in their specific network. It is possible to configure both **ais-enable** and **tunnel-fault** accept within the Epipe service. However, this may generate multiple ETH-CFM packets, or multiple actions as a result of a single failure.

The **tunnel-fault accept** service level option is also available under Epipe, VPLS and IES services hierarchy level within the CLI. This allows for a tunnel fault to share fate with these service SAPs. For the non-Epipe services, the SAP enters an operationally **down** state, and normal processing occurs as a result of the SAP transition. In order to generate any ETH-CC based fault propagation, **suspend-cmm** or **use-int-stat-tlv**, this requires service-based MEPs that are actively running CCM with a peer.

The **tunnel-fault** configuration options occur in two levels of the CLI hierarchy: service level and SAP level. Both of the levels within a service and within the SAP (whose underlying port and outer tag has a tunnel MEP) must be set to accept, in order to have the function enabled. By default the **tunnel-fault** is set to ignore at the service level and accept at the SAP level. This means that a single **tunnel-fault** accept at the service level will enable fault operations for all SAPs in the service. The operator is free to enable and disable on specific SAPs by choosing the ignore option under the individual SAP. The combination of **accept** at the service level and ignore at the SAP level prevents that specific SAP from recognizing fault. AIS generation for Epipe services is not controlled by the **tunnel-fault** configuration options.

Specific to tunnel MEPs, reception of AIS on the tunnel MEP causes AIS to be cut through to all Epipe services that have the ais-enabled command configured under the SAP. During a fault condition, it is important that the AIS configuration under the tunnel MEP not be modified. This causes increased network element CPU processing requirements and in scaled environments transitioning this command during a heavily loaded fault condition, where highly scaled SAPs are linked to the fate of the tunnel MEP, may cause the system to spend more than normal processing time to be spent dealing with this artificially induced clear and fault situation. It is not expected that operators perform these types of tasks in production networks. Reception of AIS will not trigger a fault condition or AIS to be cut through when sub second CCM intervals have been configured on the Tunnel MEP.

Service-based MEPs may also be configured as normal for all services. They perform normal processing tasks, including service-based MEP with fault propagation.

As with all other facility MEPs, use only ETH-CFM functions to cause the Tunnel MEP to enter the fault state. Tunnel MEPs support sub second ccm-intervals on selected hardware. Tunnel MEPs must be configured with a direction of down. UP MEPs are not supported as part of the facility MEP concept.

LAG-based MEPs and LAG-based tunnel MEPs cannot be configured on the same LAG. Port-based MEPs may be configured on the LAG member ports of a tunnel MEP as long as they follow the requirements for port-based MEPs on LAG member ports. All those consideration are applicable here, including nesting and port-level control only without propagation.

Port-based MEPs and port-based tunnel MEPs cannot be configured on the same port.

LAG-based tunnel MEPs are supported in Multi-Chassis LAG (MC-LAG) configuration. However, sub second CCM enabled intervals should not be configured when the LAG-based tunnel MEP utilizes the transport of an MC-LAG. Only one second and above CCM intervals should be used. Not all platforms support sub second CCM enable tunnel MEPs.

Tunnel MEPs are meant to propagate fault from one segment to the other for Epipe services. Figure 40 shows how individual Epipes have SAPs connecting to a legacy network. A MEP is configured at the tunnel level and peers with a single remote peer MEP.

Figure 40 **Tunnel Concepts and Encapsulation**

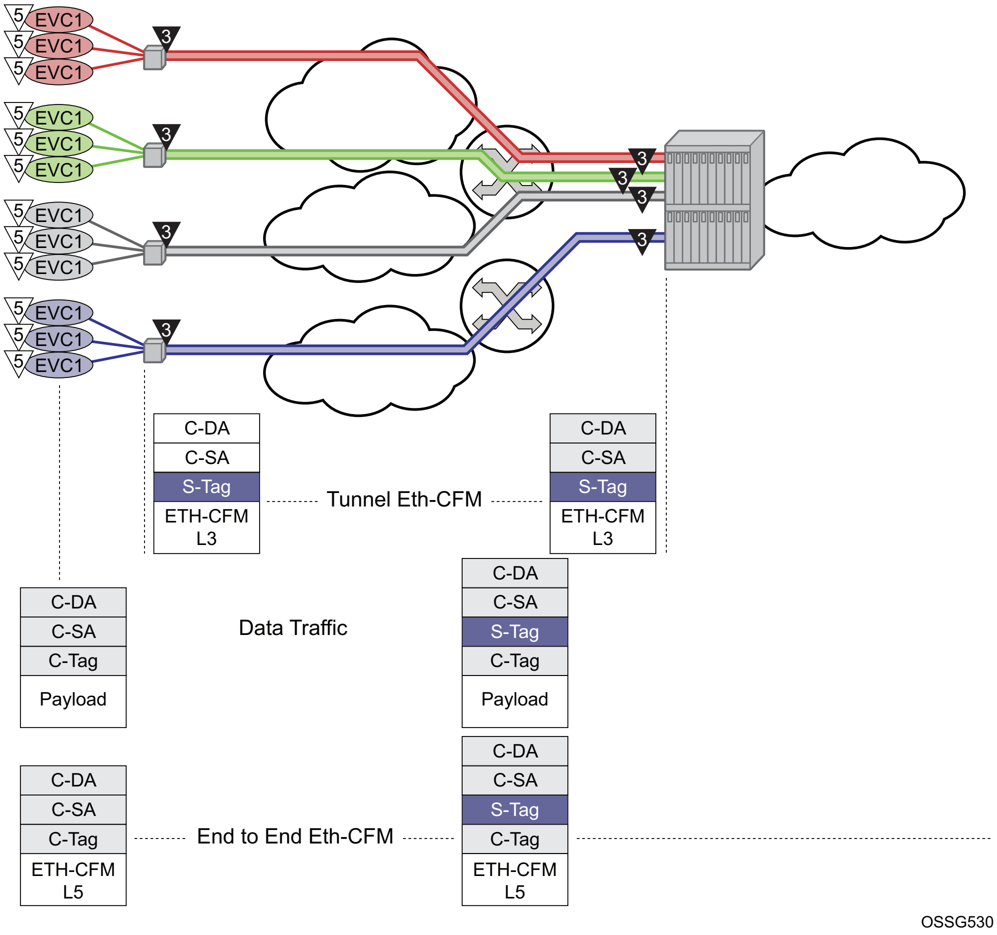

This is only one example of a tagged service. The principles of a tunnel MEP may be applied to other service as applicable. Remember that tunnel MEPs are only supported on LAGs that are configured with QinQ encapsulation and must have an outer VLAN.

Individual services can be monitored end-to-end by placing a MEP on the service endpoint at the CPE, denoted by the MEP at level 5 on the individual EVC (customer levels 5-7). The Network Interface Demarcation (NID) typically places a single tag, outer or only, on the customer traffic. This is cross connected to the proper connection in the access network and eventually arrive on the Ethernet Aggregation Switch. The connection between the legacy or access network and the aggregation switch must be either a LAG bundle or MC-LAG in order for tunnel MEPs to be configured.

Since there can be a large number of services transported by a single tunnel, the MEP executing at the tunnel-level reduces network overhead and simplifies the configuration.

Note: All services in the tunnel must share a common physical path.

A SAP is needed in order for the Tunnel MEP to extract the tunnel MEP ETH-CFM packets at the appropriate level. No SAP record is created by default. A service must already exist that includes a SAP in the form lag-id:vid.* or lag-id:vid.0 where the vid matches the outer VLAN in which the tunnel is to monitor. Since the ETH-CFM traffic arrives at the Ethernet aggregation node as a single outer tag with no inner tag, the operator may want to consider the ability to configure the lag-id:vid.0 to accept untagged only frames with the matching outer tag and no inner tag. The global command **config>system->ethernet>new-qinq-untagged-sap** is available to enable this functionality. By default both the vid.* and vid.0 accepts all packets that match the outer vid and any inner vid. If no SAP record exists for this VLAN, one must be created manually. Manually creating this SAP requires a service context. Nokia recommends that an Epipe service be configured with this single SAP, preventing any flooding of packets. It is possible to use a VPLS instance and combine many tunnel SAP records into a single service instance. However, configuration errors may result in leakage because of the multipoint nature of a VPLS service. Regardless of the service type chosen, it should be in a shutdown state. Also, normal ETH-CFM rules apply. ETH-CFM packets arriving on the SAP passes all ETH-CFM packets at and below the tunnel MEP to the ETH-CFM application for processing.

The goal of a Tunnel MEP is to validate an attachment circuit and relate the state to services that share the same LAG and outer VLAN to other services across the network. Tunnel MEPs are not intended for propagating fault between two endpoints that share the same LAG and outer VLAN. For this reason, locally switched circuits that share the same LAG and the same outer tag must not use the **ais-enable** function under those SAPs. As an example, lag-1 may have two SAPs associated with it: lag-1:1.1 and lag-1:1.2. These two SAP represent two different endpoints on the same LAG using the same outer VLAN. In this case, if the ais-enable is configured under both SAPs, AIS functionality does not work properly. Normal fault propagation could be used in this case instead. Since the tunnel MEP is validating the common physical path and these two MEPs share the common physical path, there is no reason to propagate fault. Service-based MEPs could be configured on the endpoints in order to validate the connectivity between the two endpoints when this type of model is deployed. However, two SAPs that are connected to different LAGs is a supported configuration. An example of this would be lag-1:1.1 and lag-2:1.1.

Sub second Tunnel MEPs will be monitored for every three seconds to ensure that they are not continuously bouncing and consuming an unfair allocation of ETH-CFM resources. A sub second MEP will only be allowed three operational status changes in a three second window before holding the state for the remaining time in that window. Messages will be paced from Tunnel MEPs. Fault propagation depends on factors such as how busy the node is, or how scaled the node configuration is.

Five percent of the operational/negotiated port speed not physical speed is available for Tunnel MEP control traffic. When applying this to the LAG-based Tunnel MEPs the five percent is derived from the lowest speed of a single member port in the bundle. If this bandwidth percentage required for ETH-CFM is exceeded the ETH-CFM packets will not be able to be sent and failures will occur. As an example, a physical port of 1Gb/s that has negotiated an operational speed of 100 Mb/s with a peer will be allowed to send up to a maximum of 5 Mb/s of Tunnel MEP control traffic.

Example: Tunnel MEP Configuration

Figure 41 shows how fate can be shared between the Tunnel MEP and the services configured on the same LAG and outer VLAN.

Figure 41 **Tunnel MEP Example**

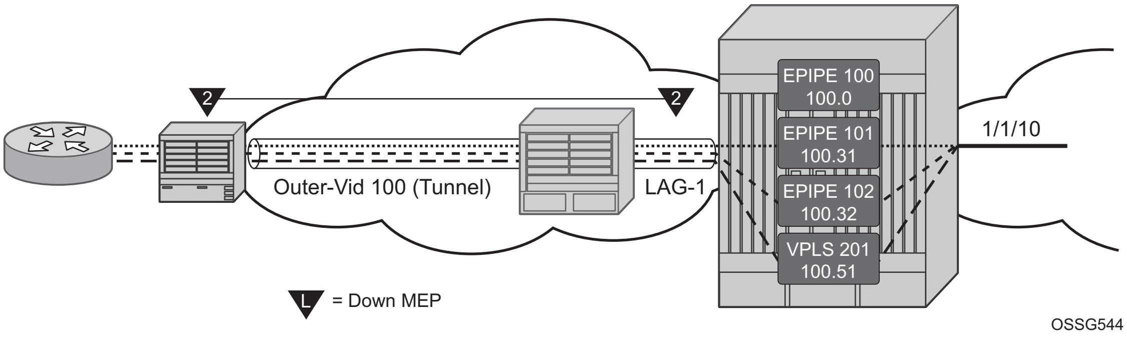

In this example, a single Tunnel, LAG-1 outer VLAN 100, carries three services. Epipe 101, Epipe 102 and VPLS 201 are the service extraction points on the aggregation node. Epipe 100 is the extraction point for the Tunnel MEP eth-cfm traffic. This is a single SAP Epipe that is operationally shutdown. One common configuration error when using Tunnel MEPs is the lack extraction on the aggregation node, causing unidirectional failures. The aggregation node is sending eth-cfm traffic to the NID, but is not extracting the eth-cfm traffic that the NID is sending.

Epipe 101 is configured to accept the tunnel MEP fate and generate AIS.

Epipe 102 is configured to accept the tunnel MEP state and apply fault propagation rules. If the network-side mate were an SDP binding, then the applicable setting of the LDP status bits are in the header. Since this example uses an Ethernet SAP as the mate, and only tunnel fault-accept is configured with no ais-enable, only the SAP flag is set to indicate an error.

VPLS 201 also shares the fate of the tunnel MEP. The tunnel-fault accept transitions the SAP to operationally down. Any configured event that occurs because of a SAP down for the VPLS also occur.

Only the configuration for the aggregation node is shown below. The NID configuration is not required to show how this function works.

Aggregation node

```
config>eth-cfm# info
----------------------------------------------
        domain 2 format none level 2
            association 1 format icc-based name "FacilityTun01"
                ccm-interval 1
                remote-mepid 101
            exit
        exit
----------------------------------------------

config>lag# info
----------------------------------------------
```

```
        mode access
        encap-type qinq
        eth-cfm
            mep 100 domain 2 association 1 vlan 100
                description "Tunnel Facility MEP - Do NOT Delete"
                ais-enable
                    client-meg-level 5
                exit
                facility-fault
                ccm-enable
                low-priority-defect allDef
                no shutdown
            exit
        exit
        port 1/1/2
        no shutdown
----------------------------------------------

config>service# info
----------------------------------------------
        customer 1 create
            description "Default customer"
        exit
        epipe 100 customer 1 create
            shutdown
            description "Tunnel Extraction Service"
            sap lag-1:100.0 create
            exit
        exit
        epipe 101 customer 1 create
            description "Customer Service 100.31"
            sap 1/1/10:100.31 create
            exit
            sap lag-1:100.31 create
                eth-cfm
                    ais-enable
                exit
            exit
            no shutdown
        exit
        epipe 102 customer 1 create
            description "Customer Service 100.32"
            eth-cfm
                tunnel-fault accept
            exit
            sap 1/1/10:100.32 create
            exit
            sap lag-1:100.32 create
            exit
            no shutdown
        exit
        vpls 201 customer 1 create
            description "Customer Service 100.51"
            stp
                shutdown
            exit
            eth-cfm
                tunnel-fault accept
            exit
            sap 1/1/10:100.51 create
            exit
            sap lag-1:100.51 create
            exit
            no shutdown
        exit
----------------------------------------------


# show eth-cfm mep 100 domain 2 association 1
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index           : 2                        Direction         : Down
Ma-index           : 1                        Admin             : Enabled
MepId              : 100                      CCM-Enable        : Enabled
Port               : lag-1                    VLAN              : 100
Description        : Tunnel Facility MEP - Do NOT Delete
FngState           : fngReset                 ControlMep        : False
LowestDefectPri    : allDef                   HighestDefect     : none
Defect Flags       : None
Mac Address        : 90:f3:ff:00:01:41        ControlMep        : False
CcmLtmPriority     : 7
CcmTx              : 3958                     CcmSequenceErr    : 0
Fault Propagation  : disabled                 FacilityFault     : Notify
MA-CcmInterval     : 1                        MA-CcmHoldTime    : 0ms
Eth-1Dm Threshold  : 3(sec)                   MD-Level          : 2
Eth-Ais:           : Enabled                  Eth-Ais Rx Ais:   : No
Eth-Ais Tx Priorit*: 7                        Eth-Ais Rx Interv*: 1
Eth-Ais Tx Interva*: 1                        Eth-Ais Tx Counte*: 175
Eth-Ais Tx Levels  : 5
Eth-Tst:           : Disabled

Redundancy:
    MC-LAG State   : n/a

CcmLastFailure Frame:
    None

XconCcmFailure Frame:
    None
===============================================================================

# show eth-cfm cfm-stack-table facility all-tunnel-meps
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM, A = AisRx
===============================================================================
CFM Facility LAG Stack Table
===============================================================================
Lag      Tunnel    Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
lag-1     100        2 Down          2          1  100 90:f3:ff:00:01:41 ------
===============================================================================

# show service sap-using eth-cfm facility
===============================================================================
Service ETH-CFM Facility Information
===============================================================================
SapId             SvcId                          SAP AIS  SAP Tunnel SVC Tunnel
                                                          Fault      Fault
-------------------------------------------------------------------------------
lag-1:100.0       100                            Disabled Accept     Ignore
lag-1:100.31      101                            Enabled  Accept     Ignore
lag-1:100.32      102                            Disabled Accept     Accept
lag-1:100.51      201                            Disabled Accept     Accept
-------------------------------------------------------------------------------
No. of Facility SAPs: 4
===============================================================================
```

When the tunnel MEP enters a fault state

- Epipe 101 will start to generate AIS out the mate sap
- Epipe 102 SAP flag will be set
- VPLS 201 SAP will go down

Output from one of the nodes is included below. Since both will react in the same manner output from both nodes is not required.

Aggregation node

```
# show eth-cfm cfm-stack-table facility all-tunnel-meps
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM, A = AisRx
===============================================================================
CFM Facility LAG Stack Table
===============================================================================
Lag      Tunnel    Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
lag-1     100        2 Down          2          1  100 90:f3:ff:00:01:41 --C---
===============================================================================

# show service sap-using eth-cfm facility tunnel 100
===============================================================================
Service ETH-CFM Facility Information
===============================================================================
SapId             SvcId                          SAP AIS  SAP Tunnel SVC Tunnel
                                                          Fault      Fault
-------------------------------------------------------------------------------
lag-1:100.0       100                            Disabled Accept     Ignore
lag-1:100.31      101                            Enabled  Accept     Ignore
lag-1:100.32      102                            Disabled Accept     Accept
lag-1:100.51      201                            Disabled Accept     Accept
-------------------------------------------------------------------------------
No. of Facility SAPs: 4
===============================================================================

# show eth-cfm mep 100 domain 2 association 1
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index           : 2                        Direction         : Down
```

```
Ma-index           : 1                        Admin            : Enabled
MepId              : 100                      CCM-Enable       : Enabled
Port               : lag-1                    VLAN             : 100
Description        : Tunnel Facility MEP - Do NOT Delete
FngState           : fngDefectReported        ControlMep       : False
LowestDefectPri    : allDef                   HighestDefect    : defRemoteCCM
Defect Flags       : bDefRemoteCCM
Mac Address        : 90:f3:ff:00:01:41        ControlMep       : False
CcmLtmPriority     : 7
CcmTx              : 4211                     CcmSequenceErr   : 0
Fault Propagation  : disabled                 FacilityFault    : Notify
MA-CcmInterval     : 1                        MA-CcmHoldTime   : 0ms
Eth-1Dm Threshold  : 3(sec)                   MD-Level         : 2
Eth-Ais:           : Enabled                  Eth-Ais Rx Ais:  : No
Eth-Ais Tx Priorit*: 7                        Eth-Ais Rx Interv*: 1
Eth-Ais Tx Interva*: 1                        Eth-Ais Tx Counte*: 215
Eth-Ais Tx Levels  : 5
Eth-Tst:           : Disabled

Redundancy:
    MC-LAG State   : n/a

CcmLastFailure Frame:
    None

XconCcmFailure Frame:
    None
===============================================================================

show service id 101 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 101                 Vpn Id            : 0
Service Type      : Epipe
Name              : (Not Specified)
Description       : Customer Service 100.31
Customer Id       : 1
Last Status Change: 02/04/2010 15:53:12
Last Mgmt Change  : 02/04/2010 16:31:00
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 2                   SDP Bind Count    : 0
Per Svc Hashing   : Disabled
Force QTag Fwd    : Disabled

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/10:100.31                        qinq         1522    1522    Up   Up
sap:lag-1:100.31                         qinq         1522    1522    Up   Up
===============================================================================

# show service id 102 base
===============================================================================
Service Basic Information
```

```
===============================================================================
Service Id        : 102                 Vpn Id            : 0
Service Type      : Epipe
Name              : (Not Specified)
Description       : Customer Service 100.32
Customer Id       : 1
Last Status Change: 02/04/2010 15:45:07
Last Mgmt Change  : 02/04/2010 16:30:43
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 2                   SDP Bind Count    : 0
Per Svc Hashing   : Disabled
Force QTag Fwd    : Disabled

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/10:100.32                        qinq         1522    1522    Up   Up
sap:lag-1:100.32                         qinq         1522    1522    Up   Up
===============================================================================

# show service id 102 sap lag-1:100.32
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 102
SAP                : lag-1:100.32             Encap             : qinq
QinQ Dot1p         : Default
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Up
Flags              : OamTunnelMEPFault
Multi Svc Site     : None
Last Status Change : 02/04/2010 15:45:07
Last Mgmt Change   : 02/04/2010 15:44:26

-------------------------------------------------------------------------------
ETH-CFM SAP specifics
-------------------------------------------------------------------------------
Tunnel Faults      : accept                   AIS               : Disabled
MC Prop-Hold-Timer : n/a
===============================================================================

*A:PE-6# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPLS
Name              : 1
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 05/08/2018 09:40:32
Last Mgmt Change  : 05/08/2018 09:40:24
Etree Mode        : Disabled
Admin State       : Up                  Oper State        : Up
```

```
MTU               : 1514
SAP Count         : 1                     SDP Bind Count    : 1
Snd Flush on Fail : Disabled              Host Conn Verify  : Disabled
SHCV pol IPv4     : None
Propagate MacFlush: Disabled              Per Svc Hashing   : Disabled
Allow IP Intf Bind: Disabled
Fwd-IPv4-Mcast-To*: Disabled              Fwd-IPv6-Mcast-To*: Disabled
Mcast IPv6 scope  : mac-based
Def. Gateway IP   : None
Def. Gateway MAC  : None
Temp Flood Time   : Disabled              Temp Flood        : Inactive
Temp Flood Chg Cnt: 0
SPI load-balance  : Disabled
TEID load-balance : Disabled
Src Tep IP        : N/A
Vxlan ECMP        : Disabled
VSD Domain        : <none>

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/c1/1:1                           q-tag        9000    9000    Up   Up
sdp:65:1 S(192.0.2.5)                    Spok         0       8974    Up   Down
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

2.15.1.6 Router Interface MEP

MEPs and associated on-demand troubleshooting functions act as router interfaces that are part of the base routing instance. This feature allows the operator to verify Layer 2 transport that connects the Layer 3 interfaces.

Router interfaces MEPs are supported for all router interface instances (null port 1/1/1, dot1q port 1/1/3:vid, null LAG-lag-id and dot1q LAG-lag-id:vid).

Example: Router MEP Configuration

The following illustration, Figure 42, shows how a Router Facility MEP can be configured on a routed interface in the base router instance.

Figure 42 **Router MEP Example**

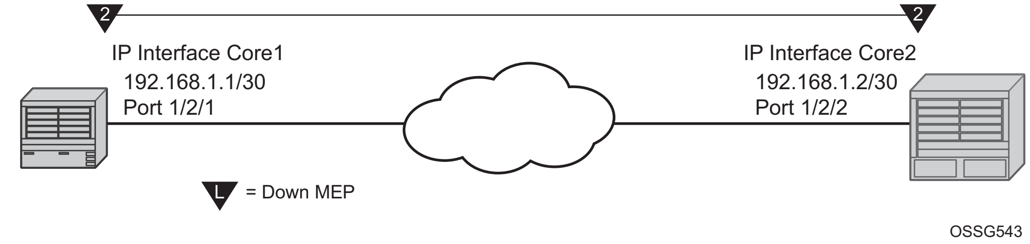

ETH-CFM tools for proactive management (ETH-CC), troubleshooting (Loopback, Linktrace, and so on) and profiling (Delay Measurement, and so on) are supported. The configuration and some ETH-CFM test commands are shown for Node1 (left). Following the on-demand test output, the configuration for Node 2 is included for completeness, without repeating the on-demand tests.

NODE1

```
config>port# info
----------------------------------------------
        ethernet
        exit
        no shutdown
----------------------------------------------

config>eth-cfm# info
----------------------------------------------
        domain 2 format none level 2
            association 2 format icc-based name "FacilityRtr01"
            exit
        exit
----------------------------------------------

config>router# info
----------------------------------------------
#--------------------------------------------------
echo "IP Configuration"
#--------------------------------------------------
        interface "Core1"
            address 192.168.1.1/30
            port 1/2/1
            eth-cfm
                mep 1 domain 2 association 2
                    mac-address d0:0d:1e:00:00:01
                    no shutdown
                exit
            exit
        exit
        interface "system"
        exit
----------------------------------------------

# show eth-cfm cfm-stack-table facility all-router-interfaces
```

```
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM, A = AisRx

===============================================================================
CFM Facility Interface Stack Table
===============================================================================
Interface          Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
Core1                2 Down          2          2    1 d0:0d:1e:00:00:01 ------
===============================================================================

# show eth-cfm cfm-stack-table facility all-router-interfaces
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM, A = AisRx

===============================================================================
CFM Facility Interface Stack Table
===============================================================================
Interface          Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
Core1                2 Down          2          2    1 d0:0d:1e:00:00:01 ------
===============================================================================

# oam eth-cfm loopback d0:0d:1e:00:00:02 mep 1 domain 2 association 2
 send-count 5
Eth-Cfm Loopback Test Initiated: Mac-Address: d0:0d:1e:00:00:02, out service: 0
Sent 5 packets, received 5 packets [0 out-of-order, 0 Bad Msdu]

# oam eth-cfm linktrace d0:0d:1e:00:00:02 mep 1 domain 2 association
2
Index Ingress Mac          Egress Mac           Relay      Action
----- -------------------- -------------------- ---------- ----------
1     D0:0D:1E:00:00:02    00:00:00:00:00:00    n/a        terminate
----- -------------------- -------------------- ---------- ----------
No more responses received in the last 6 seconds.

# oam eth-cfm two-way-delay-test d0:0d:1e:00:00:02 mep 1 domain 2 association 2
Two-Way-Delay-Test Response:
Delay 1130 microseconds        Variation 63 microseconds

# oam eth-cfm two-way-delay-test d0:0d:1e:00:00:02 mep 1 domain 2 association 2
Two-Way-Delay-Test Response:
Delay 1218 microseconds        Variation 88 microseconds
```

NODE2

```
config>port# info
----------------------------------------------
        ethernet
        exit
        no shutdown
----------------------------------------------

config>eth-cfm# info
----------------------------------------------
```

```
        domain 2 format none level 2
            association 2 format icc-based name "FacilityRtr01"
            exit
        exit
----------------------------------------------

config>router# info
----------------------------------------------
#--------------------------------------------------
echo "IP Configuration"
#--------------------------------------------------
        interface "Core2"
            address 192.168.1.2/30
            port 1/2/2
            eth-cfm
                mep 2 domain 2 association 2
                    mac-address d0:0d:1e:00:00:02
                    no shutdown
                exit
            exit
        exit
        interface "system"
        exit
----------------------------------------------
```

2.15.1.7 Hardware Support

This section applies to the 7750 SR and 7450 ESS. However, only the facility MEP has an IOM-specific requirement. SAPs and ports that are not configured as part of facility MEPs are not restricted to a specific IOM. For example, a Tunnel MEP would be required to meet the minimum IOM requirement, similar to the fated shared service SAPs. However, the mate or egress SAP or binding is not required to meet the facility MEP requirement. Of course, there may be other reasons why a mate SAP or binding requires specific IOM/IMM that are outside that of facility MEPs. Similarly, a LAG MEP requires all port members to meet the IOM/IMM requirements for facility MEPs.

Table 13 provides an overview of Facility MEP support.

Table 13 **Facility MEP Support Overview**

	Port MEPs	**Tunnel MEPs**		**LAG MEPs**	**Router MEPs**
		Port	**LAG**		
Sub Second	Yes	Yes	Yes	Yes	Yes
Port:					

Table 13 **Facility MEP Support Overview (Continued)**

	Port MEPs	Tunnel MEPs		LAG MEPs	Router MEPs
		Port	LAG		
Hybrid Network Access	Dot1q/QinQ Null/Dot1q Null/Dot1q/QinQ	QinQ no QinQ	QinQ no QinQ	Dot1q/QinQ Null/Dot1q Null/Dot1q/QinQ	Dot1q/QinQ Null/QinQ N/A
CCM	Yes	Yes	Yes	Yes	Yes
Y.1731 PM Tools	Yes	Yes	Yes	Yes	Yes
AIS Reception	No	Yes	Yes	No	No
Facility Fault	Controls port operational state Failure=Link Up Success=Up	Controls shared fate service SAPs and EPIPE AIS	Controls Shared fate service SAPs and Epipe AIS	Controls LAG operational state Failure=Oper: down, Success=Oper=up	Controls IP interface operational state in reaction to CFM state
Mutually Exclusive			Mutually Exclusive		

Sub-second CCM-enabled MEPs are platform-specific and may not be supported uniformly on the 7750 SR, 7450 ESS and 7950 XRS platforms. For those 7750 SR, 7450 ESS and 7950 XRS platforms which support sub-second CCM-enabled MEPs, this additional restriction is applied; QinQ tunnel MEPs require a minimum of SF/CPM3.

2.15.2 ETH-CFM and MC-LAG

By default, ETH-CFM Management Points (MEPs and MIPs) and MC-LAG operate independently. Nokia recommends not enabling fault propagation when the default behavior is in use. A global command is available in order to allow ETH-CFM the ability to track the state of the MC-LAG for MPs that are configured on MC-LAG ports. This feature does not allow MEPs to influence MC-LAG state. Since the MP relies heavily on the underlying MC-LAG construct, consideration must be given for the proper MC-LAG design and deployment. It is important to understand that the state of MC-LAG can be reflected in the state of the MPs which are configured on SAPs that are part MC-LAGs. For example, a SAP on a LAG that is part of an MC-LAG configuration can behave in a manner that more appropriately represents the MC-LAG.

2.15.2.1 ETH-CFM and MC-LAG Default Behavior

ETH-CFM MPs track the SAPs, bindings and facility independently. Therefore, when an MP is configured on a SAP which is not operationally up because of MC-LAG ETH-CFM defect, conditions are raised for what could be considered normal conditions. Figure 43 shows the default behavior for a point-to-point service without regard for MC-LAG. In the case below, the two up MEPs operating at level 4 on the affected SAPs set the **Interface-Status-TLV** bit in the ETH-CC header to represent the **isDown** condition, assuming ETH-CC is executing between the peer MEPs. This is the correct action based on the ETH-CFM perspective, SAPs are operationally **down**.

Figure 43 **Independent Processing UP MEP Example**

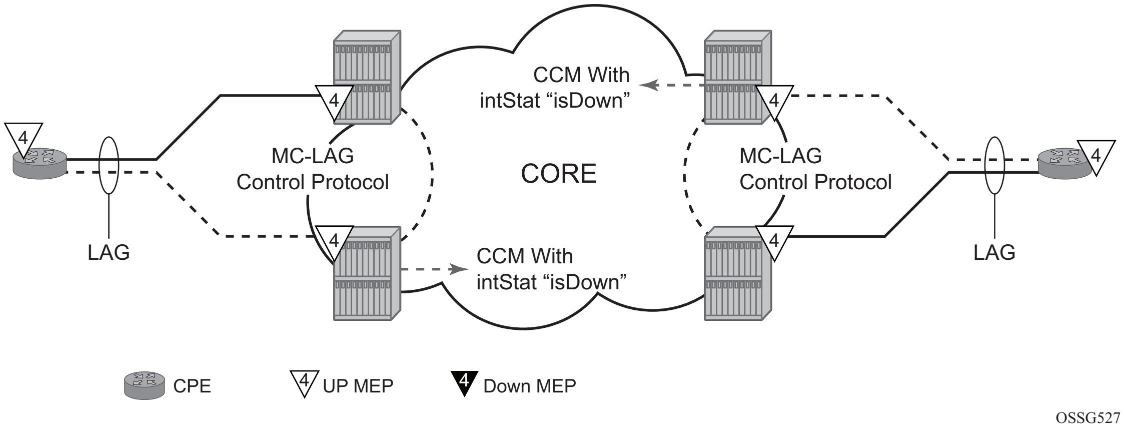

A similar condition exists if down MEPs are configured on the SAPs that are operationally down. Figure 44 shows how the same service configured with down MEPs would generate AIS, if enabled, toward the remote client at the configured client-meg-level, in the reverse direction of the MEP. This is also the proper behavior from the perspective ETH-CFM.

Figure 44 **Independent Processing Down MEP Example**

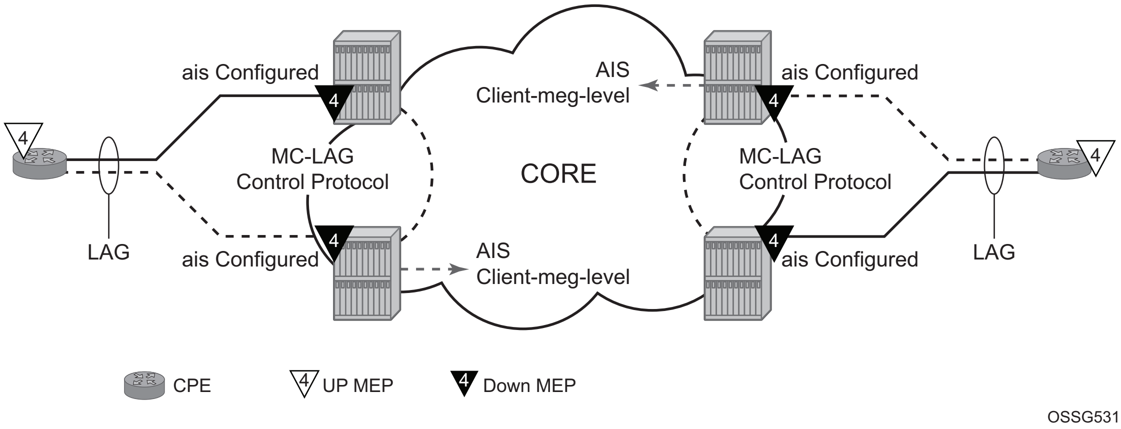

2.15.2.2 Linking ETH-CFM to MC-LAG State

Allowing ETH-CFM to understand the state of MC-LAG and adjust the behavior of the MP (MEP and MIP) according to that state has benefits.

MC-LAG represents the two upstream nodes as a single system to the node terminating a standard LAG. Linking the ETH-CFM MPs to the state of the MC-LAG allows the operator to configure MPs across the two boxes that appear the same. Under the default configuration, this would introduce various defect conditions to be raised and event conditions. However, when ETH-CFM is tracking the state of the MC-LAG, the MPs performs a role that represents the state of the resiliency mechanism. In order to enable this new behavior, configure the system-wide command **standby-mep-shutdown** under the **config>eth-cfm>redundancy>mc-lag** hierarchy.

When a MP is part of the active MC-LAG system, it performs as a normal MP: terminating, generating, responding to, and processing all appropriate ETH-CFM packets. An MP that is on the standby MC-LAG node enters a pseudo-shutdown state. These MPs terminates all ETH-CFM that are part of the regular interception process, but will not process them. They are silently discarded. Also, an MP that exists on a standby MC-LAG system does not generate any ETH-CFM packets. All proactive and on-demand functions are blocked on the standby MC-LAG node. When scheduled tests are executed through SAA these test will attempt to execute. The tests will record failures as a result of the MEP state. These failures are not representative of the network.

This feature relies on the proper configuration, design, and deployment of the MC-LAG protocol. There are numerous optimizations and configuration parameters that are available as part of the MC-LAG functions. For example, by default, when a currently active MC-LAG port transitions to standby, by any means including manual operator intervention, the remote node terminating the standard LAG sees the LAG transition because all ports in the LAG are down for an instance in time. This is standard LAG behavior does not change as a result of the linkage of MP state to MC-LAG state. This transition causes the propagation of faults for MEPs configured on that node. Normal architectural LAG design must take these types of events into consideration. MC-LAG provides numerous tuning parameters that need to be considered before deploying in the field. These include a **hold-time down** option on the node terminating the standard LAG, as well as other parameters for revertive behavior such as the **hold-time** up option. It is important to ensure that the operator's specific environment be taken into consideration when tuning the MC-LAG parameters to avoid the propagation of error conditions during normal recover events. In the case that the resumption of data forwarding exceed the timeout value of a MEP (3.5 times the CCM-Interval), the appropriate defect conditions are raised.

ETH-CFM will register a fault propagation delay timer equal to **propagate-hold-time** under the **config>eth-cfm>redundancy>mc-lag** hierarchy (default of 1s) to delay notification of an event that may be a result of MC-LAG failover. This allows the system time to coordinate events and triggers that together represent the MC-LAG transition from active to standby.

A fixed timer value of 1s will delay an UP MEP from announcing a SAP down condition through CCM Interface-Status-TLV bits, is Down. ETH-CFM maintains a status of last sent to the UP MEPs peer. When the SAP transitions either to UP or DOWN that fault will be held for the fixed 1s interval and the last Interface-Status-TLV bits will set based on the previous transmission. If the condition, different from the previous sent, still exists at the end of the 1s fixed timer and when the next CCM interval expires, the representative value of the SAP will be sent in the Interface-Status-TLV. These two timers help to smooth out network transitions at the cost of propagation and clearing of faults.

When a node with ETH-CFM linked to MC-LAG is transitioning from standby to active ETH-CFM will assume there are no underlying conditions for any of the SAPs that are now part of the newly activating MC-LAG. The initial notification to an UP MEPs peer will not include any faults. It will assume that the transitioning SAPs are stabilizing as the switchover proceeds. The fixed 1s timer will be starting and a second CCM PDU based on the UP MEPs interval will be sent without any recognition of potential fault on the SAP. However, after the expiration of the fixed timer and on the next CCM-Interval, the Interface-Status-TLV will represent the state of the SAP.

In scaled environments it is important to configure the propagation-hold-time and the CCM intervals to achieve the desired goals. If these timers are set too aggressively, then fault and defect conditions may be generated during times of network stabilization. The use of fault propagation and AIS transmission needs to be carefully considered in environments where MC-LAG protection mechanisms are deployed. Timer values do not guarantee that transitional state will not be propagated to the peer. The propagation of such state may be more taxing and disruptive that allowing the transmission states to complete. For example, if AIS generation is being used in this type of solution the operator should use a 60s AIS interval to avoid transitional state from being advertised.

AIS generation is paced in a first come first serve model not to exceed the system capability, scale is dependent on the type of system. If AIS is configured in an MC-LAG solution the operator must make sure that the same MEPs on each system are configured to generate AIS and this number does not exceed the maximum. This would require the operator to configure both nodes with the same MEPs that can generate AIS and not exceed the system capacity. If the nodes are configured differently or exceed the system scale there is a very high potential where a transition may see a different set of MEPs pacing out the AIS than the original set of MEPs. There is no synchronization of AIS state across nodes.

Administrative functions, like **admin down**, are special cases. When the administrative state changes from **up** to **down**, the timer is bypassed and communication from ETH-CFM is immediate.

When an MP is configured in an MC-LAG environment, Nokia recommends that each aspect of the MP be configured the same, including MAC address. Also, although this may be obvious, both nodes participating in the MC-LAG requiring this functionality should include the global command in the **config>eth-cfm>redundancy>mc-lag>standby-mep>shutdown** context to avoid unpredictable behavior.

In summary, a SAP with ETH-CFM tracking the state of the MC-LAG represents the state of the MC-LAG. MPs configured on the standby MC-LAG ports enters a state similar to shutdown. MPs on the MC-LAG ports on the active MC-LAG ports performs all normal processing.

Example: ETH-CFM and MC-LAG Configuration

The following illustration, shows how MEPS can be linked to MC-LAG state. In this example, a service MEP is created on the LAG SAP on NODE1 within service VPLS 100. The MEPs configured on the MC-LAG nodes within service 100 are both configured the same. Both MEPs use the same MEP-ID, the same MAC address.

Figure 45 **ETH-CFM and MC-LAG Example**

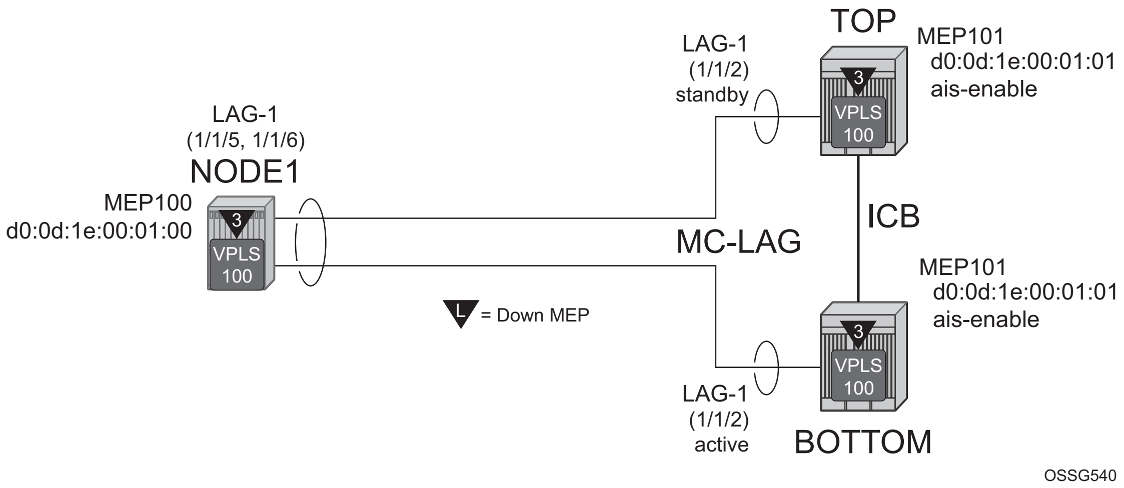

Only one of the MEPs on the MC-LAG nodes is active for VPLS service 100. The other MEP is in a shutdown mode, so that even when the MC-LAG is in standby and the port state is **Link Up**, the MEP is in a pseudo shutdown state.

The following configuration example is not meant to provide all possible MC-LAG configuration statement to tune each provider’s network. It does provide a base configuration to demonstrate the ETH-CFM feature.

NODE1

```
config>port# info (both ports)
----------------------------------------------
        ethernet
            mode access
            encap-type qinq
            autonegotiate limited
        exit
        no shutdown
----------------------------------------------


config>lag# info
----------------------------------------------
 mode access
        encap-type qinq
        access
            adapt-qos link
        exit
        port 1/1/5
        port 1/1/6
        lacp active administrative-key 32768
hold-time down 10
        no shutdown
----------------------------------------------
```

```
config>eth-cfm# info
----------------------------------------------
        domain 3 format none level 3
            association 1 format icc-based name "03-0000000100"
                bridge-identifier 100
                exit
                ccm-interval 1
                remote-mepid 101
            exit
        exit
----------------------------------------------

config>service>vpls# info
----------------------------------------------
            stp
                shutdown
            exit
            sap 1/1/3:100.100 create
            exit
            sap lag-1:100.100 create
                eth-cfm
                    mep 100 domain 3 association 1 direction down
                        ccm-enable
                        mac-address d0:0d:1e:00:01:00
                        no shutdown
                    exit
                exit
            exit
            no shutdown
----------------------------------------------


TOP (MC-LAG Standby)
config>port# info
----------------------------------------------
        ethernet
            mode access
            encap-type qinq
            autonegotiate limited
        exit
        no shutdown
----------------------------------------------

config>lag# info
----------------------------------------------
        mode access
        encap-type qinq
        access
            adapt-qos link
        exit
        port 1/1/2
        lacp active administrative-key 32768
        no shutdown
----------------------------------------------

config>router# info
----------------------------------------------
#--------------------------------------------------
echo "IP Configuration"
```

```
#---------------------------------------------------
        interface "Core2"
            address 192.168.1.2/30
            port 1/2/2
        exit
        interface "system"
        exit
----------------------------------------------

config>redundancy# info
----------------------------------------------
        multi-chassis
            peer 192.168.1.1 create
                source-address 192.168.1.2
                mc-lag
                    lag 1 lacp-key 1 system-id 00:00:00:00:00:01 system-priority
 100
                    no shutdown
                exit
                no shutdown
            exit
        exit
        synchronize boot-env
----------------------------------------------

config>eth-cfm# info
----------------------------------------------
        domain 3 format none level 3
            association 1 format icc-based name "03-0000000100"
                bridge-identifier 100
                exit
                ccm-interval 1
                remote-mepid 100
            exit
        exit
        redundancy
            mc-lag
                standby-mep-shutdown
            exit
        exit
----------------------------------------------

config>service>vpls# info
----------------------------------------------
            stp
                shutdown
            exit
            sap lag-1:100.100 create
                eth-cfm
                    mep 101 domain 3 association 1 direction down
                        exit
                        ccm-enable
                        mac-address d0:0d:1e:00:01:01
                        no shutdown
                    exit
                exit
            exit
            no shutdown
----------------------------------------------
```

```
# show lag 1
===============================================================================
Lag Data
===============================================================================
Lag-id         Adm     Opr     Port-Threshold   Up-Link-Count   MC Act/Stdby
-------------------------------------------------------------------------------
1              up      down    0                0               standby
===============================================================================

# show port
===============================================================================
Ports on Slot 1
===============================================================================
Port        Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id          State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
… snip …
1/1/2       Up    Yes  Link Up 1522 1522    1 accs qinq xcme
…snip…
==========================================================================
```

```
BOT (MC-LAG Active)
config>port# info
----------------------------------------------
        ethernet
            mode access
            encap-type qinq
            autonegotiate limited
        exit
        no shutdown
----------------------------------------------

config>lag# info
----------------------------------------------
        mode access
        encap-type qinq
        access
            adapt-qos link
        exit
        port 1/1/2
        lacp active administrative-key 32768
        no shutdown
----------------------------------------------

config>router# info
----------------------------------------------
#--------------------------------------------------
echo "IP Configuration"
#--------------------------------------------------
        interface "Core1"
            address 192.168.1.1/30
            port 1/2/1
        exit
        interface "system"
        exit
----------------------------------------------
```

```
config>redundancy# info
----------------------------------------------
        multi-chassis
            peer 192.168.1.2 create
                source-address 192.168.1.1
                mc-lag
                    lag 1 lacp-key 1 system-id 00:00:00:00:00:01 system-priority
 100
                    no shutdown
                exit
                no shutdown
            exit
        exit
        synchronize boot-env
----------------------------------------------

config>eth-cfm# info
----------------------------------------------
        domain 3 format none level 3
            association 1 format icc-based name "03-0000000100"
                bridge-identifier 100
                exit
                ccm-interval 1
                remote-mepid 100
            exit
        exit
        redundancy
            mc-lag
                standby-mep-shutdown
            exit
        exit
----------------------------------------------

config>service>vpls# info
----------------------------------------------
            stp
                shutdown
            exit
            sap lag-1:100.100 create
                eth-cfm
                    mep 101 domain 3 association 1 direction down
                        exit
                        ccm-enable
                        mac-address d0:0d:1e:00:01:01
                        no shutdown
                    exit
                exit
            exit
            no shutdown
----------------------------------------------

# show lag 1
===============================================================================
Lag Data
===============================================================================
Lag-id         Adm     Opr     Port-Threshold   Up-Link-Count   MC Act/Stdby
-------------------------------------------------------------------------------
1              up      up      0                1               active
===============================================================================
# show port
===============================================================================
Ports on Slot 1
===============================================================================
Port        Admin Link Port    Cfg  Oper LAG/ Port Port Port   C/QS/S/XFP/
Id          State      State   MTU  MTU  Bndl Mode Encp Type   MDIMDX
-------------------------------------------------------------------------------
…snip…
1/1/2       Up    Yes  Up      1522 1522    1 accs qinq xcme
…snip…

===============================================================================
```

2.15.3 ETH-CFM Features

2.15.3.1 CCM Hold Timers

In some cases the requirement exists to prevent a MEP from entering the defRemoteCCM defect, remote peer timeout, from more time than the standard 3.5 times the CCM-interval. Both the IEEE 802.1ag standard and ITU-T Y.1731 recommendation provide a non-configurable 3.5 times the CCM interval to determine a peer time out. However, when sub second CCM timers (10ms/100ms) are enabled the carrier may want to provide additional time for different network segments to converge before declaring a peer lost because of a timeout. In order to maintain compliance with the specifications the ccm-hold-timer down <delay-down> option has been introduced to artificially increase the amount of time it takes for a MEP to enter a failed state should the peer time out. This timer is only additive to CCM timeout conditions. All other CCM defect conditions, like defMACStatus, defXconCCM, and so on, will maintain their existing behavior of transitioning the MEP to a failed state and raising the proper defect condition without delay.

When the **ccm-hold-timer down** *delay-down* option is configured the following calculation is used to determine the remote peer time out (3.5 times the CCM-Interval + ccm-hold-timer delay-down).

This command is configured under the association. Only sub second CCM enabled MEPs support this hold timer. Ethernet-Tunnel Paths use a similar but slightly different approach and will continue to utilize the existing method. Ethernet-tunnels will be blocked from using this new hold timer.

It is possible to change this command on the fly without deleting it first. Simply entering the command with the new values will change to values without having to delete the command prior to the change.

It is possible to change the ccm-interval of a MEP on the fly without first deleting it. This means it is possible to change a sub second CCM enabled MEP to 1 second or above. The operator will be prevented from changing an association from a sub second CCM interval to a non-sub second CCM interval when **ccm-hold-timer** is configured in that association. The **ccm-hold-timer** must be removed using the no option prior to allowing the transition from sub second to non-sub second CCM interval.

2.15.3.2 CCM Interval

This section applies to the 7750 SR and 7450 ESS. Different service types support different ETH-CFM functionality. This is explained in the applicable service sections throughout this guide.

This feature is an enhancement that enables slow timers OAM handling of G.8032 enabling G.8032 on the 7750 SR-c4 and 7750 SR-c12 platforms. G.8032 uses the OAM for Ring Protection messages. This feature enables full G.8032 Ring support on these platforms. In addition, this feature enables Continuity Check messages (CCM) on Ring ports at 1 second intervals for all platforms where G.8032 is supported. With this feature, G.8032 can be configured on additional router platforms. CCM are optional with G.8032 but normally deployed for higher assurance of protection. The 7750 SR and 7450 ESS additionally support CCM of 100ms and 10ms. CCM is configured on a neighbor node basis so the only requirement is that neighbor switches be configured with same interval or with CCM disabled.

2.15.4 Configuring ETH-CFM Parameters

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

The configuration under the **config>eth-cfm** hierarchy defines the domains, associations, and the applicable global parameters for each of those contexts, including the linkage to the service using the bridge-identifier option. Once this configuration is complete, the Management Points (MPs = MEPs and MIPs) may be defined referencing the appropriate global context.

As described in the *7450 ESS, 7750 SR, 7950 XRS, and VSR OAM and Diagnostics Guide*, MEPs can be implemented at the service or the facility level. The focus of this guide is on how the ETH-CFM MPs are configured within the service hierarchy level. However, because of the wide range of features that the ITU-T has defined in recommendation Y.1731 (Fault Management, Performance Management and Protection Mechanisms) the features may be applied to other features and hierarchies. For example, Ethernet Ring Protection (G.8032) also make use of various ETH-CFM functions. Different section in this guide may contain ETH-CFM specific material as it applies to that specific feature.

The following is an example of how domains and associations could be constructed, illustrating how the different services are linked to the contexts.

```
config>eth-cfm# info
----------------------------------------------
        domain 3 format none level 3
            association 1 format icc-based name "03-0000000101"
                bridge-identifier 100
                exit
            exit
        exit
        domain 4 format none level 4
            association 1 format icc-based name "04-0000000102"
                bridge-identifier 100
                remote-mepid 200
                ccm-interval 60
                exit
            exit
        exit
```

The following configuration examples illustrate how different services make use of the domain and association configuration. An Epipe, VPLS, and IES service are shown in this example. Refer to the previous table that shows the supported services and the management points.

Note: The following examples cannot all be configured at the same instance because the service ID 100 cannot be spread across multiple services.

```
# configure service epipe 100 customer 1 create
* config>service>epipe# info
----------------------------------------------
            sap 1/1/2:100.31 create
                eth-cfm
                    mep 111 domain 3 association 1 direction down
 mac-address d0:0d:1e:00:01:11
                        no shutdown
                    exit
                exit
            exit
            sap 1/1/10:100.31 create
```

```
                eth-cfm
                    mep 101 domain 4 association 1 direction up
                        mac-address d0:0d:1e:00:01:01
 ccm-enable
                        no shutdown
                    exit
                exit
            exit
            no shutdown
----------------------------------------------

# configure service vpls 100 customer 1 create
* config>service>vpls# info
----------------------------------------------
            sap 1/1/2:100.31 create
                eth-cfm
                    mep 111 domain 3 association 1 direction down
 mac-address d0:0d:1e:00:01:11
                        no shutdown
                    exit
                exit
            exit
            sap 1/1/10:100.31 create
                eth-cfm
                    mep 101 domain 4 association 1 direction up
                        mac-address d0:0d:1e:00:01:01
 ccm-enable
                        no shutdown
                    exit
                exit
            exit
            no shutdown
----------------------------------------------

# configure service ies 100 customer 1 create
config>service>ies# info
----------------------------------------------
            interface "test" create
                address 10.1.1.1/30
                sap 1/1/9:100 create
                    eth-cfm
                        mep 111 domain 3 association 1 direction down
                            ccm-enable
                            no shutdown
                        exit
                    exit
                exit
            exit
            no shutdown
----------------------------------------------
```

A Virtual MEP (vMEP) is a MEP that is configured at the service level rather than on a SAP or SDP binding. A vMEP sends ETH-CFM to all the SAPs and SDP bindings in the VPLS, depending on the type of traffic. If it is multicast traffic, the packets forward out all SAPs and SDP bindings. Unicast traffic is forwarded appropriately based on the type of ETH-CFM packet and the forwarding tables. Packets inbound to a context containing a vMEP performs normal processing and forwarding through the data plane with a copying of the ETH-CFM packet delivered to the local MEP for the appropriate levels. The local MEP will determine whether or not it should process a copied inbound ETH-CFM frame acting in accordance with standard rules.

Configuring a vMEP is similar in concept to placing down MEPs on the individual SAPs and SDP bindings in the associated VPLS. This ensures that packets inbound to the service get redirected to the vMEP for processing. Proper domain nesting must be followed in order to avoid ETH-CFM error conditions.

vMEPs support VPLS, MVPLS, BVPLS, and I-VPLS contexts.

A vMEP in an I-VPLS context can only extract packets inbound on local SAP and SDP bindings. This extraction does not include packets that are mapped to the I-VPLS from an associated B-VPLS context. If this type of extraction is required in an I-VPLS context then UP MEPs are required on the appropriate SAPs and SDP bindings in the I-VPLS service.

The wider scope of the vMEP feature requires all the SAPs within the service and every network port on the node to be FP2 or higher hardware.

As with the original vMEP functionality introduced for B-VPLS contexts, DOWN MEPs are supported on the individual SAPs or SDP bindings as long as domain nesting rules are not violated. Of course, local UP MEPs are only supported at the same level as the vMEP otherwise various CCM defect conditions will be raised, assuming CCM is enabled, and leaking of ETH-CFM packets will occur (lower level ETH-CFM packets arriving on a lower level MEP). Domain nesting must be properly deployed to avoid unexpected defect conditions and leaking between ETH-CFM domains.

MIPs map be configured on the SAPs and spoke SDPs at or above level of the vMEP.

An optional **vmep-filter** provides a coarse means of silently dropping all ETH-CFM packets that would normally be redirected to the CPU following egress processing. These includes any ETH-CFM level equal to or lower than the vMEP and any level equal to and lower than any other Management Points on the same SAP or SDP binding that includes the **vmep-filter**. MIPs will automatically be deleted when they coexist on the same SAP or spoke SDP as the **vmep-filter**. Since DOWN MEPs are ingress processed they are supported in combination with a vMEP and operate normally regardless of any **vmep-filter**. Domain nesting rules must be adhered to.

If the operator requires an MP on the SAP or SDP binding an UP MEP may be created at the same level as the vMEP on the appropriate SAP or SDP binding to perform the same function as the filter but at the specific level of the MEP. Scalability needs to be clearly understood because this will redirect the ETH-CFM packets to the CPU (consider using CPU protection). Consideration must also be given to the impact this approach could have on the total number of MEPs required. There are a number of other approaches that may lend themselves to the specific network architecture.

vMEP filtering is not supported within the a PBB VPLS since it already provides separation between B-components (typically the core) and I-components (typically the customer)

vMEPs do not support any ETH-AIS functionality and do not support fault propagation functions.

The following shows a configuration sample to configure a vMEP in a VPLS context.

```
config>service# vpls 100 customer 1 create
config>service>vpls$ info
----------------------------------------------
  stp
 shutdown
  exit
    eth-cfm
     mep 100 domain 3 association 1
         mac-address d0:0d:1e:00:01:11
   ccm-enable
         no shutdown
      exit
  exit
  no shutdown
----------------------------------------------
```

2.16 Configuring NGE with CLI

NGE is fully managed by the NSP NFM-P. The NSP NFM-P ensures proper network synchronization of key groups, services, and NGE domains. Managing NGE without the NSP NFM-P is not recommended. See the *NSP NFM-P User Guide* for more information.

This section provides information about configuring NGE using the command line interface.

Topics in this chapter include:

- Basic NGE Configuration Overview
- Configuring NGE Components
- NGE Management Tasks

2.16.1 Basic NGE Configuration Overview

Use the following steps to configure NGE for an MPLS service or router interface. The steps must be performed in order.

Step 1. Configure the group encryption label. The label must be unique, and the same label must be used on all nodes in the network group.

Step 2. Create a key group, duplicating this configuration on all nodes participating in this key group.

i. Configure the encryption and authentication algorithms for the group.

ii. Configure a security association (SA) that contains the encryption and authentication keys.

iii. Configure the active outbound SA for the group.

Step 3. Select the SDPs, VPRN services, or router interfaces that require encryption.

i. For each SDP, VPRN service, or router interface, configure the outbound direction key group.

ii. For each SDP, VPRN service, or router interface, configure the inbound direction key group.

2.16.2 Configuring NGE Components

Use the CLI syntax below to configure the following NGE parameters:

- Configuring the Global Encryption Label
- Configuring a Key Group
- Assigning a Key Group to an SDP or VPRN Service

2.16.2.1 Configuring the Global Encryption Label

The global encryption label is the network-wide, unique MPLS encryption label used for all nodes in the network group. The same encryption label must be configured on each node in the group.

Use the following CLI syntax to configure the global encryption label:

CLI Syntax:

```
config>group-encryption
     group-encryption-label encryption-label
```

The following example displays global encryption label usage:

Example:

```
config# group-encryption
config>grp-encryp# group-encryption-label 34
```

The following example displays the global encryption label configuration:

```
domain1>config>grp-encryp# info
-------------------------------------------------------
     group-encryption-label 34
-------------------------------------------------------
domain1>config>grp-encryp#
```

2.16.2.2 Configuring a Key Group

To configure a key group, set the following parameters:

- encryption and authentication algorithms
- security association
- active outbound SA

The authentication and encapsulation keys must contain the exact number of hexadecimal characters required by the algorithm used. For example, using sha256 requires 64 hexadecimal characters.

Keys are entered in clear text using the **security-association** command. Once entered, they are never displayed in their original, clear text form. Keys are displayed in an encrypted form, which is indicated by the system-appended **crypto** keyword when an **info** command is run. The NGE node also includes the **crypto** keyword with an **admin**>**save** operation so that the NGE node can decrypt the keys when reloading a configuration database. For security reasons, keys encrypted on one node are not usable on other nodes (that is, keys are not exchangeable between nodes).

Use the following CLI syntax to configure key group options:

CLI Syntax:

```
config# group-encryption
     encryption-keygroup keygroup-id [create]
          description description-string
          esp-auth-algorithm {sha256|sha512}
          esp-encryption-algorithm {aes128|aes256}
          keygroup-name keygroup-name
          security-association spi spi authentication-
            key authentication-key encryption-key
            encryption-key [crypto]
          active-outbound-sa spi
```

The following example displays key group command usage:

Example:

```
config>grp-encryp# encryption-keygroup KG1_secure
config>grp-encryp>encryp-keygrp# description
   Main_secure_KG
config>grp-encryp>encryp-keygrp# esp-auth-algorithm
  sha256
config>grp-encryp>encryp-keygrp# esp-encryption-
  algorithm aes128
config>grp-encryp>encryp-keygrp# keygroup-name
  KG1_secure
config>grp-encryp>encryp-keygrp# security-association
  spi 2 authentication-key
  0x88433A6DB4FA4F8A490EF661CBE69F010BFAE9C2784BED7059E5
  ADAAB1A225C6 encryption-key
  0x63DCDD501B66F85441E4A55B597DA617
config>grp-encryp>encryp-keygrp# security-association
  spi 6 authentication-key
  0x88433A6DB4FA4F8A490EF661CBE69F010BFAE9C2784BED7059E5
  ADAAB1A225C5 encryption-key
  0x63DCDD501B66F85441E4A55B597DA616
config>grp-encryp>encryp-keygrp# active-outbound-sa 6 ]
```

The following example displays the key group configuration:

```
domain1>config>grp-encryp# info detail
----------------------------------------------
        group-encryption-label 34
        encryption-keygroup 2 create
            description "Main_secure_KG"
            keygroup-name "KG1_secure"
            esp-auth-algorithm sha256
            esp-encryption-algorithm aes128
            security-association spi 2 authentication-
key 0x78d9e66a6669bd17454fe3184 ee161315b67adb8912949ceda20b6b741eb63604abe17de478e2
4723a7d1d5f7b6ffafc encryption-
key 0x8d51db8f826239f672457442cecc73665f52cbe00aedfb4eda6166001247b4eb crypto
            security-association spi 6 authentication-key 0x7fb9fc5553630924ee29973f
7b0a48f801b0ae1cb38b7666045274476a268e8d694ab6aa7ea050b7a43cdf8d80977625 encryption-
key 0x72bd9b87841dbebcb2d114031367ab5d9153a41b7c79c8f889ac56b950d8fffa crypto
            active-outbound-sa 6
        exit
----------------------------------------------
domain1>config>grp-encryp#
```

2.16.2.3 Assigning a Key Group to an SDP or VPRN Service

A key group can be assigned to the following entities:

- SDPs
- VPRNs

NGE supports encrypting the following services when key groups are assigned to an SDP or VPRN service:

- VLL services (Epipe or BGP-VPWS)
- VPRN service using Layer 3 spoke-SDP termination
- IES service using Layer 3 spoke-SDP termination
- VPLS service using spoke and mesh SDPs
- BGP-VPLS
- routed VPLS service into a VPRN or IES
- MP-BGP based VPRNs

For services that use SDPs, all tunnels may be either MPLS LSPs (RSVP-TE, LDP, or static LSP), or GRE or MPLSoUDP tunnels.

For MP-BGP services, **auto-bind-tunnel** is supported using LDP, GRE, MPLSoUDP, RSVP-TE, or MPLS (LDP or RSVP-TE).

Use the following CLI syntax to assign a key group to an SDP or a VPRN service:

CLI Syntax:

```
config>service# sdp sdp-id [create]
     encryption-keygroup keygroup-id direction {inbound
        | outbound}
```

CLI Syntax:

```
config>service# vprn service-id
     encryption-keygroup keygroup-id direction {inbound
        | outbound}
```

The following examples display a key group assigned to an SDP or a VPRN service:

Example:

```
config>service# sdp 61 create
config>service>sdp# encryption-keygroup 4 direction
 inbound
config>service>sdp# encryption-keygroup 4 direction
 outbound
```

Example:

```
config>service# vprn 22
config>service>vprn# encryption-keygroup 2 direction
 inbound
config>service>vprn# encryption-keygroup 2 direction
 outbound
```

The following example displays key group configuration for an SDP or a VPRN service.

```
domain1>config>service# info
----------------------------------------------
...
        sdp 61 create
            shutdown
            far-end 10.10.10.10
            exit
            encryption-keygroup 4 direction inbound
            encryption-keygroup 4 direction outbound
        exit
...
        vprn 22 customer 1 create
            shutdown
            encryption-keygroup 2 direction inbound
            encryption-keygroup 2 direction outbound
        exit
...
----------------------------------------------
```

2.17 Global Service Entity Management Tasks

This section discusses global service entity management tasks.

2.17.1 Modifying Customer Accounts

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

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

To edit customer information, such as description, contact, phone, enable the parameter and then enter the new information.

CLI Syntax:

```
config>service# customer customer-id [create] contact
  contact-information
description description-string
multi-service-site customer-site-name [create]
    assignment {port port-id | card slot}
    description description-string
    egress
        agg-rate
            burst-limit size [bytes|kilobytes]
            limit-unused-bandwidth
            queue-frame-based-accounting
            rate kilobits-per-second
        policer-control-policy name
        scheduler-override
            scheduler scheduler-name [create]
                parent {[weight weight]
                        [cir-weight cir-weight]}
                rate pir-rate [cir cir-rate]
        scheduler-policy scheduler-policy-name
    ingress
        policer-control-policy name
        scheduler-override
            scheduler scheduler-name [create]
                parent {[weight weight]
                        [cir-weight cir-weight]}
                rate pir-rate [cir cir-rate]
        scheduler-policy scheduler-policy-name
phone phone-number
```

Example:

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

2.17.2 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
```

2.17.3 Modifying SDPs

To access a specific SDP, 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, the SDP encapsulation type cannot be modified.

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
```

2.17.4 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
```

2.18 NGE Management Tasks

This section discusses the following NGE management tasks:

- Modifying a Key Group
- Removing a Key Group
- Changing Key Groups
- Deleting a Key Group from an NGE Node

2.18.1 Modifying a Key Group

When modifying a key group, observe the following conditions.

- The encryption or authentication algorithm for a key group cannot be changed if there are any SAs in the key group.
- The active outgoing SA must be removed (deconfigured) before the SPI can be deleted from the SA list in the key group.
- Before the outgoing SA can be deconfigured, the key group must be removed from all services on the node that use the key group

In the following example, the active outgoing SA is deconfigured, the SAs are removed, and the encryption algorithm is changed. Then the SAs are reconfigured, followed by reconfiguration of the active outgoing SA. The output display shows the new configuration based on those shown in Configuring a Key Group.

Use the following CLI syntax to modify a key group. The first syntax deconfigures the key-group items and the second syntax reconfigures them.

CLI Syntax:

```
config# group-encryption
     encryption-keygroup keygroup-id
          no active-outbound-sa
          no security-association spi spi
     exit
```

CLI Syntax:

```
config# group-encryption
     encryption-keygroup keygroup-id
          security-association spi spi authentication-
            key auth-key encryption-key encrypt-key
          esp-encryption-algorithm {aes128|aes256}
     exit
```

Example:

```
config>grp-encryp# encryption-keygroup KG1_secure
config>grp-encryp>encryp-keygrp# no active-outbound-sa
config>grp-encryp>encryp-keygrp# no security-association
 spi 2
config>grp-encryp>encryp-keygrp# no security-association
 spi 6
```

Example:

```
config>grp-encryp# encryption-keygroup KG1_secure
config>grp-encryp>encryp-keygrp# esp-encryption-
 algorithm aes256
config>grp-encryp>encryp-keygrp# security-association
 spi 2 authentication-key
 0x0123456789012345678901234567890123456789012345678901
 234567890123 encryption-key
 0x0123456789012345678901234567890123456789012345678901
 234567890123
config>grp-encryp>encryp-keygrp# security-association
 spi 6 authentication-key
 0x0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123
 456789ABCDEF encryption-key
 0x0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123
 456789ABCDEF [crypto]
config>grp-encryp>encryp-keygrp# active-outbound-sa 2
```

The following example displays the commands used to modify a key group. The first example deconfigures the key-group items and the second example reconfigures them. The encryption algorithm is changed from 128 to 256, the keys are changed, and the active outbound SA is changed to SPI 2.

```
domain1>config>grp-encryp# info detail
----------------------------------------------
        group-encryption-label 34
        encryption-keygroup 2 create
            description "Main_secure_KG"
            keygroup-name "KG1_secure"
            esp-auth-algorithm sha256
            esp-encryption-algorithm aes128
            no security-association spi 2
            no security-association spi 6
            no active-outbound-sa
        exit
----------------------------------------------
domain1>config>grp-encryp#


domain1>config>grp-encryp# info detail
----------------------------------------------
        group-encryption-label 34
        encryption-keygroup 2 create
            description "Main_secure_KG"
            keygroup-name "KG1_secure"
            esp-auth-algorithm sha256
            esp-encryption-algorithm aes256
            security-association spi 2 authentication-
```

```
key 0x0123456789012345678901234567890123456789012345678901234567890123 encryption-
key 0x0123456789012345678901234567890123456789012345678901234567890123
          security-association spi 6 authentication-
key 0x0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF encryption-
key 0x0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF crypto
          active-outbound-sa 2
       exit
---------------------------------------------
domain1>config>grp-encryp#
```

2.18.2 Removing a Key Group

Both inbound and outbound direction key groups must be deconfigured before the key group can be removed (unbound). The inbound and outbound key groups must be deconfigured individually. Specifying a *keygroup-id* is optional.

2.18.2.1 Removing a Key Group from an SDP or VPRN Service

Use the following CLI syntax to remove a key group from an SDP or a VPRN service:

CLI Syntax:

```
config>service# sdp sdp-id
     no encryption-keygroup keygroup-id direction
        {inbound | outbound}
```

CLI Syntax:

```
config>service# vprn service-id
     no encryption-keygroup keygroup-id direction
        {inbound | outbound}
```

The following examples display a key group removed from an SDP or a VPRN service:

Example:

```
config>service# sdp 61
config>service>sdp# no encryption-keygroup 4 direction
 inbound
config>service>sdp# no encryption-keygroup 4 direction
 outbound
```

Example:

```
config>service# vprn 22
config>service>vprn# no encryption-keygroup 2 direction
 inbound
config>service>vprn# no encryption-keygroup 2 direction
 outbound
```

The following example shows that the key group configuration has been removed from an SDP or a VPRN service.

```
domain1>config>service# info
----------------------------------------------
...
        sdp 61 create
            shutdown
            far-end 10.10.10.10
            exit
        exit
...
...
        vprn 22 customer 1 create
            shutdown
        exit
...
----------------------------------------------
domain1>config>service# info
```

2.18.3 Changing Key Groups

Use the following sequence of CLI commands to change key groups:

1. Remove the inbound direction key group.
2. Change the outbound direction key group.
3. Install the new inbound direction key group.

2.18.3.1 Changing the Key Group for an SDP or VPRN Service

Changing key groups for an SDP or VPRN service must be performed on all nodes for the service.

The following CLI syntax changes the key group on an SDP. The syntax for a VPRN service is similar. In the example below, the inbound and outbound key groups are changed from key group 4 to key group 6.

CLI Syntax:

```
config>service# sdp sdp-id
      no encryption-keygroup keygroup-id direction
         {inbound|outbound}
```

Example:

```
config>service# sdp 61
config>service>sdp# no encryption-keygroup 4 direction
 inbound
```

```
config>service>sdp# encryption-keygroup 6 direction
 outbound
config>service>sdp# encryption-keygroup 6 direction
 inbound
```

The following example shows that the key group configuration has been changed for the SDP or the VPRN service.

```
domain1>config>service# info
----------------------------------------------
...
        sdp 61 create
            shutdown
            far-end 10.10.10.10
            exit
            encryption-keygroup 6 direction inbound
            encryption-keygroup 6 direction outbound
        exit
...
...
        vprn 22 customer 1 create
            shutdown
            encryption-keygroup 2 direction inbound
            encryption-keygroup 2 direction outbound
        exit
...
----------------------------------------------
domain1>config>service# info
```

2.18.4 Deleting a Key Group from an NGE Node

To delete a key group from an NGE node, the key group must be removed (unbound) from all SDPs, VPRN services, and router interfaces that use it.

To locate the key group bindings, use the CLI command **show>group-encryption> encryption-keygroup** *keygroup-id***.**

Use the following CLI syntax to delete a key group:

CLI Syntax:

```
config# group-encryption
      no encryption-keygroup keygroup-id
```

Example:

```
config>grp-encryp# no encryption-keygroup 8
```

2.19 Global Services Configuration Command Reference

This section provides the Global Services configuration command reference.

Topics include:

- Command Hierarchies
- Command Descriptions

2.19.1 Command Hierarchies

- Customer Commands
- MRP Commands
- Service System Commands
- Oper Group Commands
- Pseudowire (PW) Commands
- SDP Commands
- SAP Commands
- Ethernet Ring Commands
- ETH CFM Configuration Commands
- ETH Tunnel Commands
- Connection Profile VLAN Commands
- Network Group Encryption (NGE) Commands
- NGE Services Commands
- Model-Driven Automatic ID Commands

2.19.1.1 Customer Commands

```
config
    — service
          — customer customer-id [create] [name name]
          — no customer customer-id
                — contact contact-information
                — no contact contact-information
                — description description-string
                — no description
— multi-service-site customer-site-name [create]
— no multi-service-site customer-site-name
    — assignment {port port-id | card slot-number}
    — no assignment
    — description description-string
    — no description
    — egress
        — [no] agg-rate
            — [no] limit-unused-bandwidth
            — [no] queue-frame-based-accounting
            — rate {kilobits-per-second}
            — no rate
        — policer-control-policy policy-name
        — no policer-control-policy
        — [no] scheduler-override
            — scheduler scheduler-name [create]
            — no scheduler scheduler-name
                — parent [weight weight] [cir-weight cir-weight]
                — no parent
                — rate pir-rate [cir cir-rate]
                — no rate
        — scheduler-policy scheduler-policy-name
        — no scheduler-policy
    — ingress
        — policer-control-policy policy-name
        — no policer-control-policy
        — [no] scheduler-override
            — scheduler scheduler-name [create]
            — no scheduler scheduler-name
                — parent [weight weight] [cir-weight cir-weight]
                — no parent
                — rate pir-rate [cir cir-rate]
                — no rate
        — scheduler-policy scheduler-policy-name
        — no scheduler-policy
— [no] phone phone-number
```

2.19.1.2 MRP Commands

```
config
    — service
        — mrp
            — copy mrp-policy src-mrp-policy to dst-mrp-policy
            — mrp-policy policy-name [create]
            — no mrp-policy policy-name
                — default-action {block | allow}
                — no default-action
                — description description-string
                — no description
                — entry entry-id [create]
                — no entry entry-id
```

```
                    — action {none | block | allow | end-station}
                    — no action
                    — description description-string
                    — no description
                    — [no] match
                        — isid value [to higher-value]
                        — no isid
                        — no isid value [to higher-value]
                — renum src-entry-id to dst-entry-id
                — scope {exclusive | template}
                — no scope
```

2.19.1.3 Service System Commands

```
config
    — service
        — system
            — bgp-auto-rd-range ip-addr comm-val range to range
            — no bgp-auto-rd-range
            — gre-eth-bridged
                — [no] tunnel-termination [ip-address | ipv6-address] fpe fpe-id [create]
            — vpn-gre-source-ip ip-address
            — no vpn-gre-source-ip
```

2.19.1.4 Oper Group Commands

```
config
    — service
        — oper-group group-name [create]
        — no oper-group group-name
            — bfd-enable interface interface-name dest-ip ipv4-address [service service-id]
            — no bfd-enable
            — hold-time
                — group up time | no group up
                — group down time | no group down
```

```
config
    — service
        — ies service-id (Refer to the 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services
            Guide: IES and VPRN)
            — [no] interface ip-int-name
                — monitor-oper-group name
                — no monitor-oper-group name
```

```
config
    — service
        — vpls service-id  (Refer to the 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2
            Services and EVPN Guide: VLL, VPLS, PBB, and EVPN)
        — [no] interface ip-int-name
                    — monitor-oper-group name
                    — no monitor-oper-group

config
    — service
        — vprn service-id  (Refer to the 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3
              Services Guide: IES and VPRN)
            — site name [create]
                — monitor-oper-group name
                — no monitor-oper-group name
```

2.19.1.5 Pseudowire (PW) Commands

```
config
    — service
        — pw-routing
            — boot-timer secs
            — no boot-timer
            — local-prefix local-prefix [create]
            — no local-prefix local-prefix
                — advertise-bgp route-distinguisher rd [community community]
                — no advertise-bgp route-distinguisher rd
            — path name [create]
            — no path name
                — hop hop-index ip-address
                — no hop hop-index
                — [no] shutdown
            — retry-count [count]
            — no retry-count
            — retry-timer secs
            — no retry-timer
            — spe-address global-id:prefix
            — no spe-address
            — [no] static-route route-name

config
    — service
        — pw-template policy-id [use-provisioned-sdp | [prefer-provisioned-sdp] [auto-gre-
              sdp] ] [create] [name name]
        — no pw-template policy-id
            — accounting-policy acct-policy-id
            — no accounting-policy
            — [no] allow-fragmentation
            — [no] auto-learn-mac-protect
            — [no] block-on-peer-fault
            — [no] collect-stats
            — [no] controlword
            — [no] disable-aging
            — [no] disable-learning
            — [no] discard-unknown-source
```

- **egress**
  - **filter ipv6** *ipv6-filter-id*
  - **filter ip** *ip-filter-id*
  - **filter mac** *mac-filter-id*
  - **no filter** [**ip** *ip-filter-id*] [**mac** *mac-filter-id*] [**ipv6** *ipv6-filter-id*]
  - **filter-name ipv6** *ipv6-name*
  - **filter-name ip** *ip-name*
  - **filter mac** *mac-name*
  - **no filter** [**ip**] [**ipv6**] [**mac**]
  - **mfib-allowed-mda-destinations**
    - [**no**] **mda** *mda-id*
  - **qos** *network-policy-id* **port-redirect-group** *queue-group-name* **instance** *instance-id*
  - **qos name** *network-policy-name* **port-redirect-group** *queue-group-name* **instance** *instance-id*
  - **no qos** [*network-policy-id*]
- [**no**] **entropy-label**
- [**no**] **force-qinq-vc-forwarding**
- [**no**] **force-vlan-vc-forwarding**
- **hash-label signal-capability**
- **no hash-label**
- **igmp-snooping**
  - [**no**] **fast-leave**
  - **import** *policy-name*
  - **no import**
  - **last-member-query-interval** *interval*
  - **no last-member-query-interval**
  - **max-num-groups** *count*
  - **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**
- **ingress**
  - **filter ipv6** *ipv6-filter-id*
  - **filter ip** *ip-filter-id*
  - **filter mac** *mac-filter-id*
  - **no filter** [**ip** *ip-filter-id*] [**mac** *mac-filter-id*] [**ipv6** *ipv6-filter-id*]
  - **filter-name ipv6** *ipv6-name*
  - **filter-name ip** *ip-name*
  - **filter mac** *mac-name*
  - **no filter** [**ip**] [**ipv6**] [**mac**]
  - **qos** *network-policy-id* **fp-redirect-group** *queue-group-name* **instance** *instance-id*
  - **qos name** *network-policy-name* **fp-redirect-group** *queue-group-name* **instance** *instance-id*
  - **no qos** [*network-policy-id*]
- **l2pt-termination** [**cdp**] [**dtp**] [**pagp**] [**stp**] [**udld**] [**vtp**]
- **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
            — restrict-protected-src [alarm only | discard-frame]
            — no restrict-protected-src
            — [no] sdp-exclude group-name
            — [no] sdp-include group-name
            — split-horizon-group group-name [residential-group]
            — no split-horizon-group
                — [no] auto-learn-mac-protect
                — description description-string
                — no description
                — restrict-protected-src [alarm-only | discard-frame]
                — no restrict-protected-src
                — [no] restrict-unprotected-dst
            — stp
                — [no] auto-edge
                — [no] edge-port
                — link-type {pt-pt | shared}
                — no link-type
                — path-cost sap-path-cost
                — no path-cost
                — priority stp-priority
                — no priority
                — [no] root-guard
                — [no] shutdown
            — vc-type {ether | vlan}
            — vlan-vc-tag 0..4094
            — no vlan-vc-tag
```

2.19.1.5.1 PW Port Commands

```
config
    — service
        — sdp sdp-id [delivery-type] [create]
        — no sdp sdp-id
            — binding
                — port [port-id | lag-id]
                — no port
                — pw-port pw-port-id [vc-id vc-id] [create]
                — no pw-port pw-port-id
                    — egress
                        — [no] shaper
                            — int-dest-id int-dest-id
                            — no int-dest-id
                            — pw-sap-secondary-shaper pw-sap-secondary-
                                  shaper-name
                            — no pw-sap-secondary-shaper
                            — vport vport-name
                            — no vport
                        — vc-label vc-label
                        — no vc-label
                    — ingress
                        — vc-label vc-label
                        — no vc-label
                    — monitor-oper-group group name
                    — no monitor-oper-group
                    — [no] shutdown
                    — vc-type {ether | vlan}
                    — no vc-type
                    — vlan-vc-tag vlan-id
                    — no vlan-vc-tag
```

Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for command syntax and CLI command descriptions for the following VLL PW port commands.

```
config
    — service
        — [no] epipe service-id [customer customer-id] [test] [create] [vpn vpn-id] [vc-
            switching]
            — pw-port pw-port-id fpe fpe-id [create]
            — no pw-port
                — egress
                    — [no] shaper
                        — int-dest-id name
                        — no int-dest-id
                        — vport vport
                        — no vport
                    — monitor-oper-group group-name
                    — no monitor-oper-group
                    — [no] shutdown
```

2.19.1.6 SDP Commands

```
config
    — service
        — sdp sdp-id [delivery-type] [create]
        — no sdp sdp-id
            — accounting-policy acct-policy-id
            — no accounting-policy
            — [no] adv-mtu-override
            — [no] allow-fragmentation
            — [no] bgp-tunnel
            — booking-factor percentage
            — no booking-factor
            — class-forwarding [default-lsp lsp-name]
            — no class-forwarding
                — [no] enforce-diffserv-lsp-fc
```

```
                    — fc {fc} lsp lsp-name
                    — no fc {fc}
                    — multicast-lsp lsp-name
                    — no multicast-lsp
                    — [no] shutdown
          — [no] collect-stats
          — description description-string
          — no description
          — far-end node-id node-id [global-id global-id]
          — far-end [ip-address | ipv6-address]
          — no far-end ip-address | ipv6-address
          — keep-alive
                    — hello-time seconds
                    — no hello-time [seconds]
                    — hold-down-time seconds
                    — no hold-down-time [seconds]
                    — max-drop-count count
                    — no max-drop-count [count]
                    — message-length message-length
                    — no message-length [message-length]
                    — [no] shutdown
                    — timeout timeout
                    — no timeout
          — [no] ldp
          — local-end {ip-address | ipv6-address}
          — no local-end
          — [no] lsp lsp-name
          — metric metric
          — no metric
          — [no] mixed-lsp-mode
                    — revert-time {revert-time | infinite}
                    — no revert-time
          — network-domain network-domain-name
          — no network-domain
          — path-mtu bytes
          — no path-mtu [bytes]
          — pbb-etype type
          — no pbb-etype [type]
          — [no] shutdown
          — signaling [off | tldp | bgp]
          — source-bmac-lsb mac-lsb control-pw-vc-id vc-id
          — no source-bmac-lsb
          — [no] sr-isis
          — [no] sr-ospf
          — [no] sr-te-lsp lsp-name
          — tunnel-far-end ip-address | ipv6-address
          — no tunnel-far-end [ip-address | ipv6-address]
          — vlan-vc-etype ethernet-type
          — no vlan-vc-etype [ethernet-type]
          — [no] weighted-ecmp
— sdp-group
          — group-name group-name value group-value
          — no group-name group-name
```

2.19.1.7 SAP Commands

```
config
    — service
        — apipe (See the 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN
              Guide: VLL, VPLS, PBB, and EVPN)
            — sap sap-id [create] [no-endpoint]
            — sap sap-id [create] endpoint endpoint-name
            — no sap sap-id
        — epipe (See the 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN
              Guide: VLL, VPLS, PBB, and EVPN)
            — sap sap-id [create] [no-endpoint]
            — sap sap-id [create] endpoint endpoint-name
            — no sap sap-id
        — fpipe (See the 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN
              Guide: VLL, VPLS, PBB, and EVPN)
            — sap sap-id [create] [no-endpoint]
            — sap sap-id [create] endpoint endpoint-name
            — no sap sap-id
        — ies (See the 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES
              and VPRN)
            — sap sap-id [create]
            — no sap sap-id
        — ipipe (See the 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN
              Guide: VLL, VPLS, PBB, and EVPN)
            — sap sap-id [create] [no-endpoint]
            — sap sap-id [create] endpoint endpoint-name
            — no sap sap-id
        — vpls (See the 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN
              Guide: VLL, VPLS, PBB, and EVPN)
            — sap sap-id [split-horizon-group group-name] [create]
            — no sap sap-id
        — vprn (Refer to the 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide:
              IES and VPRN)
            — sap sap-id [create]
            — no sap sap-id
    — system
        — ethernet
            — [no] new-qinq-untagged-sap
```

2.19.1.8 Ethernet Ring Commands

```
config
    — [no] eth-ring ring-index
        — ccm-hold-time [down down-timeout] [up up-timeout]
        — no ccm-hold-time
        — compatible-version version
        — no compatible-version
        — description description-string
        — no description
        — guard-time time
        — no guard-time
        — node-id mac-address
        — no node-id
        — path {a | b} [{port-id | lag-id} raps-tag qtag1[.qtag2]]
        — no path {a | b}
            — description long-description-string
            — no description
            — eth-cfm
                — [no] mep mep-id domain md-index association ma-index
                    — alarm-notification
                        — fng-alarm-time time
                        — fng-reset-time time
                    — [no] ccm-enable
                    — ccm-ltm-priority priority
                    — no ccm-ltm-priority
                    — ccm-padding-size ccm-padding
                    — no ccm-padding-size
                    — [no] control-mep
                    — 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
                    — grace
                        — eth-ed
                            — max-rx-defect-window seconds
                            — no max-rx-defect-window
                            — priority priority
                            — no priority
                            — [no] rx-eth-ed
    — [no] tx-eth-ed
                            — eth-vsm-grace
    — [no] rx-eth-vsm-grace
    — [no] tx-eth-vsm-grace
                    — low-priority-defect {allDef | macRemErrXcon | remErrXcon |
                          errXcon | xcon | noXcon}
                    — mac-address mac-address
                    — no mac-address
                    — one-way-delay-threshold seconds
                    — [no] shutdown
            — [no] rpl-end
            — [no] shutdown
        — revert-time time
        — no revert-time
        — rpl-node {owner | nbr}
        — no rpl-node
        — [no] shutdown
        — [no] sub-ring {virtual-link | non-virtual-link}
            — [no] interconnect {ring-id ring-index | vpls}
                — [no] propagate-topology-change
```

2.19.1.9 ETH CFM Configuration Commands

```
config
    — eth-cfm
        — default-domain
            — bridge-identifier bridge-id vlan vlan-id
                — id-permission [chassis | defer]
                — no id-permission
                — mhf-creation [none | default | explicit | defer] level level
                — mip-ltr-priority priority
        — domain md-index [format {format}] [name md-name] level level [admin-name admin-
            name]
        — domain md-index
        — no domain md-index
            — association ma-index [format {format}] name ma-name [admin-name admin-
                name]
            — association ma-index
            — no association ma-index
                — [no] auto-mep-discovery
    — [no] bridge-identifier [bridge-id | bridge-name bridge-name]
                    — id-permission {chassis}
                    — no id-permission
                    — mhf-creation {none | default | explicit | static} level level
                    — no mhf-creation
                    — mip-ltr-priority priority
                    — vlan vlan-id
                    — no vlan
                — ccm-hold-time down timer
                — no ccm-hold-time
                — ccm-interval interval
                — no ccm-interval
                — facility-id-permission {chassis}
                — no facility-id-permission
                — remote-mepid mep-id remote-mac {unicast-da | default}
                — no remote-mepid mep-id
        — redundancy
            — mc-lag
                — propagate-hold-time seconds
                — no propagate-hold-time
    — [no] standby-mep-shutdown
        — slm
            — inactivity-timer timer
            — no inactivity-timer
        — system
            — [no] grace-tx-enable
            — sender-id local local-name
            — sender-id system
            — no sender-id
```

2.19.1.10 ETH Tunnel Commands

```
config
    — [no] eth-tunnel tunnel-index
        — ccm-hold-time [down down-timeout] [up up-timeout]
        — no ccm-hold-time
        — description long-description-string
        — no description
        — ethernet
            — encap-type {dot1q | qinq}
            — no encap-type
            — mac ieee-address
            — no mac
        — lag-emulation
            — access
                — adapt-qos {distribute | link | port-fair}
                — no adapt-qos
                — [no] per-fp-ing-queuing
            — path-threshold num-paths
            — no path-threshold
        — [no] path path-index
            — control-tag qtag[.qtag]
            — no control-tag
            — description description-string
            — no description
            — eth-cfm
                — [no] mep mep-id domain md-index association ma-index
                    — alarm-notification
                        — fng-alarm-time time
                        — fng-reset-time time
                    — [no] ccm-enable
                    — ccm-ltm-priority priority
                    — no ccm-ltm-priority
                    — ccm-padding-size ccm-padding
                    — no ccm-padding-size
                    — [no] control-mep
                    — 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
                    — grace
                        — eth-ed
                            — max-rx-defect-window seconds
                            — no max-rx-defect-window
                            — priority priority
                            — no priority
                            — [no] rx-eth-ed
    — [no] tx-eth-ed
                            — eth-vsm-grace
    — [no] rx-eth-vsm-grace
    — [no] tx-eth-vsm-grace
```

```
                    — low-priority-defect {allDef | macRemErrXcon | remErrXcon |
                          errXcon | xcon | noXcon}
                    — mac-address mac-address
                    — no mac-address
                    — one-way-delay-threshold seconds
                    — [no] shutdown
                — member port-id
                — no member
                — precedence {primary | secondary}
                — [no] shutdown
            — protection-type {g8031-1to1 | loadsharing}
            — revert-time time
            — no revert-time
            — [no] shutdown
```

2.19.1.11 Connection Profile VLAN Commands

```
config
    — connection-profile-vlan conn-prof-id [create]
    — no connection-profile-vlan conn-prof-id
        — description description-string
        — no description
        — vlan-range from [to to]
        — no vlan-range from
```

2.19.1.12 Network Group Encryption (NGE) Commands

```
config
    — group-encryption
        — encryption-keygroup keygroup-id [create]
        — no encryption-keygroup keygroup-id
            — active-outbound-sa spi
            — no active-outbound-sa
            — description description-string
            — no description
            — esp-auth-algorithm {sha256 | sha512}
            — no esp-auth-algorithm
            — esp-encryption-algorithm {aes128 | aes256}
            — no esp-encryption-algorithm
            — keygroup-name keygroup-name
            — no keygroup-name
            — security-association spi spi authentication-key authentication-key
                encryption-key encryption-key [crypto]
            — no security-association spi spi
        — group-encryption-label encryption-label
        — no group-encryption-label
```

2.19.1.13 NGE Services Commands

```
config
     — service
          — sdp
               — encryption-keygroup keygroup-id direction {inbound | outbound}
               — no encryption-keygroup direction {inbound | outbound}
          — vprn
               — encryption-keygroup keygroup-id direction {inbound | outbound}
               — no encryption-keygroup direction {inbound | outbound}
```

2.19.1.14 Model-Driven Automatic ID Commands

```
config
     — service
          — md-auto-id
               — customer-id-range start customer-id end customer-id
               — no customer-id-range
               — pw-template-id-range start pw-template-id end pw-template-id
               — no pw-template-id-range
               — service-id-range start service-id end service-id
               — no service-id-range
```

2.19.2 Command Descriptions

This section provides CLI command descriptions and output. The command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

Topics in this section include:

- Generic Commands
- Customer Management Commands
- Service System Commands
- MRP Commands
- Oper Group Commands
- Pseudowire (PW) Commands
- PW Port Commands
- SDP Commands
- Ethernet Ring Commands
- ETH CFM Configuration Commands

- ETH-Tunnel Commands
- Connection Profile VLAN Commands
- Network Group Encryption (NGE) Commands
- NGE Services Commands
- Model-Driven Automatic ID Commands

2.19.2.1 Generic Commands

description

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

Context config>service>cust
config>service>cust>multi-service-site
config>service>mrp>mrp-policy
config>service>mrp>mrp-policy>entry
config>service>pw-template
config>service>pw-template>split-horizon-group
config>service>sdp
config>eth-ring
config>eth-ring>path
config>eth-ring>path>eth-cfm>mep
config>eth-tunnel
config>eth-tunnel>path
config>eth-tunnel>path>eth-cfm>mep
config>connection-profile-vlan
config>grp-encryp>encryp-keygrp

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, and so on), the entire string must be enclosed within double quotes.

shutdown

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

Context config>eth-cf>mep
config>service>sdp
config>service>sdp>class-forwarding
config>service>sdp>keep-alive
config>service>pw-routing>hop
config>service>pw-template>stp
config>service>sdp>binding>pw-port
config>eth-ring>path
config>eth-tunnel
config>eth-tunnel>path
config>eth-tunnel>path>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.

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.

new-qinq-untagged-sap

Syntax [**no**] **new-qinq-untagged-sap**

Context config>system>ethernet

Description This command controls the behavior of QinQ SAP y.0 (for example, 1/1/1:3000.0). If the flag is not enabled (no new-qinq-untagged-sap), the y.0 SAP works the same as the y.* SAP (for example, 1/1/1:3000.*); all frames tagged with outer VLAN y and no inner VLANs or inner VLAN x where inner VLAN x is not specified in a SAP y.x configured on the same port (for example, 1/1/1:3000.10).

If the flag is enabled, then the following new behavior immediately applies to all existing and future y.0 SAPs: the y.0 SAP maps all the ingress frames tagged with outer tag VLAN-id of y (qinq-etype) and no inner tag or with inner tag of VLAN-id of zero (0). When the flag is disabled, there is no disruption for existing usage of this SAP type.

Default no new-qinq-untagged-sap.

2.19.2.2 Customer Management Commands

customer

Syntax **customer** *customer-id* [**create**] [**name** *name*]
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.

An optional customer **name** can be specified and is tied to the **customer-name** in the customer context (setting either **customer-name** or **name** will cause the other to change as well).

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.

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

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

Values *customer-id*: 1 to 2147483647
customer-name: 64 characters maximum

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.

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

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

If a name is not specified at creation time, then SR OS assigns a string version of the *customer-id* as the name.

Values *name*: 64 characters maximum

contact

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

Context config>service>cust

Description This command configures contact information for a customer.

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

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

Default no contact

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

multi-service-site

Syntax **multi-service-site** *customer-site-name* [**create**]
no multi-service-site *customer-site-name*

Context config>service>cust

Description This command creates a new customer site or edits an existing customer site with the *customer-site-name* parameter. A customer site is an anchor point to create an ingress and egress virtual scheduler hierarchy. When a site is created, it must be assigned to a chassis slot or port with the exception of the 7450 ESS-1 in which the slot is set to 1. When scheduler policies are defined for ingress and egress, the scheduler names contained in each policy are created according to the parameters defined in the policy. Multi-service customer sites exist for the sole purpose of creating a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs).

The scheduler policy association with the customer site normally prevents the scheduler policy from being deleted until after the scheduler policy is removed from the customer site. The multi-service-site object will generate a log message indicating that the association was deleted due to scheduler policy removal.

When the multi-service customer site is created, an ingress and egress scheduler policy association does not exist. This does not prevent the site from being assigned to a chassis slot or prevent service SAP assignment. After the site has been created, the ingress and egress scheduler policy associations can be assigned or removed at any time.

Parameters *customer-site-name* — Specifies the customer site name. Each customer site must have a unique name within the context of the customer. If *customer-site-name* already exists for the customer ID, the CLI context changes to that site name for the purpose of editing the site scheduler policies or assignment. Any modifications made to an existing site will affect all SAPs associated with the site. Changing a scheduler policy association may cause new schedulers to be created and existing policers and queues on the SAPs to no longer be orphaned. Existing schedulers on the site may cease to exist, causing policers and queues relying on that scheduler to be orphaned.

If the *customer-site-name* does not exist, it is assumed that an attempt is being made to create a site of that name in the customer ID context. The success of the command execution depends on the following:

- The maximum number of customer sites defined for the chassis has not been met.
- The *customer-site-name* is valid.
- The **create** keyword is included in the command line syntax (if the system requires).

When the maximum number of customer sites has been exceeded a configuration error occurs; the command will not execute and the CLI context will not change.

If the *customer-site-name* is invalid, a syntax error occurs; the command will not execute and the CLI context will not change.

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

assignment

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

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

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

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

Default None

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

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

Values For the 7950 XRS:

port-id	slot/mda/port [.channel]		
	eth-sat-id	esat-id/slot/port	
		esat	keyword
		id: 1 to 20	
	pxc-id	psc-id.sub-port	
		pxc psc-id.sub-port	
		pxc	keyword
		id: 1 to 64	

	sub-port: a, b	
lag		keyword
id	1 to 800	1 to 800

For the 7750 SR and the 7450 ESS:

port-id	*slot*/*mda*/*port*[.*channel*]		
	pxc-id	psc-id.sub-port	
		pxc psc-id.sub-port	
		pxc	keyword
		id: 1 to 64	
		sub-port: a, b	
	aps-id	aps-*group-id*[.*channel*]	
		aps keyword	
		group-id	1 to 64
		group-id	1 to 16
	bundle-*type-slot/mda*.*bundle-num*		
		bundle	keyword
		type	ima, ppp
		bundle-num	1 to 256
	bpgrp-id:	bpgrp-*type-bpgrp-num*	
		bpgrp	keyword
		type	ima
		bpgrp-num	1 to 1280
	ccag-id	- ccag-<id>.<path-id>[cc-type]	
		ccag	keyword
		id	1 to 8
		path-id	a, b
		cc-type[.sap-net | .net-sap]	
	lag-id	lag-*id*	
		lag	keyword
		id	1 to 800

slot-number — Assigns the multi-service customer site to the slot-number given. When the multi-service customer site is assigned to a specific slot in the chassis, all SAPs associated with this customer site must be on a service owned by the customer and created on the defined chassis slot. The defined slot must already have been pre-provisioned on the system but need not be installed when the customer site assignment is made.

Values Any pre-provisioned slot number for the chassis type that allows SAP creation.

1 to 10

egress

Syntax **egress**

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

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

ingress

Syntax **ingress**

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

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

agg-rate

Syntax [**no**] **agg-rate**

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

Description This command enables the context to control an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: **rate**, **limit-unused-bandwidth**, **and queue-frame-based-accounting**.

The **no** form of the command disables the aggregate rate limit parameters.

limit-unused-bandwidth

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

Context config>service>cust>multi-service-site>egress>agg-rate

Description This command is used to enable aggregate rate overrun protection.

The **no** form of the command disables aggregate rate overrun protection.

Default no limit-unused-bandwidth

queue-frame-based-accounting

Syntax [**no**] **queue-frame-based-accounting**

Context config>service>cust>multi-service-site>egress>agg-rate

Description This command enables frame based accounting on all policers and queues associated with the agg-rate context. Only supported on Ethernet ports. Not supported on HSMDA Ethernet ports. Packet byte offset settings are not included in the applied rate when queue frame based accounting is configured, however the offsets are applied to the statistics.

The **no** form of the command disables frame based accounting.

Default no queue-frame-based-accounting

rate

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

Context config>service>cust>multi-service-site>egress>agg-rate

Description This command defines the enforced aggregate rate for all queues associated with the agg-rate context. A rate must be specified for the agg-rate context to be considered to be active on the context’s object (SAP, subscriber, Vport, and so on).

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

Default no rate

policer-control-policy

Syntax **policer-control-policy** *policy-name*
no policer-control-policy

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

Description This command, within the QoS CLI node, is used to create, delete or modify policer control policies. A policer control policy is very similar to the scheduler-policy which is used to manage a set of queues by defining a hierarchy of virtual schedulers and specifying how the virtual schedulers interact to provide an aggregate SLA. In a similar fashion, the policer-control-policy controls the aggregate bandwidth available to a set of child policers. Once created, the policy can be applied to ingress or egress SAPs.

Policer Control Policy Instances

On the SAP side, an instance of a policy is created each time a policy is applied. When applied to a 7750 SR or 7450 ESS sub-profile, an instance of the policy is created each time a subscriber successfully maps one or more hosts to the profile per ingress SAP.

Each instance of the policer-control-policy manages the policers associated with the object that owns the policy instance (SAP or subscriber). If a policer on the object is parented to an appropriate arbiter name that exists within the policy, the policer will be managed by the instance. If a policer is not parented or is parented to a non-existent arbiter, the policer will be orphaned and not subject to bandwidth control by the policy instance.

Maximum Rate and Root Arbiter

The policer-control-policy supports an overall maximum rate (max-rate) that defines the total amount of bandwidth that may be distributed to all associated child policers. By default, that rate is set to max which provides an unlimited amount of bandwidth to the policers. Once the policy is created, an actual rate should be configured in order for the policy instances to be effective. At the SAP level, the maximum rate may be overridden on a per instance basis. For 7750 SR or 7450 ESS subscribers, the maximum rate may only be overridden on the subscriber profile which will then be applied to all instances associated with the profile.

The maximum rate is defined within the context of the root arbiter which is always present in a policer-control-policy. The system creates a parent policer which polices the output of all child policers attached to the policy instance to the configured rate. Child policers may be parented directly to the root arbiter (parent root) or parented to one of the tiered arbiters (parent arbiter-name). Since each tiered arbiter must be parented to either another tiered arbiter or the root arbiter (default), every parented child policer is associated with the root arbiter and thus the root arbiter’s parent policer.

Parent Policer PIR Leaky Bucket Operation

The parent policer is a single leaky bucket that monitors the aggregate throughput rate of the associated child policers. Forwarded packets increment the bucket by the size of each packet. The rate of the parent policer is implemented as a bucket decrement function which attempts to drain the bucket. If the rate of the packets flowing through the bucket is less than the decrement rate, the bucket does not accumulate depth. Each packet that flows through the bucket is accompanied by a derived discard threshold. If the current depth of the bucket is less than the discard threshold, the packet is allowed to pass through, retaining the colors derived from the packet's child policer. If the current depth is equal to or greater than the threshold value, the packet is colored red and the bucket depth is not incremented by the packet size. Also, any increased bucket depths in the child policer are canceled making any discard event an atomic function between the child and the parent.

Due to the fact that multiple thresholds are supported by the parent policer, the policer control policy is able to protect the throughput of higher priority child policers from the throughput of the lower priority child policers within the aggregate rate.

Tier 1 and Tier 2 Arbiters

As stated above, each child is attached either to the always available root arbiter or to an explicitly created tier 1 or tier 2 arbiter. Unlike the hardware parent policer based root arbiter, the arbiters at tier 1 and tier 2 are only represented in software and are meant to provide an arbitrary hierarchical bandwidth distribution capability. An arbiter created on tier 2 must parent to either to an arbiter on tier 1 or to the root arbiter. Arbiters created on tier 1 always parent to the root arbiter. In this manner, every arbiter ultimately is parented or grand-parented by the root arbiter.

Each tiered arbiter supports an optional rate parameter that defines a rate limit for all child arbiters or child policers associated with the arbiter. Child arbiters and policers attached to the arbiter have a level attribute that defines the strict level at which the child is given bandwidth by the arbiter. Level 8 is the highest and 1 is the lowest. Also a weight attribute defines each child's weight at that strict level in order to determine how bandwidth is distributed to multiple children at that level when insufficient bandwidth is available to meet each child's required bandwidth.

Fair and Unfair Bandwidth Control

Each child policer supports three leaky buckets. The PIR bucket manages the policer's peak rate and maximum burst size, the CIR leaky bucket manages the policer's committed rate and committed burst size. The third leaky bucket is used by the policer control policy instance to manage the child policer's fair rate (FIR). When multiple child policers are attached to the root arbiter at the same priority level, the policy instance uses each child's FIR bucket rate to control how much of the traffic forwarded by the policer is fair and how much is unfair.

In the simplest case where all the child policers in the same priority level are directly attached to the root arbiter, each child's FIR rate is set according to the child's weight divided by the sum of the active children's weights multiplied by the available bandwidth at the priority level. The result is that the FIR bucket will mark the appropriate amount of traffic for each child as fair based on the weighted fair output of the policy instance.

The fair/unfair forwarding control in the root parent policer is accomplished by implementing two different discard thresholds for the priority. The first threshold is discard-unfair and the second is discard-all for packet associated with the priority level. As the parent policer PIR bucket fills (due the aggregate forwarded rate being greater than the parent policers PIR decrement rate) and the bucket depth reaches the first threshold, all unfair packets within the priority are discarded. This leaves room in the bucket for the fair packets to be forwarded.

In the more complex case where one or more tiered arbiters are attached at the priority level, the policer control policy instance must consider more than just the child policer weights associated with the attached arbiter. If the arbiter is configured with an aggregate rate limit that its children cannot exceed, the policer control policy instance will switch to calculating the rate each child serviced by the arbiter should receive and enforces that rate using each child policers PIR leaky bucket.

When the child policer PIR leaky bucket is used to limit the bandwidth for the child policer and the child's PIR bucket discard threshold is reached, packets associated with the child policer are discarded. The child policer's discarded packets do not consume depth in the child policer's CIR or FIR buckets. The child policers discarded packets are also prevented from impacting the parent policer and will not consume the aggregate bandwidth managed by the parent policer.

Parent Policer Priority Level Thresholds

As stated above, each child policer is attached either to the root arbiter or explicitly to one of the tier 1 or tier 2 arbiters. When attached directly to the root arbiter, its priority relative to all other child policers is indicated by the parenting level parameter. When attached through one of the tiered arbiters, the parenting hierarchy of the arbiters must be traced through to the ultimate attachment to the root arbiter. The parenting level parameter of the arbiter parented to the root arbiter defines the child policer's priority level within the parent policer.

The priority level is important since it defines the parent policer discard thresholds that will be applied at the parent policer. The parent policer has 8 levels of strict priority and each priority level has its own discard-unfair and discard-all thresholds. Each priority's thresholds are larger than the thresholds of the lower priority levels. This ensures that when the parent policer is discarding, it will be priority sensitive.

To visualize the behavior of the parent policer, picture that when the aggregate forwarding rate of all child policers is currently above the decrement rate of the parent PIR leaky bucket, the bucket depth will increase over time. As the bucket depth increases, it will eventually cross the lowest priority's discard-unfair threshold. If this amount of discard sufficiently lowers the remaining aggregate child policer rate, the parent PIR bucket will hover around this bucket depth. If however, the remaining aggregate child rate is still greater than the decrement rate, the bucket will continue to rise and eventually reach the lowest priority's discard-all threshold which will cause all packets associated with the priority level to be discarded (fair and unfair). Again, if the remaining aggregate child rate is less than or equal to the bucket decrement rate, the parent PIR bucket will hover around this higher bucket depth. If the remaining aggregate child rate is still higher than the decrement rate, the bucket will continue to rise through the remaining priority level discards until equilibrium is achieved.

As noted above, each child's rate feeding into the parent policer is governed by the child policer's PIR bucket decrement rate. The amount of bandwidth the child policer offers to the parent policer will not exceed the child policer's configured maximum rate.

Root Arbiter's Parent Policer's Priority Aggregate Thresholds

Each policer-control-policy root arbiter supports configurable aggregate priority thresholds which are used to control burst tolerance within each priority level. Two values are maintained per priority level; the shared-portion and the fair-portion. The shared-portion represents the amount of parent PIR bucket depth that is allowed to be consumed by both fair and unfair child packets at the priority level. The fair-portion represents the amount of parent PIR bucket depth that only the fair child policer packets may consume within the priority level. It should be noted that the fair and unfair child packets associated with a higher parent policer priority level may also consume the bucket depth set aside for this priority.

While the policy maintains a parent policer default or explicit configurable values for shared-portion and fair-portion within each priority level, it is possible that some priority levels will not be used within the parent policer. Most parent policer use cases require fewer than eight strict priority levels.

In order to derive the actual priority level discard-unfair and discard-all thresholds while only accounting for the actual in-use priority levels, the system maintains a child policer to parent policer association counter per priority level for each policer control policy instance. As a child policer is parented to either the root or a tiered arbiter, the system determines the parent policer priority level for the child policer and increments the association counter for that priority level on the parent policer instance.

The shared-portion for each priority level is affected by the parent policer global min-thresh-separation parameter that defines the minimum separation between any in-use discard thresholds. When more than one child policer is associated with a parent policer priority level, the shared-portion for that priority level will be the current value of min-thresh-separation. When only a single child policer is associated, the priority level’s shared-portion is zero since all packets from the child will be marked fair and the discard-unfair threshold is meaningless. When the association counter is zero, both the shared-portion and the fair-portion for that priority level are zero since neither discard thresholds will be used. Whenever the association counter is greater than 0, the fair-portion for that priority level will be derived from the current value of the priority’s mbs-contribution parameter and the global min-thresh-separation parameter.

Each priority level’s discard-unfair and discard-all thresholds are calculated based on an accumulation of lower priorities shared-portions and fair-portions and the priority level’s own shared-portion and fair-portion. The base threshold value for each priority level is equal to the sum of all lower priority level’s shared-portions and fair-portions. The discard-unfair threshold is the priority level’s base threshold plus the priority level’s shared-portion. The discard-all threshold for the priority level is the priority level’s base threshold plus both the shared-portion and fair-portion values of the priority. As can be seen, an in-use priority level’s thresholds are always greater than the thresholds of lower priority levels.

Policer Control Policy Application

A policer-control-policy may be applied on any Ethernet ingress or egress SAP that is associated with a port (or ports in the case of LAG).

The **no** form of the command removes a non-associated policer control policy from the system. The command will not execute when policer-name is currently associated with any SAP context.

Parameters *policy-name* — Specifies the policy name up to 32 characters in length. Each policer-control-policy must be created with a unique policy name. The name must adhere to the system policy ASCII naming requirements. If the defined policy name already exists, the system will enter that policy’s context for editing purposes. If policy name does not exist, the system will attempt to create a policy with the specified name.

scheduler-override

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

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

Description This command specifies the set of attributes whose values have been overridden by management on this virtual scheduler. Clearing a given flag will return the corresponding overridden attribute to the value defined on the SAP's ingress and egress scheduler policy.

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

scheduler

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

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

Description This command override specifics attributes of the specified scheduler name.

A scheduler defines bandwidth controls that limit each child (other schedulers, policers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have policers, queues or other schedulers defined as child associations. The scheduler can be a child which takes bandwidth from a scheduler in a higher tier. A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers.

Each scheduler must have a unique name within the context of the scheduler policy; however the same name can be reused in multiple scheduler policies. If *scheduler-name* already exists within the policy tier level (regardless of the inclusion of the keyword create), the context changes to that scheduler name for the purpose of editing the scheduler parameters. Modifications made to an existing scheduler are executed on all instantiated schedulers created through association with the policy of the edited scheduler. This can cause policer, queues or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce service level agreements (SLAs).

If the *scheduler-name* exists within the policy on a different tier (regardless of the inclusion of the keyword create), an error occurs and the current CLI context will not change.

If the *scheduler-name* does not exist in this or another tier within the scheduler policy, it is assumed that an attempt is being made to create a scheduler of that name. The success of the command execution is dependent on the following:

Step 1. The maximum number of schedulers has not been configured.

Step 2. The provided *scheduler-name* is valid.

Step 3. The **create** keyword is entered with the command if the system is configured to require it (enabled in the **environment create** command).

When the maximum number of schedulers has been exceeded on the policy, a configuration error occurs and the command will not execute, nor will the CLI context change.

If the provided scheduler-name is invalid according to the criteria below, a name syntax error will occur, the command will not execute, and the CLI context will not change.

The **no** form of the command disables the scheduler override.

Parameters *scheduler-name* — Specifies the name of the scheduler.

Values Valid names consist of any string up to 32 characters in length, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

Default **None.** Each scheduler must be explicitly created.

create — This optional keyword explicitly specifies that it is acceptable to create a scheduler with the given *scheduler-name*. If the **create** keyword is omitted, **scheduler-name** is not created when the system environment variable create is set to true. This safeguard is meant to avoid accidental creation of system objects (such as schedulers) while attempting to edit an object with a mistyped name or ID. The keyword has no effect when the object already exists.

parent

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

Context config>service>cust>multi-service-site>ingress>sched-override>scheduler
config>service>cust>multi-service-site>egress>sched-override>scheduler

Description This command overrides the scheduler’s parent weight and CIR weight information. The weights apply to the associated level or cir-level configured in the applied scheduler policy. The scheduler name must exist in the scheduler policy applied to the ingress or egress of the SAP or multi-service site.

The override weights are ignored if the scheduler does not have a **parent** command configured in the scheduler policy. This allows the parent of the scheduler to be removed from the scheduler policy without having to remove all of the SAP/MSS overrides. If the parent scheduler does not exist causing the configured scheduler to be fostered on an egress port scheduler, the override weights will be ignored and the default values used; this avoids having non-default weightings for fostered schedulers.

The **no** form of the command returns the scheduler's parent weight and CIR weight to the value configured in the applied scheduler policy.

Default no parent

Parameters *weight* — Defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same strict **level** defined by the level parameter in the applied scheduler policy. Within the level, all weight values from active children at that level are summed and the ratio of each active child's weight to the total is used to distribute the available bandwidth at that level. A weight is considered to be active when the policer, queue or scheduler the weight pertains to has not reached its maximum rate and still has packets to transmit.
A 0 (zero) weight value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict level.

Values 0 to 100

Default 1

cir-weight — Defines the relative weight of this scheduler in comparison to other child schedulers and queues at the same *cir-level* defined by the **cir-level** parameter in the applied scheduler policy. Within the strict **cir-level**, all **cir-weight** values from active children at that level are summed and the ratio of each active child's **cir-weight** to the total is used to distribute the available bandwidth at that level. A **cir-weight** is considered to be active when the policer, queue or scheduler that the **cir-weight** pertains to has not reached the CIR and still has packets to transmit.
A 0 (zero) **cir-weight** value signifies that the child scheduler will receive bandwidth only after bandwidth is distributed to all other non-zero weighted children in the strict cir-level.

Values 0 to 100

Default 0

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>service>cust>multi-service-site>ingress>sched-override>scheduler
config>service>cust>multi-service-site>egress>sched-override>scheduler

Description This command overrides specific attributes of the specified scheduler rate.

The **rate** command defines the maximum bandwidth that the scheduler can offer its child policers, queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the scheduler's amount of bandwidth to be considered during the parent schedulers 'within CIR' distribution phase.

The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child policers or queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.

When a scheduler is defined without specifying a rate, the default rate is **max**. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler will assume that an infinite amount of bandwidth is available and allow all child queues and schedulers to operate at their maximum rates.

The **no** form of this command returns the scheduler's to the PIR and CIR parameters to the value configured in the applied scheduler policy.

Parameters *pir-rate* — Specifies the PIR rate.

Values 1 to 3200000000, **max**

Default **max**

cir-rate — Specifies the CIR rate.

If the *cir-rate* is set to **max**, then the CIR rate is set to infinity. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers, policers or queues.

Values 0 to 3200000000, **max**, **sum**

Default **sum**

scheduler-policy

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

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

Description This command applies an existing scheduler policy to an ingress or egress scheduler used by SAP queues or, at egress only, policers associated with this multi-service customer site. The schedulers defined in the scheduler policy can only be created once the customer site has been appropriately assigned to a chassis port, channel or slot. Scheduler policies are defined in the **config>qos>scheduler-policy** *scheduler-policy-name* context.

The **no** form of this command removes the configured ingress or egress scheduler policy from the multi-service customer site. When the policy is removed, the schedulers created due to the policy are removed also making them unavailable for the SAP policers and queues associated with the customer site. Policers and queues that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler.

The SAPs that have ingress queues reliant on the removed schedulers enter into an operational state depicting the orphaned status of one or more policers and queues. When the **no scheduler-policy** command is executed, the customer site ingress or egress node will not contain an applied scheduler policy.

Parameters *scheduler-policy-name* — Applies an existing scheduler policy that was created in the **config>qos>scheduler-policy** *scheduler-policy-name* context to create the hierarchy of ingress or egress virtual schedulers. The scheduler names defined within the policy are created and made available to any ingress or egress queues and egress policers managed by HQoS created on associated SAPs.

Values Any existing valid scheduler policy name up to 32 characters in length.

phone

Syntax [**no**] **phone** *phone-number*

Context config>service>cust

Description This command adds telephone number information for a customer ID. The **no** form of this command removes the phone number value from the customer ID.

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

2.19.2.3 MRP Commands

mrp

Syntax **mrp**

Context config>service

Description This command configures a Multi-service Route Processor (MRP).

copy

Syntax **copy** *source-name* **to** *dest-name*

Context config>service>mrp

Description This command copies existing **mrp-policy** list entries for a specific policy name to another policy name. The copy command is a configuration level maintenance tool used to create a new **mrp-policy** using an existing **mrp-policy**.

An error will occur if the destination policy name exists.

Parameters **mrp-policy** — Indicates that source-name and dest-name are MRP policy names.

source-name — Identifies the source **mrp-policy** from which the copy command will attempt to copy. The **mrp-policy** with this name must exist for the command to be successful.

dest-name — Identifies the destination **mrp-policy** to which the copy command will attempt to copy. If the **mrp-policy** with dest-name exist within the system an error message is generated.

mrp-policy

Syntax **mrp-policy** *policy-name* [**create**]
[**no**] **mrp-policy** *policy-name*

Context config>service>mrp

Description This command enables the context to configure MRP policy parameters. The **mrp-policy** specifies either a forward or a drop action for the Group BMAC attributes associated with the ISIDs specified in the match criteria. The **mrp-policy** can be applied to multiple BVPLS services as long as the scope of the policy is template.

Any changes made to the existing policy, using any of the sub-commands, will be applied immediately to all services where this policy is applied. For this reason, when many changes are required on a mrp-policy, it is recommended that the policy be copied to a work area. That work-in-progress policy can be modified until complete and then written over the original mrp-policy. Use the config mrp-policy copy command to maintain policies in this manner.

The **no** form of the command deletes the mrp-policy. An MRP policy cannot be deleted until it is removed from all the SAPs or SDPs where it is applied.

Default no mrp-policy is defined

Parameters *policy-name* — Specifies the redirect policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

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.

default-action

Syntax **default-action** {**block | allow**}

Context config>service>mrp>mrp-policy

Description This command specifies the action to be applied to the MMRP attributes (Group BMACs) whose ISIDs do not match the specified criteria in all of the entries of the mrp-policy.

When multiple default-action commands are entered, the last command will overwrite the previous command.

Default default-action allow

Parameters **block** — Specifies that all MMRP attributes will not be declared or registered unless there is a specific mrp-policy entry which causes them to be allowed on this SAP or SDP.

allow — Specifies that all MMRP attributes will be declared and registered unless there is a specific mrp-policy entry which causes them to be blocked on this SAP or SDP.

entry

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

Context config>service>mrp>mrp-policy

Description This command creates or edits an mrp-policy entry. Multiple entries can be created using unique entry ID numbers within the policy. The implementation exits the policy on the first match found and executes the actions in accordance with the accompanying action command. For this reason, entries must be sequenced correctly from most to least explicit. An entry may not have any match criteria defined (in which case, everything matches) but must have at least the keyword action for it to be considered complete. Entries without the action keyword will be considered incomplete and hence will be rendered inactive.

The **no** form of the command removes the specified entry from the MRP policy. Entries removed from the MRP policy are immediately removed from all services where the policy is applied.

Parameters *entry-id* — Specifies an entry ID. It is recommended that multiple entries be given entry IDs in staggered increments. This allows users to insert a new entry in an existing policy without requiring renumbering of all the existing entries.

Values 1 to 65535

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.

action

Syntax **action** {*action*}
no action

Context config>serv>mrp>mrp-policy>entry

Description This command specifies the action to be applied to the MMRP attributes (Group BMACs) whose ISIDs match the specified ISID criteria in the related entry.

The action keyword must be entered for the entry to be active. Any filter entry without the action keyword will be considered incomplete and will be inactive. If neither keyword is specified (no action is used), this is considered a No-Op policy entry used to explicitly set an entry inactive without modifying match criteria or removing the entry itself. Multiple action statements entered will overwrite previous actions parameters when defined. To remove a parameter, use the **no** form of the action command with the specified parameter.

The **no** form of the command removes the specified action statement. The entry is considered incomplete and hence rendered inactive without the action keyword.

Default no action

Parameters *action* — Specifies the action for the MRP policy entry.

block — Specifies that the matching MMRP attributes will not be declared or registered on this SAP or SDP.

allow — Specifies that the matching MMRP attributes will be declared and registered on this SAP or SDP.

end-station — Specifies that an end-station emulation is present on this SAP or SDP for the MMRP attributes related with matching ISIDs. Equivalent action with the block keyword on that SAP or SDP. The attributes associated with the matching ISIDs are not declared or registered on the SAP or SDP. The matching attributes on the other hand are mapped as static MMRP entries on the SAP or SDP which implicitly instantiates in the data plane as a MFIB entry associated with that SAP or SDP for the related Group BMAC. For the other SAPs/SDPs in the BVPLS with MRP enabled (no shutdown). This means that the permanent declaration of the matching attributes, as in the case when the IVPLS instances associated with these ISIDs were locally configured.

If an MRP policy has end-station action in one entry, the only default action allowed in the policy is block. Also no other actions are allowed to be configured in other entry configured under the policy.

This policy will apply even if the MRP is shutdown on the local SAP or SDP or for the whole BVPLS to allow for manual creation of MMRP entries in the data plane. Specifically the following rules apply:

- If **service vpls mrp shutdown** is executed, and the MMRP on all SAP or SDPs is shutdown, then MRP PDUs pass-through transparently.
- If **service vpls mrp no shutdown**, and the **endstation** statement (even with no ISID values in the related match statement) is used in an MRP policy applied to SAP or SDP, then no declaration is sent on SAP or SDP. The provisioned ISIDs in the match statement are registered on that SAP or SDP and are propagated on all the other MRP enabled endpoints.

match

Syntax [**no**] **match**

Context config>serv>mrp>mrp-policy>entry

Description This command enables the context to configure match criteria for the MRP policy. When the match criteria have been satisfied the action associated with the match criteria is executed. In the current implementation just one match criteria (ISID-based) is possible in the entry associated with the MRP policy. Only one match statement can be entered per entry.

The **no** form of the command removes the match criteria for the entry ID.

isid

Syntax **isid** *value* [**to** *higher-value*]
no isid
no isid *value* [**to** *higher-value*]

Context config>serv>mrp>mrp-policy>entry>match

Description This command configures an ISID value or a range of ISID values to be matched by the mrp-policy parent when looking at the related MMRP attributes (Group BMACs). The pbb-etype value for the related SAP (inherited from the ethernet port configuration) or for the related SDP binding (inherited from SDP configuration) will be used to identify the ISID tag.

Multiple ISID statements are allowed under a match node. The following rules govern the usage of multiple ISID statements:

- Overlapping values are allowed:
  - isid from 1 to 10
  - isid from 5 to 15
  - isid 16
- The minimum and maximum values from overlapping ranges are considered and displayed. The above entries will be equivalent with the “isid from 1 to 16” statement.

- There is no consistency check with the content of ISID statements from other entries. The entries are evaluated in the order of their IDs and the first match causes the implementation to execute the associated action for that entry and then to exit the mrp-policy.
- If there are no ISID statements under a match criteria but the **mac-filter** type is **isid** the following behaviors apply for different actions:
  - For **end-station**, it treats any ISID value as no match and goes to next entry or default action which must be “block” in this case
  - For **allow**, it treats any ISID value as a match and allows it
  - For **block**, it treats any ISID value as a match and blocks it

The **no** form of the command can be used in two ways:

no isid removes all the previous statements under one match node.

no isid *value* **| from** *value* **to** *higher-value* removes a specific ISID value or range. It must match a previously used positive statement: for example if the command **isid 16 to 100** was used using **no isid 16 to 50** will not work but **no isid 16 to 100** will be successful.

Default no isid

Parameters *value or higher-value* — Specifies the ISID value in 24 bits. When just one value is present, it identifies a particular ISID to be used for matching.

Values 0 to 16777215

from value to higher-value — Identifies a range of ISIDs to be used as matching criteria.

renum

Syntax **renum** *src-entry-id* **to** *dts-entry-id*

Context config>service>mrp>mrp-policy

Description This command renumbers existing MRP policy entries to properly sequence policy entries. This may be required in some cases since the implementation exits when the first match is found and executes the actions according to the accompanying action command. This requires that entries be sequenced correctly from most to least explicit.

Parameters *src-entry-id* — Specifies the entry number of an existing entry.

Values 1 to 65535

dts-entry-id — Specifies the new entry number to be assigned to the old entry. If the new entry exists, an error message is generated.

scope

Syntax **scope** {**exclusive | template**}

no scope

Context config>service>mrp>mrp-policy

Description This command configures the filter policy scope as exclusive or template. If the scope of the policy is template and is applied to one or more services, the scope cannot be changed.

The **no** form of the command sets the scope of the policy to the default of template.

Default scope template

Parameters **exclusive** — When the scope of a policy is defined as exclusive, the policy can only be applied to a single entity (SAP or SDP). Attempting to assign the policy to a second entity will result in an error message. If the policy is removed from the entity, it will become available for assignment to another entity.

template — When the scope of a policy is defined as template, the policy can be applied to multiple SAPs or network ports.

2.19.2.4 Service System Commands

bgp-auto-rd-range

Syntax **bgp-auto-rd-range** *ip-address* **comm-val** *comm-val* **to** *comm-val*
no bgp-auto-rd-range

Context config>service>system

Description This command defines the type-1 route-distinguisher ipv4 address and community value range within which the system will select a route-distinguisher for the bgp-enabled services using auto-rd.

Default no bgp-auto-rd-range

Parameters *ip-address* — Specifies the IPv4 address used in the first 4 octets of all the type-1 auto route-distinguishers selected by the system.

comm-val — Specifies the community value of the type-1 auto route-distinguisher.

Values 1 to 65535

Interactions: This command is used along with the *route-distinguisher auto-rd* command supported in VPLS, VPRN and Epipe services. The system forces the user to create a *bgp-auto-range* before the *auto-rd* option can be used in the services.

Note: The system will keep allocating values for services configured with *route-distinguisher auto-rd* as long as there are available community values within the configured range.

Once the command is added, the following changes are allowed:

- The *ip-address* can be changed without changing the *comm-val* range, even if there are services using auto-rd. The affected routes will be withdrawn and re-advertised with the new route-distinguishers.
- The *comm-val* range can be modified as long as there are not existing conflicting values in the new range. For instance, the user may expand the range as long as the new range does not overlap with existing manual route-distinguishers. The user may also reduce the range as long as the new range can accommodate the already allocated auto-RDs.

gre-eth-bridged

Syntax **gre-eth-bridged**

Context config>service>system

Description This command provides the context for configuring parameters related to termination of a GRE tunnel carrying Ethernet payload onto a PW port by using Forwarding Path Extensions (FPE).

tunnel-termination

Syntax **tunnel-termination** [*ip-address* **|** *ipv6-address*] **fpe** *fpe-id* [**create**]
no tunnel-termination [*ip-address* **|** *ipv6-address*]

Context config>service>system>gre-eth-bridged

Description This command configures an end-point IP address for a GRE tunnel carrying Ethernet payload that is to be terminated on a PW SAP. It also associates this IP address with the FPE object that is providing cross-connect logic between the tunnel and the PW port. This IP address fully supports ICPM protocols such as PING, traceroute, and others.

Parameters *ip-address* — The tunnel end-point IP address in the SR OS node.

ipv6-address — The tunnel end-point IPv6 address in the SR OS node.

fpe *id* — The FPE ID that is providing cross-connect service between the GRE tunnel and the PW port.

Values 1 to 64

vpn-gre-source-ip

Syntax **vpn-gre-source-ip** *ip-address*
no vpn-gre-source-ip

Context config>service>system

Description This command configures a single system-wide alternate source IPv4 address of the GRE tunnels in all VPRN services using the **auto-bind-tunnel** or an explicit SDP binding (**config**>**service**>**vprn**>**spoke-sdp**) with a tunnel of encapsulation GRE.

A change to the value of the **vpn-gre-source-ip** parameter can be performed without disabling the service. Once the new value is configured, the system address is not used in services which bind to the GRE tunnel.

The primary IPv4 address of any local network IP interface, loopback or otherwise, may be used.

The address of the following interfaces are not supported, and the configuration will be rejected:

- unnumbered network IP interface
- IES interface
- VPRN interface
- CSC VPRN interface

The **vpn-gre-source-ip** parameter value adheres to the following rules:

- This single source address counts towards the maximum of 15 distinct address values per system that are used by all GRE SDPs under the **configure**>**service**>**sdp**>**local-end** context and all L2oGRE SDPs under the **config**>**service**>**system**>**gre-eth-bridged**>**tunnel-termination** context.
- The same source address can be used in both **vpn-gre-source-ip** and **configure**>**service**>**sdp**>**local-end** contexts.
- The same source address cannot be used in both **vpn-gre-source-ip** and **config**>**service**>**system**>**gre-eth-bridged**>**tunnel-termination** contexts because an address configured for a L2oGRE SDP matches an internally created interface which is not available to other applications.
- The **vpn-gre-source-ip** address, when different from system, need not match the primary address of an interface which has the MPLS-over-GRE termination subnet configured, unless a GRE SDP or tunnel from the far-end router terminates on this address.

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

Default vpn-gre-source-ip ip-address (System interface primary IPv4 address)

Parameters *ip-address* — Specifies the IPv4 address (a.b.c.d).

2.19.2.5 Oper Group Commands

oper-group

Syntax **oper-group** *group-name* [**create**]
no oper-group *group-name*

Context config>service

Description This command creates a system-wide group (operational group) name which can be used to associate a number of service objects (for example, SAPs or pseudowires). The status of the group is derived from the status of its members. The status of the group can then be used to influence the status of non-member objects. For example, when a group status is marked as down, the object(s) that monitor the group change their status accordingly.

The **no** form of the command removes the group. All the object associations need to be removed before the no form of the command can be executed.

no oper-group

Parameters *group-name* — Specifies the operational group identifier up to 32 characters in length.

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.

bfd-enable

Syntax **bfd-enable interface** *interface-name* **dest-ip** *ipv4-address* [**service** *service-id*]
no bfd-enable

Context config>service>oper-group

Description This command associates a BFD sessions with the named oper-group so that if the BFD session fails then the oper-group is changed to operationally down and all monitoring interfaces should also be brought operationally down.

Parameters *interface-name* — Specifies the source interface, up to 32 characters in length, for the BFD sessions to be monitored for the associated oper-group.

ipv4-address — Specifies the destination IPv4 address for the BFD sessions to be monitored for the associated oper-group.

service-id — Specifies the service ID, up to 64 characters in length, in which the BFD session exists if it is not in the base routing context.

hold-time

Syntax **hold-time**

Context config>service>oper-group

Description This command enables the context to configure hold time information.

group up

Syntax **group up** *time* **| no group up**

Context config>service>oper-group>hold-time

Description This command configures the number of seconds to wait before notifying clients monitoring this group when its operational status transitions from down to up. A value of zero indicates that transitions are reported immediately to monitoring clients. The up time option is a must to achieve fast convergence: when the group comes up, the monitoring MH site which tracks the group status may wait without impacting the overall convergence; there is usually a pair MH site that is already handling the traffic.

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

Default group-up 4

Parameters *time* — Specifies the group up time value.

Values 0 to 3600

group down

Syntax **group down** *time* **| no group down**

Context config>service>oper-group>hold-time

Description This command configures the number of seconds to wait before notifying clients monitoring this group when its operational status transitions from up to down.

The **no** form of the command sets the values back to the default.

Default group down 0

2.19.2.6 Pseudowire (PW) Commands

pw-routing

Syntax **pw-routing**

Context config>service

Description This command enables the context to configure dynamic multi-segment pseudowire (MS-PW) routing. Pseudowire routing must be configured on each node that will be a T-PE or an S-PE.

Default disabled

boot-timer

Syntax **boot-timer** *secs*
no boot-timer

Context config>service>pw-routing

Description This command configures a hold-off timer for MS-PW routing advertisements and signaling and is used at boot time.

The **no** form of this command removes a previously configured timer and restores it to its default.

Default 10

Parameters *timer-value* — Specifies the value of the boot timer in seconds.

Values 0 to 600

local-prefix

Syntax **local-prefix** *local-prefix* [**create**]
no local-prefix *local-prefix*

Context config>service>pw-routing

Description This command configures one or more node prefix values to be used for MS-PW routing. At least one prefix must be configured on each node that is an S-PE or a T-PE.

The **no** form of this command removes a previously configured prefix, and will cause the corresponding route to be withdrawn if it has been advertised in BGP.

Default no local-prefix.

Parameters *local-prefix* — Specifies a 32 bit prefix for the AII. One or more prefix values, up to a maximum of 16, may be assigned to the 7450 ESS, 7750 SR, or 7950 XRS node. The global ID can contain the 2-octet or 4-octet value of the provider's Autonomous System Number (ASN). The presence of a global ID based on the provider's ASN ensures that the AII for spoke-SDPs configured on the node will be globally unique.

Values <global-id>:<ip-addr>|<raw-prefix>

ip-addr	a.b.c.d
raw-prefix	1 to 4294967295
global-id	1 to 4294967295

advertise-bgp

Syntax **advertise-bgp route-distinguisher** *rd* [**community** *community*]
no advertise-bgp route-distinguisher *rd*

Context config>service>pw-routing>local-prefix

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

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

Default no advertise-bgp

Parameters *rd* — Specifies an 8-octet route distinguisher associated with the prefix. Up to 4 unique route distinguishers can be configured and advertised for a given prefix though multiple instances of the advertise-bgp command. This parameter is mandatory.

Values (6 bytes, other 2 Bytes of type will be automatically generated)
asn:number1 (RD Type 0): 2bytes ASN and 4 bytes locally administered number
ip-address:number2 (RD Type 1): 4bytes IPv4 and 2 bytes locally administered number;

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

Values

community	{2-byte-as-number:comm-va1}
2-byte-asnumber	0 to 65535
comm.-val	0 to 65535

path

Syntax **path** *name* [**create**]
no path *name*

Context config>service>pw-routing

Description This command configures an explicit path between this T-PE and a remote T-PE. For each path, one or more intermediate S-PE hops must be configured. A path can be used by multiple multi-segment pseudowires. Paths are used by a 7450 ESS, 7750 SR, or 7950 XRS T-PE to populate the list of Explicit Route TLVs included in the signaling of a dynamic MS-PW.

A path may specify all or only some of the hops along the route to reach a T-PE.

The **no** form of the command removes a specified explicit path from the configuration.

Default no path

Parameters *path-name* — Specifies a locally-unique case-sensitive alphanumeric name label for the MS-PW path of up to 32 characters in length.

hop

Syntax **hop** *hop-index ip-address*
no hop *hop-index*

Context config>service>pw-routing>path

Description This command configures each hop on an explicit path that can be used by one or more dynamic MS-PWs. It specifies the IP addresses of the hops that the MS-PE should traverse. These IP addresses can correspond to the system IP address of each S-PE, or the IP address on which the T-LDP session to a given S-PE terminates.

The **no** form of this command deletes hop list entries for the path. All the MS-PWs currently using this path are unaffected. Additionally, all services actively using these MS-PWs are unaffected. The path must be shutdown first in order to delete the hop from the hop list. The '**no hop hop-index**' command will not result in any action, except for a warning message on the console indicating that the path is administratively up.

Default no hop

Parameters *hop-index* — Specifies a locally significant numeric identifier for the hop. The hop index is used to order the hops specified. The LSP always traverses from the lowest hop index to the highest. The hop index does not need to be sequential.

Values 1 to 1024

ip-address — Specifies the system IP address or terminating IP address for the T-LDP session to the S-PE corresponding to this hop. For a given IP address on a hop, the system will choose the appropriate SDP to use.

retry-count

Syntax **retry-count** [*count*]
no retry-count

Context config>service>pw-routing

Description This optional command specifies the number of attempts software should make to re-establish the spoke SDP after it has failed. After each successful attempt, the counter is reset to zero.

When the specified number is reached, no more attempts are made and the spoke SDP is put into the shutdown state.

Use the **no shutdown** command to bring up the path after the retry limit is exceeded.

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

Default 30

Parameters *count* — Specifies the maximum number of retries before putting the spoke SDP into the shutdown state.

Values 10 to 10000

retry-timer

Syntax **retry-timer** *secs*
no retry-timer

Context config>service>pw-routing

Description This command configures a retry-timer for the spoke-SDP. This is a configurable exponential back-off timer that determines the interval between retries to re-establish a spoke-SDP if it fails and a label withdraw message is received with the status code “All unreachable”.

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

Default 30

Parameters *secs* — Specifies initial retry-timer value in seconds.

Values 10 to 480

spe-address

Syntax **spe-address** *global-id:prefix*
no spe-address

Context config>service>pw-routing

Description This command configures a single S-PE Address for the node to be used for dynamic MS-PWs. This value is used for the pseudowire switching point TLV used in LDP signaling, and is the value used by pseudowire status signaling to indicate the PE that originates a pseudowire status message. Configuration of this parameter is mandatory to enable dynamic MS-PW support on a node.

If the S-PE Address is not configured, spoke-sdps that use dynamic MS-PWs and pw-routing local-prefixes cannot be configured on a T-PE. Furthermore, the node will send a label release for any label mappings received for FEC129 AII type 2.

The S-PE Address cannot be changed unless the dynamic ms-pw configuration is removed. Furthermore, changing the S-PE Address will also result in all dynamic MS-PWs for which this node is an S-PE being released. It is recommended that the S-PE Address should be configured for the life of an MS-PW configuration after reboot of the router.

The **no** form of this command removes the configured S-PE Address.

Default no spe-address

Parameters *global-id* — Specifies a 4-octet value that is unique to the service provider. For example, the global ID can contain the 2-octet or 4-octet value of the provider's Autonomous System Number (ASN).

Values

<global-id:prefix>:	<global-id>:{<prefix>|<ipaddress>}
global-id	1 to 4294967295
prefix	1 to 4294967295
ipaddress	a.b.c.d

static-route

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

Context config>service>pw-routing

Description This command configures a static route to a next hop S-PE or T-PE. Static routes may be configured on either S-PEs or T-PEs.

A default static route is entered as follows:

static-route 0:0:next_hop_ip_addresss

or

static-route 0:0.0.0.0:next_hop_ip_address

The **no** form of this command removes a previously configured static route.

Default no static-route

Parameters *route-name* — Specifies the static pseudowire route.

Values

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

pw-template

Syntax **pw-template** *policy-id* [**use-provisioned-sdp | [prefer-provisioned-sdp**] [**auto-gre-sdp**]] [**create**] [**name** *name*]

no **pw-template** *policy-id*

Context config>service

Description This command configures an SDP template.

Default auto-mpls-sdp

Parameters *policy-id* — Specifies a number that uniquely identifies a template for the creation of an SDP.

Values *policy-id*: 1 to 2147483647

use-provisioned-sdp — Specifies whether to use an already provisioned SDP. When specified, the tunnel manager is consulted for an existing active SDP (with a matching far-end address), and the SDP with the lowest metric is chosen. If there are multiple SDPs with the same metric, then the highest SDP identifier that is oper-up is chosen. The choice of SDP can be configured by applying **sdp-include/exclude** in the PW template together with an sdp-group in the provisioned SDPs. This option, and the **auto-gre-sdp** option, are mutually exclusive.

prefer-provisioned-sdp — Specifies that if an existing matching SDP that conforms to any restrictions defined in the **pw-template** is found (for example, **sdp-include/ exclude** *group*), then it will be used, following the same logic as for the **use-provisioned-sdp** parameter. Otherwise, the command will automatically create an SDP in the same manner as if the user did not specify any option. This option and the **use-provisioned-sdp** option are mutually exclusive.

auto-gre-sdp — Specifies that an SDP should automatically be created using a GRE tunnel. This option and the **use-provisioned-sdp** option are mutually exclusive. The PW template parameters **hash-label**, **entropy-label** and **sdp-include/exclude** are ignored when an GRE SDP is auto-created.

auto-mpls-sdp — Specifies that an SDP should automatically be created using an MPLS tunnel. This is the default.

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.

name *name* — A name of the operator's choice, up to 64 characters. The name is saved as part of the configuration.

If a name is not specified at creation time, then SR OS assigns a string version of the policy-id as the name.

Values *name*: 64 characters maximum

accounting-policy

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

Context 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 accounting policy reverts to the default.

Default no accounting-policy

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

Values 1 to 99

auto-learn-mac-protect

Syntax [**no**] **auto-learn-mac-protect**

Context config>service>pw-template
config>service>pw-template>split-horizon-group

Description This command specifies whether to enable automatic population of the MAC protect list with source MAC addresses learned on the associated with this SHG. For more information about auto-learn MAC protect, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN*.

The **no** form of the command disables the automatic population of the MAC protect list.

Default auto-learn-mac-protect

block-on-peer-fault

Syntax [**no**] **block-on-peer-fault**

Context config>service>pw-template

Description When enabled, this command blocks the transmit direction of a pseudowire when any of the following pseudowire status codes is received from the far end PE:

0x00000001	Pseudowire Not Forwarding
0x00000002	Local Attachment Circuit (ingress) Receive Fault
0x00000004	Local Attachment Circuit (egress) Transmit Fault
0x00000008	Local PSN-facing PW (ingress) Receive Fault
0x00000010	Local PSN-facing PW (egress) Transmit Fault

The transmit direction is unblocked when the following pseudowire status code is received:

0x00000000	Pseudowire forwarding (clear all failures)

This command is mutually exclusive with **no pw-status-signaling**, and **standby-signaling-slave**. It is not applicable to spoke SDPs forming part of an MC-LAG or spoke SDPs in an endpoint.

Default no block-on-peer-fault

collect-stats

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

Context config>service>pw-template

Description This command enables accounting and statistical data collection for either the PW template. 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 or XCM 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

controlword

Syntax [**no**] **controlword**

Context config>service>pw-template

Description This command enables the use of the control word on pseudowire packets in VPLS and VPWS 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/VPWS 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 controlword

disable-aging

Syntax [**no**] **disable-aging**

Context config>service>pw-template

Description This command disables MAC address aging across a service.

The **no** form of this command enables aging.

Default no disable-aging

disable-learning

Syntax [**no**] **disable-learning**

Context config>service>pw-template

Description This command enables learning of new MAC addresses.

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

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

Default no disable-learning (Normal MAC learning is enabled)

discard-unknown-source

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

Context config>service>pw-template

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-source

egress

Syntax **egress**

Context config>service>pw-template

Description This command enables the context to configure spoke SDP binding egress filter parameters.

filter

Syntax **filter ip** *ip-filter-id*
filter ipv6 *ipv6-filter-id*
filter mac *mac-filter-id*
no filter [**ip** *ip-filter-id*] [**mac** *mac-filter-id*] [**ipv6** *ipv6-filter-id*]

Context config>service>pw-template>egress
config>service>pw-template>ingress

Description This command associates an IP filter policy or MAC filter policy on egress or ingress. 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.

This command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **filter-name** command can be used in all configuration modes.

This command is mutually exclusive with the **filter-name** command. Only one or the other can be configured.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Parameters

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

Values 1 to 65535

ipv6-filter-id — Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 to 65535

mac-filter-id — Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.

Values 1 to 65535

filter-name

Syntax

filter-name ip *ip-name*
filter-name ipv6 *ipv6-name*
filter-name mac *mac-name*
no filter-name [**ip**] [**ipv6**] [**mac**]

Context

config>service>pw-template>egress
config>service>pw-template>ingress

Description

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

This command is mutually exclusive with the **filter** command. Only one or the other can be configured.

The **no** form of this command removes any configured filter name association with the SAP or IP interface. The filter name itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter name and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Parameters *ip-name* — Specifies the IP filter policy. The filter name must already exist within the created IP filters, up to 64 characters.

ipv6-name — Specifies the IPv6 filter policy. The filter name must already exist within the created IPv6 filters, up to 64 characters.

mac-name — Specifies the MAC filter policy. The specified filter name must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters, up to 64 characters.

mfib-allowed-mda-destinations

Syntax **mfib-allowed-mda-destinations**

Context config>service>pw-template>egress

Description This command enables the context to configure MFIB-allowed XMA or MDA destinations.

The allowed-mda-destinations node and the corresponding **mda** command are used on spoke and mesh SDP bindings to provide a list of XMA or MDA destinations in the chassis that are allowed as destinations for multicast streams represented by [*,g] and [s,g] multicast flooding records on the VPLS service. The XMA or MDA list only applies to IP multicast forwarding when IGMP snooping is enabled on the VPLS service. The XMA or MDA list has no effect on normal VPLS flooding such as broadcast, Layer 2 multicast, unknown destinations or non-snooped IP multicast.

At the IGMP snooping level, a spoke or mesh SDP binding is included in the flooding domain for an IP multicast stream when it has either been defined as a multicast router port, received a IGMP query through the binding or has been associated with the multicast stream through an IGMP request by a host over the binding. Due to the dynamic nature of the way that a spoke or mesh SDP binding is associated with one or more egress network IP interfaces, the system treats the binding as appearing on all network ports. This causes all possible network destinations in the switch fabric to be included in the multicast streams flooding domain. The XMA or MDA destination list provides a simple mechanism that narrows the IP multicast switch fabric destinations for the spoke or mesh SDP binding.

If no XMAs or MDAs are defined within the allowed-mda-destinations node, the system operates normally and will forward IP multicast flooded packets associated with the spoke or mesh SDP binding to all switch fabric taps containing network IP interfaces.

The XMA or MDA inclusion list should include all XMAs or MDAs that the SDP binding may attempt to forward through. A simple way to ensure that an XMA or MDA that is not included in the list is not being used by the binding is to define the SDP the binding is associated with as MPLS and use an RSVP-TE LSP with a strict egress hop. The XMA or MDA associated with the IP interface defined as the strict egress hop should be present in the inclusion list.

If the inclusion list does not currently contain the XMA or MDA that the binding is forwarding through, the multicast packets will not reach the destination represented by the binding. By default, the XMA or MDA inclusion list is empty.

If an XMA or MDA is removed from the list, the XMA or MDA is automatically removed from the flooding domain of any snooped IP multicast streams associated with a destination on the XMA or MDA unless the XMA or MDA was the last XMA or MDA on the inclusion list. Once the inclusion list is empty, all XMAs or MDAs are eligible for snooped IP multicast flooding for streams associated with the SDP binding.

mda

Syntax [**no**] **mda** *mda-id*

Context config>service>pw-template>egress>mfib-mda

Description This command specifies an MFIB-allowed media adapter destination for an SDP binding configured in the system.

Parameters *mda-id* — Specifies an MFIB-allowed media adapters destination.

Values 1, 2

qos

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

Context config>service>pw-template>egress

Description This command is used to redirect PW packets to an egress port queue-group for the purpose of shaping.

The egress PW shaping provisioning model allows the mapping of one or more PWs to the same instance of queues, or policers and queues, that are defined in the queue-group template.

Operationally, the provisioning model consists of the following steps:

1. Create an egress queue-group template and configure queues only, or policers and queues for each FC that needs to be redirected.
2. Apply the queue-group template to the network egress context of all ports where there exists a network IP interface that the PW packets can be forwarded on. This creates one instance of the template on the egress of the port. One or more instances of the same template can be created.
3. Configure FC-to-policer or FC-to-queue mappings together with the redirect to a queue-group in the egress context of a network QoS policy. No queue-group name is specified in this step, which means the same network QoS policy can redirect different PWs to different queue-group templates.
4. Apply this network QoS policy to the egress context of a spoke-SDP inside a service, or to the egress context of a PW template and specify the redirect queue-group name.

One or more spoke-SDPs can have their FCs redirected to use queues only, or queues and policers in the same queue-group instance.

The following are the constraints and rules of this provisioning model.

1. When a PW FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name does not exist, the association is failed at the time the user associates the egress context of a spoke-SDP to the named queue-group. In such a case, the PW packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the PW packet is forwarded on. This queue can be a queue-group queue or the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port. This is the existing implementation and default behavior for a PW packet.
2. When a PW FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name exists but the policer-id and/or the queue-id is not defined in the queue-group template, the association is failed at the time the user associates the egress context of a spoke-SDP to the named queue-group. In such a case, the PW packet will be fed directly to the corresponding egress queue for that FC used by the IP network interface the PW packet is forwarded on.
3. When a PW FC is redirected to use a queue or a policer and a queue in a queue-group and the queue-group name exists and the policer-id or policer-id plus queue-id exist, it is not required to check that an instance of that queue-group exists in all egress network ports that have network IP interfaces. The handling of this is dealt with in the data path as follows:
   - When a PW packet for that FC is forwarded and an instance of the referenced queue-group name exists on that egress port, the packet is processed by the queue-group policer and will then be fed to the queue-group queue.
   - When a PW packet for that FC is forwarded and an instance of the referenced queue-group name does not exist on that egress port, the PW packet will be fed directly to the corresponding egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.
4. If a network QoS policy is applied to the egress context of a PW, any PW FC that is not explicitly redirected in the network QoS policy will have the corresponding packets feed directly the corresponding the egress shared queue for that FC defined in the network-queue policy applied to the egress of this port.

When the queue-group name the PW is redirected to exists and the redirection succeeds, the marking of the packet's DEI/dot1p/DSCP and the tunnel's DEI/dot1p/DSCP/EXP is performed according to the relevant mappings of the {FC, profile} in the egress context of the network QoS policy applied to the PW. This is true regardless of whether an instance of the queue-group exists or not on the egress port the PW packet is forwarded to. If the packet's profile value changed due to egress child policer CIR profiling, the new profile value is used to mark the packet's DEI/dot1p and the tunnel's DEI/dot1p/EXP, and the DSCP/prec will be remarked if **enable-dscp-prec-marking** is enabled under the policer.

When the queue-group name the PW is redirected does not exist, the redirection command is failed. In this case, the marking of the packet's DEI/dot1p/DSCP and the tunnel's DEI/dot1p/DSCP/EXP fields is performed according to the relevant commands in the egress context of the network QoS policy applied to the network IP interface the PW packet is forwarded to.

The **no** version of this command removes the redirection of the PW to the queue-group.

Parameters *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

This variant of the command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **qos name** *network-policy-name* variant can be used in all configuration modes.

Values 1 to 65535

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

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

Values 1 to 65535

name *network-policy-name* — Specifies the network policy name. The value uniquely identifies the policy on the system, up to 64 characters.

entropy-label

Syntax [**no**] **entropy-label**

Context config>service>pw-template

Description This command enables or disables the use of an entropy label for the service, spoke SDP or SDPs to which the pseudowire template applies.

If **entropy-label** is configured, the entropy label and ELI are inserted in packets for which at least one LSP in the stack for the far end of the tunnel used by the service has advertised entropy label capability. If the tunnel type is RSVP, **entropy-label** can also be controlled under the **config>router>mpls** or **config>router>mpls>lsp** contexts.

The entropy label and hash label features are mutually exclusive. The entropy label cannot be configured on a spoke SDP or service where the hash label has already been configured.

force-qinq-vc-forwarding

Syntax [**no**] **force-qinq-vc-forwarding**

Context config>service>epipe>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>pw-template

Description This command forces two VLAN tags to be inserted and removed for spoke and mesh SDPs that have either **vc-type ether** or **vc-type vlan**. The use of this command is mutually exclusive with the **force-vlan-vc-forwarding** command.

The VLAN identifiers and dot 1p/DE bits inserted in the two VLAN tags are taken from the inner tag received on a qinq SAP or qinq mesh/spoke SDP, or from the VLAN tag received on a dot1q SAP or mesh/spoke SDP (with **vc-type vlan** or **force-vlan-vc-forwarding**), or taken from the outer tag received on a qtag.* SAP or 0 if there is no service delimiting VLAN tag at the ingress SAP or mesh/spoke SDP. The VLAN identifiers in both VLAN tags can be set to the value configured in the **vlan-vc-tag** parameter in the **pw-template** or under the mesh/spoke SDP configuration. In the received direction, the VLAN identifiers are ignored and the dot1p/DE bits are not used for ingress classification. However, the inner dot1p/DE bits are propagated to the egress QoS processing.

The Ether type inserted and used to determine the presence of a received VLAN tag for both VLAN tags is 0x8100. A different Ether type can be used for the outer VLAN tag by configuring the PW template with **use-provisioned-sdps** and setting the Ether type using the SDP **vlan-vc-etype** parameter (this Ether type value is then used for all mesh/spoke SDPs using that SDP).

The **no** version of this command sets default behavior.

force-vlan-vc-forwarding

Syntax [**no**] **force-vlan-vc-forwarding**

Context config>service>pw-template

Description This command forces vc-vlan-type forwarding in the data path for spoke and mesh SDPs which have **ether** vc-type. This command is not allowed on vlan-vc-type SDPs.

The system expects a symmetrical configuration with its peer, specifically it expects to remove the same number of VLAN tags from received traffic as it adds to transmitted traffic. As some of the related configuration parameters are local and not communicated in the signaling plane, an asymmetrical behavior cannot always be detected and so cannot be blocked. Consequently, protocol extractions will not necessarily function for asymmetrical configurations as they would with a symmetrical configurations resulting in an unexpected operation.

The **no** version of this command sets default behavior.

Default no force-vlan-vc-forwarding

hash-label

Syntax **hash-label** [**signal-capability**]
no hash-label

Context config>service>pw-template

Description This command enables the use of the hash label on a VLL, VPRN or VPLS service bound to any MPLS type encapsulated SDP as well as to a VPRN service using the **auto-bind-tunnel** with the **resolution-filter** set to any MPLS tunnel type. This feature is not supported on a service bound to a GRE SDP or for a VPRN service using the autobind mode with the gre option. This feature is also not supported on multicast packets forwarded using RSVP P2MP LSP or mLDP LSP in both the base router instance and in the multicast VPN (mVPN) instance. It is, however, supported when forwarding multicast packets using an IES/VPRN spoke-interface.

When this feature is enabled, the ingress data path is modified such that the result of the hash on the packet header is communicated to the egress data path for use as the value of the label field of the hash label. The egress data path appends the hash label at the bottom of the stack (BoS) and sets the S-bit to one (1).

In order to allow applications where the egress LER infers the presence of the hash label implicitly from the value of the label, the Most Significant Bit (MSB) of the result of the hash is set before copying into the Hash Label. This means that the value of the hash label will always be in the range [524,288 - 1,048,575] and will not overlap with the signaled/static LSP and signaled/static service label ranges. This also guarantees that the hash label will not match a value in the reserved label range.

The (unmodified) result of the hash continues to be used for the purpose of ECMP and LAG spraying of packets locally on the ingress LER. However, for VLL services, the result of the hash is overwritten and the ECMP and LAG spraying will be based on service-id when ingress SAP shared queuing is not enabled. However, the hash label will still reflect the result of the hash such that an LSR can use it to perform fine grained load balancing of VLL pseudowire packets.

Packets generated in CPM and that are forwarded labeled within the context of a service (for example, OAM packets) must also include a Hash Label at the BoS and set the S-bit accordingly.

The TTL of the hash label is set to a value of 0.

The user enables the signaling of the hash-label capability under a VLL spoke-sdp, a VPLS spoke-sdp or mesh-sdp, or an IES/VPRN spoke interface by adding the **signal-capability** option. In this case, the decision whether to insert the hash label on the user and control plane packets by the local PE is solely determined by the outcome of the signaling process and can override the local PE configuration. The following are the procedures:

- The local PE will insert the flow label interface parameters sub-TLV with F=1 in the PW ID FEC element in the label mapping message for that spoke-sdp or mesh-sdp.
- If the remote PE includes this sub-TLV with F=1 or F=0, then local PE must insert the hash label in the user and control plane packets.
- If remote PE does not include this sub-TLV (for example, it does not support it, or it is supported but the user did not enable the **hash-label** option or the **signal-capability** option), then the local PE establishes the pseudowire but must not insert the hash label in the user and control packets over that spoke-sdp or mesh-sdp. If the remote PE does not support the **signal-capability** option, then there are a couple of possible outcomes:
  - If the **hash-label** option was enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the pseudowire packets received by the local PE will have the hash label included. These packets must be dropped. The only way to solve this is to disable the signaling capability option on the local node which will result in the insertion of the hash label by both PE nodes.
  - If the **hash-label** option is not supported or was not enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the pseudowire received by the local PE will not have the hash label included.
- The user can enable or disable the signal-capability option in CLI as needed. When doing so, the router must withdraw the label it sent to its peer and send a new label mapping message with the new value of the F bit in the flow label interface parameters sub-TLV of the PW ID FEC element.

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

Default no hash-label

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

igmp-snooping

Syntax **igmp-snooping**

Context config>service>pw-template

Description This command enables the Internet Group Management Protocol (IGMP) snooping context.

Default none

fast-leave

Syntax [**no**] **fast-leave**

Context config>service>pw-template>igmp-snooping

Description This command enables fast leave.

When IGMP fast leave processing is enabled, the 7750 SR will immediately remove a SAP or SDP from the IP multicast group when it detects an IGMP **leave** on that SAP or SDP. Fast leave processing allows the switch to remove a SAP or SDP that sends a **leave** from the forwarding table without first sending out group-specific queries to the SAP or SDP, and thus speeds up the process of changing channels (zapping).

Fast leave should only be enabled when there is a single receiver present on the SAP or SDP.

When fast leave is enabled, the configured last-member-query-interval value is ignored.

Default no fast-leave

import

Syntax **import** *policy-name*
no import

Context config>service>pw-template>igmp-snooping

Description This command specifies the import routing policy to be used for IGMP packets. Only a single policy can be imported at a time.

The **no** form of the command removes the policy association.

Default no import

Parameters *policy-name* — Specifies the import policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes. Routing policies are configured in the config>router>policy-options context The router policy must be defined before it can be imported.

last-member-query-interval

Syntax **last-member-query-interval** *interval*
no last-member-query-interval

Context 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 last-member-query-interval 10

Parameters *interval* — Specifies the frequency, in tenths of seconds, at which query messages are sent.

Values 1 to 50

max-num-groups

Syntax **max-num-groups** *count*
no max-num-groups

Context config>service>pw-template>igmp-snooping

Description This command defines the maximum number of multicast groups that can be joined. If the router 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.

Values 1 to 1000

query-interval

Syntax **query-interval** *seconds*
no query-interval

Context 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 query-interval 125

Parameters *seconds* — Specifies the time interval, in seconds, that the router transmits general host-query messages.

Values 2 to 1024

query-response-interval

Syntax **query-response-interval** *seconds*

no query-response-interval

Context 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 IGMPv2/v3 queries.

The configured query-response-interval must be smaller than the configured query-interval.

If send-queries is not enabled on this SAP or SDP, the configured query-response-interval value is ignored.

Default query-response-interval 10

Parameters *seconds* — Specifies the length of time to wait to receive a response to the host-query message from the host.

Values 1 to 1023

robust-count

Syntax **robust-count** *robust-count*
no robust-count

Context config>service>pw-template>igmp-snooping

Description If the **send-queries** command is enabled, this parameter allows tuning for the expected packet loss. The robust-count variable allows tuning for the expected packet loss on a subnet and is comparable to a retry count.

If send-queries is not enabled, this parameter will be ignored.

Default robust-count 2

Parameters *robust-count* — Specifies the robust count for the SAP or SDP.

Values 2 to 7

send-queries

Syntax [**no**] **send-queries**

Context config>service>pw-template>igmp-snooping

Description This command specifies whether to send IGMP general query messages.

When send-queries is configured, all type of queries generate ourselves are of the configured version. If a report of a version higher than the configured version is received, the report will get dropped and a new wrong version counter will get incremented.

If send-queries is not configured, the version command has no effect. The version used on that SAP or SDP will be the version of the querier. This implies that, for example, when we have a v2 querier, we will never send out a v3 group or group-source specific query when a host wants to leave a certain group.

Default no send-queries

version

Syntax **version** *version*
no version

Context config>service>pw-template>igmp-snooping

Description This command specifies the version of IGMP. This object can be used to configure a router capable of running either value. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP on that LAN.

When the **send-query** command is configured, all type of queries generate ourselves are of the configured **version**. If a report of a version higher than the configured version is received, the report gets dropped and a new “wrong version” counter is incremented.

If the **send-query** command is not configured, the **version** command has no effect. The version used on that SAP or SDP will be the version of the querier. This implies that, for example, when there is a v2 querier, a v3 group or group-source specific query when a host wants to leave a certain group will never be sent.

Default version 3

Parameters *version* — Specifies the IGMP version.

Values 1, 2, 3

ingress

Syntax **ingress**

Context config>service>pw-template

Description This command enables the context to configure spoke SDP binding ingress filter parameters.

qos

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

Context config>service>pw-template>ingress

Description This command is used to redirect pseudowire packets to an ingress forwarding plane queue-group for the purpose of rate-limiting.

The ingress pseudowire rate-limiting feature uses a policer in queue-group provisioning model. This model allows the mapping of one or more pseudowires to the same instance of policers which are defined in a queue-group template.

Operationally, the provisioning model in the case of the ingress pseudowire shaping feature consists of the following steps:

1. Create an ingress queue-group template and configure policers for each FC which needs to be redirected and optionally for each traffic type (unicast or multicast).
2. Apply the queue-group template to the network ingress forwarding plane where there exists a network IP interface which the pseudowire packets can be received on. This creates one instance of the template on the ingress of the FP. One or more instances of the same template can be created.
3. Configure FC-to-policer mappings together with the policer redirect to a queue-group in the ingress context of a network QoS policy. No queue-group name is specified in this step which means the same network QoS policy can redirect different pseudowires to different queue-group templates.
4. Apply this network QoS policy to the ingress context of a spoke-sdp inside a service or to the ingress context of a pseudowire template and specify the redirect queue-group name.

One or more spoke-sdps can have their FCs redirected to use policers in the same policer queue-group instance.

The following are the constraints and rules of this provisioning model when used in the ingress pseudowire rate-limiting feature:

1. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name does not exist, the association is failed at the time the user associates the ingress context of a spoke-sdp to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the media adapters and FP.
2. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists but the policer-id is not defined in the queue-group template, the association is failed at the time the user associates the ingress context of a spoke-sdp to the named queue-group. In such a case, the pseudowire packet will feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the media adapters and FP.
3. When a pseudowire FC is redirected to use a policer in a named policer queue-group and the queue-group name exists and the policer-id is defined in the queue-group template, it is not required to check that an instance of that queue-group exists in all ingress FPs which have network IP interfaces. The handling of this is dealt with in the data path as follows:

- When a pseudowire packet for that FC is received and an instance of the referenced queue-group name exists on that FP, the packet is processed by the policer and will then feed the per-FP ingress shared queues referred to as “policer-output-queues”.
- When a pseudowire packet for that FC is received and an instance of the referenced queue-group name does not exist on that FP, the pseudowire packets will be fed directly into the corresponding ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the media adapters and FP.

4. If a network QoS policy is applied to the ingress context of a pseudowire, any pseudowire FC which is not explicitly redirected in the network QoS policy will have the corresponding packets feed directly the ingress network shared queue for that FC defined in the network-queue policy applied to the ingress of the media adapters and FP.
5. If no network QoS policy is applied to the ingress context of the pseudowire, then all packets of the pseudowire will feed:
   - the ingress network shared queue for the packet’s FC defined in the network-queue policy applied to the ingress of the media adapters and FP. This is the default behavior.
   - a queue-group policer followed by the per-FP ingress shared queues referred to as “policer-output-queues” Good received is redirected to a queue-group. The only exceptions to this behavior are for packets received from a IES/VPRN spoke interface and from a R-VPLS spoke-sdp which is forwarded to the R-VPLS IP interface. In these two cases, the ingress network shared queue for the packet’s FC defined in the network-queue policy applied to the ingress of the XMA, MDA, or FP is used. When a pseudowire is redirected to use a policer queue-group, the classification of the packet for the purpose of FC and profile determination is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the pseudowire. This is true regardless if an instance of the named policer queue-group exists on the ingress FP the pseudowire packet is received on. The user can apply a QoS filter matching the dot1p in the VLAN tag corresponding to the Ethernet port encapsulation, the EXP in the outer label when the tunnel is an LSP, the DSCP in the IP header if the tunnel encapsulation is GRE, and the DSCP in the payload’s IP header if the user enabled the ler-use-dscp option and the pseudowire terminates in IES or VPRN service (spoke-interface).

When the policer queue-group name the pseudowire is redirected does not exist, the redirection command is failed. In this case, the packet classification is performed according to default classification rule or the QoS filters defined in the ingress context of the network QoS policy applied to the network IP interface the pseudowire packet is received on.

The no version of this command removes the redirection of the pseudowire to the queue-group.

Parameters *network-policy-id* — Specifies the network policy identification. The value uniquely identifies the policy on the system.

This variant of the command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **qos name** *network-policy-name* variant can be used in all configuration modes.

Values 1 to 65535

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

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

Values 1 to 65535

name *network-policy-name* — Specifies the network policy name. The value uniquely identifies the policy on the system, up to 64 characters.

l2pt-termination

Syntax **l2pt-termination** [**cdp**] [**dtp**] [**pagp**] [**stp**] [**udld**] [**vtp**]
no l2pt-termination

Context config>service>pw-template

Description This command enables Layer 2 Protocol Tunneling (L2PT) termination on a given SAP or spoke SDP. L2PT termination will be supported only for STP BPDUs. PDUs of other protocols will be discarded.

This feature can be enabled only if STP is disabled in the context of the given VPLS service.

Default no l2pt-termination

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 virtual trunk protocol.

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 limit-mac-move 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.

mac-pinning

Syntax [**no**] **mac-pinning**

Context config>service>pw-template

Description Enabling this command will disable re-learning of MAC addresses on other SAPs within the service. The MAC address will remain attached to a given SAP for duration of its age-timer.

The age of the MAC address entry in the FDB is set by the age timer. If **mac-aging** is disabled on a given VPLS service, any MAC address learned on a SAP or SDP with **mac-pinning** enabled will remain in the FDB on this SAP or SDP forever. Every event that would otherwise result in re-learning will be logged (MAC address; original-SAP; new-SAP).

Note: For 7750 SR and 7450 ESS, MAC addresses learned during DHCP address assignment (DHCP snooping enabled) are not impacted by this command. MAC-pinning for such addresses is implicit.

Default When a SAP or spoke SDP is part of a Residential Split Horizon Group (RSHG), MAC pinning is activated at creation of the SAP. Otherwise MAC pinning is not enabled by default.

max-nbr-mac-addr

Syntax **max-nbr-mac-addr** *table-size*
no max-nbr-mac-addr

Context config>service>pw-template

Description This command specifies the maximum number of FDB entries for both learned and static MAC addresses for this SAP or spoke SDP.

When the configured limit has been reached, and **discard-unknown-source** has been enabled for this SAP or spoke SDP (see discard-unknown-source), 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 to 511999

restrict-protected-src

Syntax **restrict-protected-src** [**alarm-only** | **discard-frame**]
no restrict-protected-src

Context config>service>pw-template
config>service>pw-template>split-horizon-group

Description This command indicates how the agent will handle relearn requests for protected MAC addresses, either manually added using the mac-protect command or automatically added using the auto-learn-mac-protect command. While enabled all packets entering the configured SAP, spoke SDP, mesh SDP, or any SAP that is part of the configured split horizon group (SHG) will be verified not to contain a protected source MAC address. If the packet is found to contain such an address, the action taken depends on the parameter specified on the restrict-protected-src command, namely:

- No parameter

  The packet will be discarded, an alarm will be generated and the SAP, spoke SDP or mesh SDP will be set operationally down. The SAP, spoke SDP or mesh SDP must be shutdown and enabled (**no shutdown**) for this state to be cleared.
- alarm-only

  The packet will be forwarded, an alarm will be generated but the source MAC is not learned on the SAP, spoke SDP or mesh SDP.
- discard-frame

  The packet will be discarded and an alarm generated. The frequency of alarm generation is fixed to be at most one alarm per MAC address per FP per 10 minutes in a given VPLS service. This parameter is only applicable to automatically protected MAC addresses.

When the **restrict-protected-src** is enabled on an SHG, the action only applies to the associated SAPs (no action is taken by default for spoke SDPs in the SHG) and is displayed in the SAP show output as the oper state unless it is overridden by the configuration of **restrict-protected-src** on the SAP itself. In order to enable this function for spoke SDPs within a SHG, the **restrict-protected-src** must be enabled explicitly under the spoke SDP. If required, **restrict-protected-src** can also be enabled explicitly under specific SAPs within the SHG.

When this command is applied or removed, with either the alarm-only or discard-frame parameters, the MAC addresses are cleared from the related object.

The use of **restrict-protected-src discard-frame** is mutually exclusive with the configuration of manually protected MAC addresses within a given VPLS.

The **alarm-only** parameter is not supported on the 7750 SR-a, 7750 SR-1e/2e/3e, or 7950 XRS.

Default no restrict-protected-src

Parameters **alarm-only** — Specifies that the packet will be forwarded, an alarm will be generated but the source MAC is not learned on the SAP/spoke SDP/mesh SDP. This parameter is not supported on the 7950 XRS.

Default no alarm-only

discard-frame — Specifies that the packet will be discarded and an alarm generated. The frequency of alarm generation is fixed to be at most one alarm per FP per MAC address per 10 minutes within a given VPLS service.

Default no discard-frame

sdp-exclude

Syntax [**no**] **sdp-exclude** *group-name*

Context config>service>pw-template

Description This command configures SDP admin group constraints for a pseudowire template.

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

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

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

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

In the SDP selection process, all provisioned SDPs with the correct far-end IP address, the correct tunnel-far-end IP address, and the correct service label signaling are considered. The SDP with the lowest admin metric is selected. If more than one SDP with the same lowest metric are found then the SDP with the highest sdp-id is selected. The type of SDP, GRE or MPLS (BGP/RSVP/LDP) is not a criterion in this selection.

The selection rule with SDP admin groups is modified such that the following admin-group constraints are applied upfront to prune SDPs that do not comply:

- if one or more **sdp-include** statement is part of the pw-template, then an SDP that is a member of one or more of the included groups will be considered. With the **sdp-include** statement, there is no preference for an SDP that belongs to all included groups versus one that belongs to one or fewer of the included groups. All SDPs satisfying the admin-group constraint will be considered and the selection above based on the lowest metric and highest sdp-id is applied.
- if one or more **sdp-exclude** statement is part of the pw-template, then an sdp that is a member of any of the excluded groups will not be considered.

SDP admin group constraints can be configured on all router services that makes use of the pseudowire template (BGP-AD VPLS service, BGP-VPLS service, and FEC129 VLL service). In the latter case, only support at a T-PE node is provided.

The **no** form of this command removes the SDP admin group constraints from the pseudowire template.

Default none

Parameters *group-name* — Specifies the name of the SDP admin group. A maximum of 32 characters can be entered.

sdp-include

Syntax [**no**] **sdp-include** *group-name*

Context config>service>pw-template

Description This command configures SDP admin group constraints for a pseudowire template.

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

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

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

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

In the SDP selection process, all provisioned SDPs with the correct far-end IP address, the correct tunnel-far-end IP address, and the correct service label signaling are considered. The SDP with the lowest admin metric is selected. If more than one SDP with the same lowest metric are found then the SDP with the highest sdp-id is selected. The type of SDP, GRE or MPLS (BGP/RSVP/LDP) is not a criterion in this selection.

The selection rule with SDP admin groups is modified such that the following admin-group constraints are applied upfront to prune SDPs that do not comply:

- if one or more **sdp-include** statement is part of the pw-template, then an SDP that is a member of one or more of the included groups will be considered. With the **sdp-include** statement, there is no preference for an SDP that belongs to all included groups versus one that belongs to one or fewer of the included groups. All SDPs satisfying the admin-group constraint will be considered and the selection above based on the lowest metric and highest sdp-id is applied.
- if one or more **sdp-exclude** statement is part of the pw-template, then an sdp that is a member of any of the excluded groups will not be considered.

SDP admin group constraints can be configured on all router services that make use of the pseudowire template (BGP-AD VPLS service, BGP-VPLS service, and FEC129 VLL service). In the latter case, only support at a T-PE node is provided.

The **no** form of this command removes the SDP admin group constraints from the pseudowire template.

Default none

Parameters *group-name* — Specifies the name of the SDP admin group. A maximum of 32 characters can be entered.

split-horizon-group

Syntax **split-horizon-group** *group-name*
no split-horizon-group

Context config>service>pw-template

Description This command creates a new split horizon group (SGH).

Comparing a “residential” SGH and a “regular” SHG is that a residential SHG:

- Has different defaults for the SAP or SDP that belong to this group (ARP reply agent enabled (SAP only), MAC pinning enabled). These can be disabled in the configuration.
- Does not allow enabling spanning tree (STP) on a SAP. It is allowed on an SDP.
- Does not allow for downstream broadcast (broadcast/unknown unicast) on a SAP. It is allowed on an SDP.
- On a SAP, downstream multicast is only allowed when IGMP is enabled (for which an MFIB state exists; only IP multicast); on a SDP, downstream mcast is allowed.

When the feature was initially introduced, residential SHGs were also using ingress shared queuing by default to increase SAP scaling.

A residential SAP (SAP that belongs to a RSHG) is used to scale the number of SAPs in a single VPLS instance. The limit depends on the hardware used and is higher for residential SAPs (where there is no need for egress multicast replication on residential SAPs) than for regular SAPs. Therefore, residential SAPs are useful in residential aggregation environments (for example, triple play networks) with a VLAN/subscriber model.

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

Default A split horizon group is by default not created as a residential-group.

Parameters *group-name* — Specifies the name of the split horizon group to which the SDP belongs.

residential-group — Defines a split horizon group as a residential split horizon group (RSHG). Doing so entails that:

- SAPs which are members of this Residential Split Horizon Group will have:
  - Double-pass queuing at ingress as default setting (can be disabled)
  - STP disabled (cannot be enabled)
  - ARP reply agent enabled per default (can be disabled)
  - MAC pinning enabled per default (can be disabled)
  - Downstream Broadcast packets are discarded thus also blocking the unknown, flooded traffic
  - Downstream Multicast packets are allowed when IGMP snooping is enabled
- Spoke SDPs which are members of this Residential Split Horizon Group will have:
  - Downstream multicast traffic supported
  - Double-pass queuing is not applicable
  - STP is disabled (can be enabled)
  - ARP reply agent is not applicable on the 7750 SR and 7450 ESS (dhcp-lease-states are not supported on spoke SDPs)
  - MAC pinning enabled per default (can be disabled)

auto-learn-mac-protect

Syntax [**no**] **auto-learn-mac-protect**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls >mesh-sdp
config>service>vpls>split-horizon-group
config>service>vpls>endpoint
config>service>pw-template
config>service>pw-template>split-horizon-group

Description This command enables the automatic protection of source MAC addresses learned on the associated object. MAC protection is used in conjunction with restrict-protected-src, restrict-unprotected-dst and mac-protect. When this command is applied or removed, the MAC addresses are cleared from the related object.

When the auto-learn-mac-protect is enabled on an SHG the action only applies to the associated SAPs (no action is taken by default for spoke SDPs in the SHG). In order to enable this function for spoke SDPs within a SHG, the auto-learn-mac-protect must be enabled explicitly under the spoke-SDP. If required, auto-learn-mac-protect can also be enabled explicitly under specific SAPs within the SHG. For more information about auto-learn MAC protect, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN*.

Default no auto-learn-mac-protect

restrict-protected-src

Syntax **restrict-protected-src** [**alarm-only** | **discard-frame**]
no restrict-protected-src

Context config>service>pw-template
config>service>pw-template>split-horizon-group

Description This command indicates how the agent will handle relearn requests for protected MAC addresses, either manually added using the mac-protect command or automatically added using the auto-learn-mac-protect command. While enabled all packets entering the configured SAP, spoke SDP, mesh SDP, or any SAP that is part of the configured split horizon group (SHG) will be verified not to contain a protected source MAC address. If the packet is found to contain such an address, the action taken depends on the parameter specified on the restrict-protected-src command, namely:

- No parameter

  The packet will be discarded, an alarm will be generated and the SAP, spoke SDP or mesh SDP will be set operationally down. The SAP, spoke SDP or mesh SDP must be shutdown and enabled (no shutdown) for this state to be cleared.
- alarm-only

  The packet will be forwarded, an alarm will be generated but the source MAC is not learned on the SAP/spoke SDP/mesh SDP.
- discard-frame

  The packet will be discarded and an alarm generated. The frequency of alarm generation is fixed to be at most one alarm per MAC address per FP per 10 minutes in a given VPLS service. This parameter is only applicable to automatically protected MAC addresses.

When the **restrict-protected-src** is enabled on an SHG, the action only applies to the associated SAPs (no action is taken by default for spoke SDPs in the SHG) and is displayed in the SAP show output as the oper state unless it is overridden by the configuration of **restrict-protected-src** on the SAP itself. In order to enable this function for spoke SDPs within a SHG, the **restrict-protected-src** must be enabled explicitly under the spoke SDP. If required, **restrict-protected-src** can also be enabled explicitly under specific SAPs within the SHG.

When this command is applied or removed, with either the alarm-only or discard-frame parameters, the MAC addresses are cleared from the related object.

The use of **restrict-protected-src discard-frame** is mutually exclusive with the configuration of manually protected MAC addresses within a given VPLS.

The **alarm-only** parameter is not supported on the 7750 SR-a, 7750 SR-1e/2e/3e, or 7950 XRS.

Default no restrict-protected-src

Parameters **alarm-only** — Specifies that the packet will be forwarded, an alarm will be generated but the source MAC is not learned on the SAP, spoke SDP and mesh SDP. This parameter is not supported on the 7950 XRS.

Default no alarm-only

discard-frame — Specifies that the packet will be discarded and an alarm generated. The frequency of alarm generation is fixed to be at most one alarm per FP per MAC address per 10 minutes within a given VPLS service.

Default no discard-frame

restrict-unprotected-dst

Syntax [**no**] **restrict-unprotected-dst**

Context config>service>pw-template>split-horizon-group

Description This command indicates how the system will forward packets destined to an unprotected MAC address, either manually added using the mac-protect command or automatically added using the auto-learn-mac-protect command. While enabled all packets entering the configured SAP or SAPs within a split-horizon-group (but not spoke or mesh-SDPs) will be verified to contain a protected destination MAC address. If the packet is found to contain a non-protected destination MAC, it will be discarded. Detecting a non-protected destination MAC on the SAP will not cause the SAP to be placed in the operationally down state. No alarms are generated.

If the destination MAC address is unknown, even if the packet is entering a restricted SAP, with restrict-unprotected-dst enabled, it will be flooded.

Default no restrict-unprotected-dst

stp

Syntax **stp**

Context config>service>pw-template

Description This command enables the context to configure the Spanning Tree Protocol (STP) parameters. The 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 service routers should not be blocked, the root path is calculated from the core perspective.

auto-edge

Syntax [**no**] **auto-edge**

Context config>service>pw-template>stp

Description This command configures automatic detection of the edge port characteristics of the SAP or spoke SDP.

If auto-edge is enabled, and STP concludes there is no bridge behind the spoke SDP, the OPER_EDGE variable will dynamically be set to true. If auto-edge is enabled, and a BPDU is received, the OPER_EDGE variable will dynamically be set to true (see edge-port).

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>pw-template>stp

Description This command configures the SAP or SDP as an edge or non-edge port. If **auto-edge** is enabled for the SAP, this value will be used only as the initial value.

Note: On the 7750 SR and the 7950 XRS, the function of the **edge-port** command is similar to the **rapid-start** command. It tells RSTP that it is on the edge of the network (for example, there are no other bridges connected to that port) and, as a consequence, it can immediately transition to a forwarding state if the port becomes available.

RSTP, however, can detect that the actual situation is different from what **edge-port** may indicate.

Initially, the value of 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

link-type

Syntax **link-type** {**pt-pt** | **shared**}
no link-type

Context config>service>pw-template>stp

Description This command instructs STP on the maximum number of bridges behind this SAP or spoke SDP. If there is only a single bridge, transitioning to forwarding state will be based on handshaking (fast transitions). If more than two bridges are connected via a shared media, their SAP or spoke SDPs should all be configured as shared, and timer-based transitions are used.

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

Default link-type pt-pt

path-cost

Syntax **path-cost** *sap-path-cost*
no path-cost

Context config>service>pw-template>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 or spoke SDP. When BPDUs are sent out other egress SAPs or spoke SDPs, the newly calculated root path cost is used. These are the values used for CIST when running MSTP.

STP suggests that the path cost is defined as a function of the link bandwidth. Since SAPs and spoke SDPs are controlled by complex queuing dynamics, the STP path cost is a purely static configuration.

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

Default path-cost 10

Parameters *path-cost* — Specifies the path cost for the SAP or spoke SDP.

Values 1 to 200000000 (1 is the lowest cost)

Default 10

priority

Syntax **priority** *bridge-priority*
no priority

Context config>service>pw-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* — Specifies the bridge priority for the STP instance.

Values Allowed values are integers in the range of 4096 to 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.

root-guard

Syntax [**no**] **root-guard**

Context config>service>pw-template>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

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 **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.

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 top VLAN tag, if a VLAN tag is present, is stripped from traffic received on the pseudowire, and a vlan-tag is inserted when forwarding into the pseudowire. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings.

Note: The system expects a symmetrical configuration with its peer, specifically it expects to remove the same number of VLAN tags from received traffic as it adds to transmitted traffic. As some of the related configuration parameters are local and not communicated in the signaling plane, an asymmetrical behavior cannot always be detected and so cannot be blocked. Consequently, protocol extractions will not necessarily function for asymmetrical configurations as they would with a symmetrical configurations resulting in an unexpected operation.

vlan-vc-tag

Syntax **vlan-vc-tag** *vlan-id*
no vlan-vc-tag

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 *vlan-id* — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

Values 0 to 4094

2.19.2.6.1 PW Port Commands

binding

Syntax **binding**

Context config>service>sdp

Description The command enables the context to configure SDP bindings.

port

Syntax **port** [*port-id* **|** *lag-id*]
no port

Context config>service>sdp>binding

Description This command specifies the port or lag identifier, to which the pseudowire ports associated with the underlying SDP are bound. If the underlying SDP is re-routed to a port or lag other than the specified one, the pseudowire ports on the SDP are operationally brought down.

The **no** form of the command removes the value from the configuration.

Default none

Parameters *port-id* — Specifies the identifier of the port in the slot/mda/port format.

port-id	*slot*/*mda*/*port*[.*channel*]	
	pxc-id	psc-id.sub-port
		pxc psc-id.sub-port
		pxc: keyword
		id: 1 to 64
		sub-port: a, b
	aps-id	aps-*group-id*[.*channel*]

	aps keyword	
	group-id	1 to 64
	group-id	1 to 16
bundle-*type-slot/mda*.*bundle-num*		
	bundle	keyword
	type	ima, ppp
	bundle-num	1 to 256
bpgrp-id:	bpgrp-*type-bpgrp-num*	
	bpgrp	keyword
	type	ima
	bpgrp-num	1 to 1280
ccag-id	- ccag-<id>.<path-id>[cc-type]	
	ccag	keyword
	id	1 to 8
	path-id	a, b
	cc-type[.sap-net | .net-sap]	
lag-id	lag-*id*	
	lag	keyword
	id	1 to 800

lag-id — Specifies the LAG identifier.

pw-port

Syntax **pw-port** *pw-port-id* [**vc-id** *vc-id*] [**create**]
no pw-port *pw-port-id*

Context config>service>sdp>binding

Description This command creates a pseudowire port.

The **no** form of the command removes the pseudowire port ID from the configuration.

Default none

Parameters *pw-port-id* — Specifies a unique identifier of the pseudowire port.

Values 1 to 10239

vc-id — Specifies a virtual circuit identifier signaled to the peer.

Values 1 to 4294967295

create — This keyword is required when a new pseudowire is being created.

egress

Syntax **egress**

Context config>service>sdp>binding>pw-port

Description This command enters egress configuration context for the vport.

Default none

shaper

Syntax [**no**] **shaper**

Context config>service>sdp>binding>pw-port>egress

Description This command enables the egress shaping option for use by a pseudowire port.

The **no** form of the command disables the egress shaping option.

Default no shaper

int-dest-id

Syntax **int-dest-id** *int-dest-id*
no int-dest-id

Context config>service>sdp>binding>pw-port>egress>shaper

Description This command configures an intermediate destination identifier applicable to ESM PW SAPs.

The **no** form of the command removes the intermediate destination identifier from the configuration.

Parameters *int-dest-id* — Specifies the intermediate destination ID.

pw-sap-secondary-shaper

Syntax **pw-sap-secondary-shaper** *pw-sap-sec-shaper-name*
no pw-sap-secondary-shaper

Context config>service>sdp>binding>pw-port>egress>shaper

Description This command configures a default secondary shaper applicable to pw-saps under normal interfaces.

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

vport

Syntax **vport** *vport-name*
no vport

Context config>service>sdp>binding>pw-port>egress>shaper

Description This command configures a virtual port applicable to all pw-saps.

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

Parameters *vport-name* — Specifies the name up to 32 characters in length.

vc-label

Syntax **vc-label** *vc-label*
no vc-label

Context config>service>sdp>binding>pw-port>egress

Description This command configures the egress VC label for the PW representing the PW-port.

Default no vc-label

Parameters *vc-label* — Specifies the VC egress value that indicates a specific connection.

Values 16 to 1048575

ingress

Syntax **ingress**

Context config>service>sdp>binding>pw-port

Description This command configures ingress parameters for the PW port.

vc-label

Syntax **vc-label** *ingress-vc-label*
no vc-label

Context config>service>sdp>binding>pw-port>ingress

Description This command configures the ingress VC label used for the PW representing the PW port.

Note that the maximum value of the vc-label that may be configured is limited by the **config**>**router**>**mpls-labels**>**static-label-range** command.

Default no vc-label

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

Values 32 to 18431

monitor-oper-group

Syntax **monitor-oper-group** *group name*
no monitor-oper-group

Context config>service>sdp>binding>pw-port

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

The **no** form of the command removes the association from the configuration.

Default no monitor-oper-group

Parameters *name* — Specifies a character string of maximum 32 ASCII characters identifying the group instance.

vc-type

Syntax **vc-type** {**ether | vlan**}
no vc-type

Context config>service>sdp>binding>pw-port

Description This command sets the forwarding mode for the pseudowire port. The vc-type is signaled to the peer, and must be configured consistently on both ends of the pseudowire. vc-type VLAN is only configurable with dot1q encapsulation on the pseudowire port. The tag with vc-type vlan only has significance for transport, and is not used for service delineation or ESM. The top (provider tag) is stripped while forwarding out of the pseudowire, and a configured vlan-tag (for vc-type vlan) is inserted when forwarding into the pseudowire. With vc-type ether, the tags if present (max 2), are transparently preserved when forwarding in our out of the pseudowire.

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

Default vc-type ether

Parameters **ether** — Specifies **ether** as the virtual circuit (VC) associated with the SDP binding.

vlan — Specifies **vlan** as the virtual circuit (VC) associated with the SDP binding.

vlan-vc-tag

Syntax **vlan-vc-tag** *vlan-id*
no vlan-vc-tag

Context config>service>sdp>binding>pw-port

Description This command sets tag relevant for vc-type vlan mode. This tag is inserted in traffic forwarded into the pseudowire.

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

Default no vlan-vc-tag

Parameters *vlan-id* — Specifies the VLAN ID value.

Values 0 to 4094

2.19.2.7 SDP Commands

sdp

Syntax **sdp** *sdp-id* [*delivery-type*] [**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 router to a particular service without having to specifically define far-end SAPs. Each SDP represents a method to reach another router.

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

The second method is Multi-Protocol Label Switching (MPLS) encapsulation. A router 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.

Segment routing is another MPLS tunnel type and is used to allow service binding to an SR tunnel programmed in TTM by OSPF or IS-IS. The SDP of type **sr-isis** or **sr-ospf** can be used with the **far-end** option. The **tunnel-farend** option is not supported. In addition, the **mixed-lsp-mode** option does not support the **sr-isis** and **sr-ospf** tunnel types.

L2TPv3-over-IPv6 transport is also an option for 7750 SR and 7950 XR Ethernet Pipe (Epipe) Services. Like GRE, L2TPv3 is stateless in the core of the network, as well as on the service nodes as the L2TPv3 control plane functionality is disabled for this SDP type. A unique source and destination IPv6 address combined with TX and RX Cookie values are used to ensure that the SDP is bound to the correct service.

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 the *sdp-id* does not exist, a new SDP is created. When creating an SDP, either the **gre, mpls, or l2tpv3** keyword must be specified. SDPs are created in the admin down state (**shutdown**) and the **no shutdown** command must be executed once all relevant parameters are defined and before the SDP can be used.

If *sdp-id* exists, the current CLI context is changed to that SDP for editing and modification. For editing an existing SDP, neither the **gre, mpls**, or **l2tpv3** 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* — Specifies the SDP identifier.

Values 1 to 17407

gre — Specifies the SDP will use GRE to reach the far-end router. The GRE encapsulation of the MPLS service packet uses the base 4-byte header as per RFC 2890. The optional fields Checksum (plus Reserved field), Key, and Sequence Number are not inserted. Only one GRE SDP can be created to a given destination address. Multiple GRE SDPs to a single destination address serve no purpose as the path taken to reach the far end is determined by the IGP which will be the same for all SDPs to a given destination and there is no bandwidth reservation in GRE tunnels.

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

l2tpv3 — Specifies the SDP will use L2TPv3-over-IPv6 encapsulation for the 7750 SR or 7950 XRS. One SDP is created per service, regardless of whether the far-end node is common or not. Unique local and far-end addresses are configured for every L2TPv3 SDP type. The local address must exist on the local node.

eth-gre-bridged — Configures the SDP as an L2oGRE tunnel that is terminated on an FPE-based PW port. Only the end-points of such a tunnel (the far-end IPv4/IPv6 address or local-end IPv4/IPv6 address) are allowed to be configured under this SDP.

accounting-policy

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

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

Description This command creates the accounting policy context that can be applied to an SDP. An accounting policy must be defined before it can be associated with a SDP. If the *acct-policy-id* does not exist, an error message is generated.

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

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

Default no accounting-policy

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

Values 1 to 99

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 L2 services using this SDP-ID. When enabled, the router signals a VC MTU equal to the service MTU, which includes the Layer 2 header. It also allows this router to accept an MTU advertised by the far-end PE which value matches either its advertised MTU or its advertised MTU minus the L2 headers.

By default, the router advertises a VC-MTU equal to the L2 service MTU minus the Layer 2 header and always matches its advertised MTU to that signaled by the far-end PE router, otherwise the spoke-sdp goes operationally down.

When this command is enabled on the SDP, it has no effect on a spoke-sdp of an IES/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 - L2 headers}. The router also continues to make sure that the advertised 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

allow-fragmentation

Syntax [**no**] **allow-fragmentation**

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

Description This command disables the setting of the **do-not-fragment** bit in the IP header of GRE encapsulated service traffic. This feature is only applicable to GRE SDPs and will be applied to all service traffic using the associated GRE SDP.

The **no** form of this command removes the command from the active configuration and returns the associated SDP to its default which is to set the **do-not-fragment** bit in all GRE encapsulated service traffic.

Default no allow-fragmentation

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 BGP, SR-ISIS, or SR-OSPF. This restriction is relaxed for some combinations of the transport methods when the mixed-lsp-mode option is enabled within the SDP.

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)

booking-factor

Syntax **booking-factor** *percentage*
no booking-factor

Context config>service>sdp

Description This command specifies the booking factor applied against the maximum SDP available bandwidth by the VLL CAC feature.

The service manager keeps track of the available bandwidth for each SDP. The maximum value is the sum of the bandwidths of all constituent LSPs in the SDP. The SDP available bandwidth is adjusted by the user configured booking factor. A value of 0 means no VLL can be admitted into the SDP.

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

Default no booking-factor

Parameters *percentage* — Specifies the percentage of the SDP maximum available bandwidth for VLL call admission. When the value of this parameter is set to zero (0), no new VLL spoke SDP bindings with non-zero bandwidth are permitted with this SDP. Overbooking, >100% is allowed.

Values 0 to 1000%

Default 100

class-forwarding

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

Context config>service>sdp

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

VLL service packets are forwarded based on their forwarding class only if shared queuing is enabled on the ingress SAP. Shared queuing must be enabled on the VLL ingress SAP if class-forwarding is enabled on the SDP the service is bound to. Otherwise, the VLL packets will be forwarded to the LSP which is the result of hashing the VLL service ID. Since there are eight entries in the ECMP table for an SDP, one LSP ID for each forwarding class, the resulting load balancing of VLL service ID is weighted by the number of times an LSP appears on that table. For instance, if there are eight LSPs, the result of the hashing will be similar to when class based forwarding is disabled on the SDP. If there are fewer LSPs, then the LSPs which were mapped to more than one forwarding class, including the default LSP, will have proportionally more VLL services forwarding to them.

Class-based forwarding is not supported on a spoke SDP used for termination on an IES or VPRN service. All packets are forwarded over the default LSP.

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

Default no class-forwarding

Parameters **default-lsp** *lsp-name* — Specifies the default LSP for the SDP. This LSP name must exist and must have been associated with this SDP using the *lsp-name* configured in the **config>service>sdp>lsp** context. The default LSP is used to forward packets when there is no available LSP which matches the packet's forwarding class. This could be because the LSP associated with the packet's forwarding class is down, or that the user did not configure a mapping of the packet's forwarding class to an LSP using the **config>service>sdp>class-forwarding>fc** command. The default LSP is also used to forward VPLS service multicast/broadcast packets in the absence of a user configuration indicating an explicit association to one of the SDP LSPs.

Note: When the default LSP is down, the SDP is also brought down. The user will not be able to enter the class-forwarding node if the default LSP was not previously specified. In other words, the class-forwarding for this SDP will remain shutdown.

enforce-diffserv-lsp-fc

Syntax [**no**] **enforce-diffserv-lsp-fc**

Context config>service>sdp>class-forwarding

Description This command enables checking by RSVP that a Forwarding Class (FC) mapping to an LSP under the SDP configuration is compatible with the Diff-Serv Class Type (CT) configuration for this LSP.

When the user enables this option, the service manager inquires with RSVP if the FC is supported by the LSP. RSVP checks if the FC maps to the CT of the LSP, for example, the default class-type value or the class-type value entered at the LSP configuration level.

If RSVP did not validate the FC, then the service manager will return an error and the check has failed. In this case, packets matching this FC will be forwarded over the default LSP. Any addition of an LSP to an SDP that will not satisfy the FC check will also be rejected.

The service manager does no validate the default-lsp FC-to-CT mapping. Whether or not the FC is validated, the default-lsp will always end up being used in this case.

RSVP will not allow the user to change the CT of the LSP until no SDP with class-based forwarding enabled and the **enforce-diffserv-lsp-fc** option enabled is using this LSP. All other SDPs using this LSP are not concerned by this rule.

The SDP will continue to enforce the mapping of a single LSP per FC. However, when **enforce-diffserv-lsp-fc** enabled, RSVP will also enforce the use of a single CT per FC as per the user configured mapping in RSVP.

If class-forwarding is enabled but **enforce-diffserv-lsp-fc** is disabled, forwarding of the service packets will continue to be based on the user entered mapping of FC to LSP name without further validation as per the existing implementation. The CT of the LSP does not matter in this case.

If class-forwarding is not enabled on the SDP, forwarding of the service packets will continue to be based on the ECMP/LAG hash routine. The CT of the LSP does not matter in this case.

The **no** form of this command reverts to the default value which is to use the user entered mapping of FC to LSP name.

Default no enforce-diffserv-lsp-fc

fc

Syntax **fc** {*fc*} **lsp** *lsp-name*
no fc {fc}

Context config>service>sdp>class-forwarding

Description This command makes an explicit association between a forwarding class and an LSP. The LSP name must exist and must have been associated with this SDP using the command config>service>sdp>lsp. Multiple forwarding classes can be associated with the same LSP. However, a forwarding class can only be associated with a single LSP in a given SDP. All subclasses will be assigned to the same LSP as the parent forwarding class.

Default none

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

fc — Specifies a forwarding class to LSP mapping.

Values be, l2, af, 1, h2, ef, h1, nc

multicast-lsp

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

Context config>service>sdp>class-forwarding

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

Default no multicast-lsp — traffic mapped to default-lsp *name*

collect-stats

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

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

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

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM or XCM 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

far-end

Syntax **far-end node-id** *node-id* [**global-id** *global-id*]
far-end [*ip-address* | *ipv6-address*]
no far-end *ip-address* | *ipv6-address*

Context config>service>sdp

Description This command configures the system IP address of the far-end destination 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 that of an SR OS and for a GRE SDP it must match the system IP address of the far end router.

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

If the SDP uses MPLS encapsulation, the **far-end** 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** address, the LSP will not be added to the SDP and an error will be generated. Alternatively, an SDP that uses MPLS can have an MPLS-TP node with an MPLS-TP node-id and (optionally) a global ID. In this case, the SDP must use an MPLS-TP LSP and the SDP **signaling** parameter must be set to **off**.

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

On a tunnel configured as SDP with delivery type of eth-gre-bridged, this command designates L2oGRE tunnel end points. This is the only configuration option allowed for this type of SDP.

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 **far-end** — Specifies the far-end termination point for the GRE tunnel.

ip-address | *ipv6-address* — Specifies a IPv4 or IPv6 address of the far-end SR OS for the SDP in dotted decimal notation.

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

global-id — Specifies a MPLS-TP Global ID of the far-end system for the SDP, in an unsigned 32-bit integer (0 to 4294967295). This parameter is optional for an SDP using an MPLS-TP LSP. If not 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.

keep-alive

Syntax **keep-alive**

Context config>service>sdp

Description This command enables the context to configure SDP connectivity monitoring keepalive messages for the SDP ID.

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

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

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

Table 14 describes the keepalive interpretation of SDP echo reply response conditions and the effect on the SDP ID operational status.

Table 14 **Keepalive Interpretation and Effect of SDP Echo Reply**

Result of Request	Stored Response State	Operational State
keepalive request timeout without reply	Request Timeout	Down

Table 14 **Keepalive Interpretation and Effect of SDP Echo Reply (Continued)**

Result of Request	Stored Response State	Operational State
keepalive request not sent due to non-existent *orig-sdp-id* (This condition should not occur)	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)

hello-time

Syntax [**no**] **hello-time** *seconds*

Context config>service>sdp>keep-alive

Description This command configures the time period between SDP keepalive messages on the SDP-ID for the SDP connectivity monitoring messages.

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

Default hello-time 10

Parameters *seconds* — Specifies the time period in seconds between SDP keepalive messages, expressed as a decimal integer.

Values 1 to 3600

hold-down-time

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

Context config>service>sdp>keep-alive

Description This command 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

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

Values 1 to 5

message-length

Syntax **message-length** *message-length*
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 no message-length — 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.

Parameters *message-length* — Specifies 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 to 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 timeout 5

Parameters **timeout** — Specifies the timeout time, in seconds.

Values 1 to 10

ldp

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

Context config>service>sdp

Description This command enables LDP-signaled LSPs on MPLS-encapsulated SDPs.

In MPLS SDP configurations *either* one or more LSP names can be specified *or* LDP can be enabled. The SDP **ldp** and **lsp** commands are mutually exclusive except if the mixed-lsp-mode option is also enabled. 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 or the mixed-lsp-mode option is also enabled.

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)

local-end

Syntax **local-end** {*ip-address* | *ipv6-address*}
no local-end

Context config>service>sdp

Description This command configures the local-end address of the following SDP encapsulation types:

- IPv6 address of the termination point of a SDP of encapsulation **l2tpv3** (L2TP v3 tunnel).
- IPv4/IPv6 source address of a SDP of encapsulation **eth-gre-bridged** (L2oGRE SDP).
- IPv4 source address of a SDP of encapsulation **gre** (GRE SDP).

A change to the value of the local-end parameter requires that the SDP be shut down.

When used as the source address of a SDP of encapsulation **gre** (GRE SDP), the primary IPv4 address of any local network IP interface, loopback or otherwise, may be used.

The address of the following interfaces are not supported:

- unnumbered network IP interface
- IES interface
- VPRN interface
- CSC VPRN interface

The **local-end** parameter value adheres to the following rules:

- A maximum of 15 distinct address values can be configured for all GRE SDPs under the **configure**>**service**>**sdp**>**local-end** context, and all L2oGRE SDPs under the **config**>**service**>**system**>**gre-eth-bridged**>**tunnel-termination** context.
- The same source address cannot be used in both contexts since an address configured for a L2oGRE SDP matches an internally created interface that is not available to other applications.
- The **local-end** address of a GRE SDP, when different from system, need not match the primary address of an interface that has the MPLS-over-GRE termination subnet configured, unless a GRE SDP or tunnel from the far-end router terminates on this address.

The **no** form of the command removes the address from the local-end configuration.

Parameters *ip-address* | *ipv6-address* — Specifies a IPv4 or IPv6 address for local-end of an SDP in dotted decimal notation.

Values

ip-address	a.b.c.d
ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0..FFFF]H
	d - [0..255]D

lsp

Syntax [**no**] **lsp** *lsp-name*

Context config>service>sdp

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

In MPLS SDP configurations *either* one or more LSP names can be specified *or* LDP can be enabled. The SDP **ldp** and **lsp** commands are mutually exclusive except if the mixed-lsp-mode option is also enabled. 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 or the mixed-lsp-mode option is also enabled. 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).

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

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

Default none

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

metric

Syntax **metric** *metric*
no metric

Context config>service>sdp

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

Parameters *metric* — Specifies the SDP metric.

Values 0 to 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:

- RSVP LSP type. Up to 16 RSVP LSPs can be entered by the user and programmed by the service manager in ingress linecard to load balance service packets. This is the highest priority LSP type.
- LDP LSP type. One LDP FEC programmed by the service manager but the ingress card can use up to 16 LDP ECMP paths for the FEC to load balance service packets when ECMP is enabled on the node.
- BGP LSP type. One RFC 3107-labeled BGP prefix programmed by the service manager. The ingress card can use more than one next-hop for the prefix.

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 card 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 card 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: 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 **config>router>ldp>tunnel-down-damp-time** command.

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.

If class based forwarding is enabled for this SDP, the forwarding of the packets over the RSVP LSPs will be based on the FC of the packet as in current implementation. When the SDP activates the LDP LSP type, then packets are forwarded over the LDP ECMP paths using the regular hash routine.

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 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 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 no revert-time

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

Values 0 to 600

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

network-domain

Syntax **network-domain** *network-domain-name*
no network-domain

Context config>service>sdp

Description This command assigns a given SDP to a given network-domain. The network-domain is then taken into account during sap-ingress queue allocation for VPLS SAP.

The network-domain association can only be done in a base-routing context. Associating a network domain with an loop-back or system interface will be rejected. Associating a network-domain with an interface that has no physical port specified will be accepted, but will have no effect as long as a corresponding port, or LAG, is undefined.

A single SDP can only be associated with a single network-domain.

Default network-domain "default"

path-mtu

Syntax **path-mtu** [*bytes*]
no path-mtu *bytes*

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

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

Default no path-mtu — The default path-mtu defined on the system for the type of SDP is used.

pbb-etype

Syntax **pbb-etype** *type*
no pbb-etype [*type*]

Context config>service>sdp

Description This command configures the Ethertype used for PBB.

Default 0x88E7

Parameters *type* — Specifies the Ethertype.

Values 0x0600..0xffff or 1536 to 65535 (accepted in decimal or hex)

signaling

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

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 signaling value can only be changed while the administrative status of the SDP is down. Additionally, the signaling can only be changed on an SDP if that SDP is not in use by BGP-AD or BGP-VPLS. BGP signaling can only be enabled if that SDP does not already have pseudowires signaled over it. Also, BGP signaling is not supported with mixed mode LSP SDPs.

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 signaling tldp

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

tldp — Ingress and egress pseudowire signaling using T-LDP is enabled. Default value used when BGP AD automatically instantiates the SDP.

bgp — Ingress and egress pseudowire signaling using BGP is enabled. Default value used when BGP VPLS automatically instantiates the SDP.

source-bmac-lsb

Syntax **source-bmac-lsb** *mac-lsb* **control-pw-vc-id** *vc-id*
no source-bmac-lsb

Context config>service>sdp

Description This command defines the 16 least significant bits (lsb) which, when combined with the 32 most significant bits of the PBB **source-bmac**, are used as the virtual backbone MAC associated with this SDP. The virtual backbone MAC is used as the source backbone MAC for traffic received on a PBB EPIPE spoke-SDP with **use-sdp-bmac** configured (that is, a redundant pseudowire) and forwarded into the B-VPLS domain.

The control-pw-vc-id defines VC identifier of the spoke-SDP relating to the control pseudowire whose status is to be used to determine whether SPBM advertises this virtual backbone MAC. This is a mandatory parameter when the **source-bmac-lsb** is added or changed. The spoke SDP must have the parameter **use-sdp-bmac** for the control pseudowire to be active.

Default no source-bmac-lsb

Parameters *mac-lsb* — Specifies the 16 least significant bits of the virtual backbone MAC associated with this SDP.

Values 1 to 65535 or xx-xx or xx:xx

control-pw-vc-id *vc-id* — Specifies the VC identifier of the control pseudowire.

Values 1 to 4294967295

sr-isis

Syntax [**no**] **sr-isis**

Context config>service>sdp

Description This command configures an MPLS SDP of LSP type ISIS Segment Routing. The SDP of LSP type sr-isis can be used with the far-end option. The signaling protocol for the service labels for an SDP using an SR tunnel can be configured to static (off), T-LDP (tldp), or BGP (bgp).

sr-ospf

Syntax [**no**] **sr-ospf**

Context config>service>sdp

Description This command configures an MPLS SDP of LSP type OSPF Segment Routing. The SDP of LSP type sr-ospf can be used with the far-end option. The signaling protocol for the service labels for an SDP using an SR tunnel can be configured to static (off), T-LDP (tldp), or BGP (bgp).

sr-te-lsp

Syntax [**no**] **sr-te-lsp** *lsp-name*

Context config>service>sdp

Description This command configures an MPLS SDP of LSP type SR-TE.

The user can specify up to 16 SR-TE LSP names. The destination address of all LSPs must match that of the SDP far-end option. Service data packets are sprayed over the set of LSPs in the SDP using the same procedures as for tunnel selection in ECMP. Each SR-TE LSP can, however, have up to 32 next-hops at the ingress LER when the first segment is a node SID-based SR tunnel. Thus service data packet will be forwarded over one of a maximum of 16x32 next-hops.

The **tunnel-far-end** option is not supported. In addition, the **mixed-lsp-mode** option does not support the **sr-te** tunnel type.

The signaling protocol for the service labels for an SDP using a SR-TE LSP can be configured to static (**off**), T-LDP (**tldp**), or BGP (**bgp**).

tunnel-far-end

Syntax **tunnel-far-end** *ip-address* | *ipv6-address*
no tunnel-far-end [*ip-address* | *ipv6-address*]

Context config>service>sdp

Description This command enables the user to specify an SDP tunnel destination address that is different from the configuration in the SDP far-end option.
The SDP must be shutdown first to add or change the configuration of the **tunnel-far-end** option.

When this option is enabled, service packets are encapsulated using an LDP LSP with a FEC prefix matching the value entered in ip-address. By default, service packets are encapsulated using an LDP LSP with a FEC prefix matching the address entered in the SDP far-end option.

The T-LDP session to the remote PE is still targeted to the address configured under the **far-end option**. This means that targeted hello messages are sent to the far-end address, which is also the LSR-ID of the remote node. TCP based LDP messages, such as initialization and label mapping messages, are sent to the address specified in the transport-address field of the "hello" message received from the remote PE. This address can be the same as the remote PE LSR-ID, or a different address. This feature works, however, if the signaling option in the SDP is set to off instead of tldp, in which case, the service labels are statically configured.

This feature operates on an SDP of type LDP only. It can be used with VLL, VPLS, and VPRN services when an explicit binding to an SDP with the **tunnel-far-end** is specified. It also operates with a spoke interface on an IES or VPRN service. Finally, this feature operates with a BGP AD based VPLS service when the **use-provisioned-sdp** option is enabled in the pseudowire template.

This feature is not supported in an SDP of type MPLS when an RSVP LSP name is configured under the SDP. It also does not work with a mixed-lsp SDP.

The **no** form of this command disables the use of the **tunnel-far-end** option and returns to using the address specified in the far-end.

Default no tunnel-far-end

Parameters *ip-address* **|** *ipv6-address* — Specifies the system address of the far-end router for the SDP in dotted decimal notation.

vlan-vc-etype

Syntax **vlan-vc-etype** *ethernet-type*
no vlan-vc-etype [*ethernet-type*]

Context config>service>sdp

Description This command configures the VLAN VC EtherType.

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

Default no vlan-vc-etype

Parameters *ethernet-type* — Specifies a valid VLAN etype identifier.

Values 0x0600 to 0xffff

weighted-ecmp

Syntax [**no**] **weighted-ecmp**

Context config>service>sdp

Description This command enables weighted ECMP on an SDP. When weighted ECMP is enabled, packets from services using the SDP are sprayed across LSPs in the ECMP set to the SDP far end according to the outcome of the hash algorithm and the configured load-balancing weight of each LSP.

The **no** version of this command disables weighted ECMP for next-hop tunnel selection.

Default no weighted-ecmp

sdp-group

Syntax **sdp-group**

Context config>service>sdp

Description This command configures the SDP membership in admin groups.

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

The **no** form of this command removes this SDP membership to the specified admin group.

Default none

group-name

Syntax **group-name** *group-name* **value** *group-value*
no group-name *group-name*

Context config>service>sdp-group

Description This command defines SDP administrative groups, referred to as SDP admin groups.

SDP admin groups provide a way for services using a pseudowire template to automatically include or exclude specific provisioned SDPs. SDPs sharing a specific characteristic or attribute can be made members of the same admin group. When users configure a pseudowire template, they can include and/or exclude one or more admin groups. When the service is bound to the pseudowire template, the SDP selection rules will enforce the admin group constraints specified in the **sdp-include** and **sdp-exclude** commands.

A maximum of 32 admin groups can be created. The group value ranges from zero (0) to 31. It is uniquely associated with the group name at creation time. If the user attempts to configure another group name for a group value that is already assigned to an existing group name, the SDP admin group creation is failed. The same happens if the user attempts to configure an SDP admin group with a new name but associates it to a group value already assigned to an existing group name.

The **no** option of this command deletes the SDP admin group but is only allowed if the group-name is not referenced in a pw-template or SDP.

Default none

Parameters *group-name* — Specifies the name of the SDP admin group. A maximum of 32 characters can be entered.

group-value — Specifies the group value associated with this SDP admin group. This value is unique within the system.

Values 0 to 31

2.19.2.8 Ethernet Ring Commands

eth-ring

Syntax **eth-ring** *ring-index*
no eth-ring

Context config

Description This command configures a G.8032 protected Ethernet ring. G.8032 Rings may be configured as major rings with two paths (a&b) or as Sub-Rings with two paths or in the case of an interconnection node a single path.

The **no** form of this command deletes the Ethernet ring specified by the ring-id.

Default no eth-ring

Parameters *ring-index* — Specifies the ring ID.

Values 1 to 128

ccm-hold-time

Syntax **ccm-hold-time** [**down** *down-timeout*] [**up** *up-timeout*]
no ccm-hold-time

Context config>eth-ring

Description This command configures eth-ring dampening timers. See the **down** and **up** commands for more information.

The **no** form of the command sets the up and down timers to the default values.

Parameters *down-timeout* — Specifies the down timeout, in centiseconds.

Values 0 to 5000

up-timeout — Specifies the hold-time for reporting the recovery, in deciseconds.

Values 0 to 5000

compatible-version

Syntax **compatible-version** *version*
no compatible-version

Context config>eth-ring

Description This command configures eth-ring compatibility version for the G.8032 state machine and messages. The default is version 2 and all router switches use version 2. If there is a need to interwork with third party devices that only support version 1 this can be set to version 1.

The **no** form of this command set the compatibility version to 2.

Default compatible-version 2

Parameters *version* — Specifies the version of the G.8032 state machine.

Values 1, 2

guard-time

Syntax **guard-time** *time*
no guard-time

Context config>eth-ring

Description This command configures the guard time for an Eth-Ring. The guard timer is standard and is configurable from “x” ms to 2 seconds.

The **no** form of this command restores the default guard-time.

Default no guard-time

Parameters *value* — Specifies the guard-time, in deciseconds.

Values 1 to 20

Default 5

node-id

Syntax **node-id** *mac-address*
no node-id

Context config>eth-ring

Description This optional command configures the MAC address of the RPL control. The default is to use the chassis MAC for the ring control. This command allows the chassis MAC to be overridden with another MAC address.

The **no** form of the command removes the RPL link.

Default no node-id

Parameters *mac-address* — xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

path

Syntax **path** {**a** | **b**} [{*port-id* | *lag-id*} **raps-tag** *qtag1*[.*qtag2*]]
no path {**a** | **b**}

Context config>eth-ring

Description This command assigns the ring (major or sub-ring) path to a port and defines the Ring APS tag. Rings typically have two paths a and b.

The **no** form of this command removes the path a or b.

Default no path

Parameters *port-id* — Specifies the port ID.

Values *slot*/*mda*/*port*

lag-id — Specifies the LAG ID.

Values **lag** — Keyword.
id — Specifies the LAG ID number.

raps-tag — Specifies the member's encapsulation.

qtag1 — Specifies the top or outer VLAN ID.

Values 1 to 4094

qtag2 — Specifies the bottom or inner VLAN ID.

Values 1 to 4094

eth-cfm

Syntax **eth-cfm**

Context config>eth-ring>path

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

mep

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

Context config>eth-ring>path>eth-cfm

Description This command provisions an 802.1ag maintenance endpoint (MEP).

The **no** form of the command deletes the MEP.

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

Values 1 to 81921

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

Values 1 to 4294967295

ma-index — Specifies the MA index value.

Values 1 to 4294967295

alarm-notification

Syntax **alarm-notification**

Context config>eth-ring>path>eth-cfm>mep

Description This command enables the context to configure the MEP alarm notification parameters.

fng-alarm-time

Syntax **fng-alarm-time** *time*

Context config>eth-ring>path>eth-cfm>mep>alarm-notification
config>eth-tunnel>path>eth-cfm>mep>alarm-notification

Description This command configures the Fault Notification Generation (FNG) alarm time.

Default 0

Parameters *time* — Specifies the FNG alarm time in centi-seconds.

Values 0, 250, 500, 1000

fng-reset-time

Syntax **fng-reset-time** *time*

Context config>eth-ring>path>eth-cfm>mep>alarm-notification
config>eth-tunnel>path>eth-cfm>mep>alarm-notification

Description This command configures the Fault Notification Generation (FNG) reset time.

Parameters *time* — Specifies the FNG reset time in centi-seconds.

Values 0,250,500,1000

ccm-enable

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

Context config>eth-ring>path>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>eth-ring>path>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 to 7

ccm-padding-size

Syntax **ccm-padding-size** *ccm-padding*

no ccm-padding-size

Context config>eth-ring>path>eth-cfm>mep

Description This command inserts additional padding in the CCM packets.

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

Parameters **ccm-padding** — Specifies the additional padding in the CCM packets.

Values 3 to 1500 octets

control-mep

Syntax [**no**] **control-mep**

Context config>eth-ring>path>eth-cfm>mep

Description This command enables the Ethernet ring control on the MEP. The use of control-mep command is mandatory for a ring. MEP detection of failure using CCM may be enabled or disabled independently of the control mep.

The **no** form of this command disables Ethernet ring control.

Default no control-mep

eth-test-enable

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

Context config>eth-ring>path>eth-cfm>mep

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

oam eth-cfm eth-test *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.

bit-error-threshold

Syntax **bit-error-threshold** *bit-errors*

Context config>eth-ring>path>eth-cfm>mep>eth-test-enable

Description This command specifies the lowest priority defect that is allowed to generate a fault alarm.

Default bit-error-threshold 1

Parameters *bit-errors* — Specifies the lowest priority defect.

Values 0 to 11840

test-pattern

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

Context config>eth-ring>path>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.

Default test-pattern all-zeros

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

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

crc-enable — Generates a CRC checksum.

grace

Syntax **grace**

Context config>eth-ring>path>eth-cfm>mep
config>eth-tunnel>path>eth-cfm>mep

Description This command enables the context to configure Nokia ETH-CFM Grace and ITU-T Y.1731 ETH-ED expected defect functional parameters.

eth-ed

Syntax **eth-ed**

Context config>eth-ring>path>eth-cfm>mep>grace
config>eth-tunnel>path>eth-cfm>mep>grace

Description This command enables the context to configure ITU-T Y.1731 ETH-ED expected defect functional parameters.

max-rx-defect-window

Syntax **max-rx-defect-window** *seconds*
no max-rx-defect-window

Context config>eth-ring>path>eth-cfm>mep>grace>eth-ed
config>eth-tunnel>path>eth-cfm>mep>grace>eth-ed

Description This command limits the duration of the received ETH-ED expected defect window to the lower value of either the received value from the peer or this parameter.

The **no** form of the command removes the limitation, and any valid defect window value received from a peer MEP in the ETH-ED PDU will be used.

Default no max-rx-defect-window

Parameters *seconds* — Specifies the duration, in seconds, of the maximum expected defect window.

Values 1 to 86400

priority

Syntax **priority** *priority*
no priority

Context config>eth-ring>path>eth-cfm>mep>grace>eth-ed
config>eth-tunnel>path>eth-cfm>mep>grace>eth-ed

Description This command sets the priority bits and determines the forwarding class based on the mapping of priority to FC.

The **no** form of the command disables the local priority configuration and sets the priority to the **ccm-ltm-priority** associated with this MEP.

Default no priority

Parameters *priority* — Specifies the priority bit.

Values 0 to 7

rx-eth-ed

Syntax [**no**] **rx-eth-ed**

Context config>eth-ring>path>eth-cfm>mep>grace>eth-ed
config>eth-tunnel>path>eth-cfm>mep>grace>eth-ed

Description This command enables the reception and processing of the ITU-T Y.1731 ETH-ED PDU on the MEP.

The **no** form of the command disables the reception of the ITU-T Y.1731 ETH-ED PDU on the MEP.

Default no rx-eth-ed

tx-eth-ed

Syntax [**no**] **tx-eth-ed**

Context config>eth-ring>path>eth-cfm>mep>grace>eth-ed
config>eth-tunnel>path>eth-cfm>mep>grace>eth-ed

Description This command enables the transmission of the ITU-T Y.1731 ETH-ED PDU from the MEP when a system soft reset notification is received for one or more cards.

The **config**>**eth-cfm**>**system**>**grace-tx-enable** command must be configured to instruct the system that the node is capable of transmitting expected defect windows to the peers. Only one form of ETH-CFM grace (Nokia ETH-CFM Grace or ITU-T Y.1731 ETH-ED) may be transmitted.

The **no** form of the command disables the transmission of the ITU-T Y.1731 ETH-ED PDU from the MEP.

Default no tx-eth-ed

eth-vsm-grace

Syntax **eth-vsm-grace**

Context config>eth-ring>path>eth-cfm>mep>grace
config>eth-tunnel>path>eth-cfm>mep>grace

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

rx-eth-vsm-grace

Syntax [**no**] **rx-eth-vsm-grace**

Context config>eth-ring>path>eth-cfm>mep>grace>eth-vsm-grace
config>eth-tunnel>path>eth-cfm>mep>grace>eth-vsm-grace

Description This command enables the reception and processing of the Nokia ETH-CFM Grace PDU on the MEP.

The Nokia Grace function is a vendor-specific PDU that informs MEP peers that the local node may be entering a period of expected defect.

The **no** form of the command disables the reception of the Nokia ETH-CFM Grace PDU on the MEP.

Default rx-eth-vsm-grace

tx-eth-vsm-grace

Syntax [**no**] **tx-eth-vsm-grace**

Context config>eth-ring>path>eth-cfm>mep>grace>eth-vsm-grace
config>eth-tunnel>path>eth-cfm>mep>grace>eth-vsm-grace

Description This command enables the transmission of the Nokia ETH-CFM Grace PDU from the MEP when a system soft reset notification is received for one or more cards.

The Nokia Grace function is a vendor-specific PDU that informs MEP peers that the local node may be entering a period of expected defect.

The **config**>**eth-cfm**>**system**>**grace-tx-enable** command must be configured to instruct the system that the node is capable of transmitting expected defect windows to the peers. Only one form of ETH-CFM grace (Nokia ETH-CFM Grace or ITU-T Y.1731 ETH-ED) may be transmitted.

The **no** form of the command disables the transmission of the Nokia ETH-CFM Grace PDU from the MEP.

Default tx-eth-vsm-grace

low-priority-defect

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

Context config>eth-ring>path>eth-cfm>mep
config>eth-tunnel>path>eth-cfm>mep

Description This command specifies the lowest priority defect that is allowed to generate a fault alarm.

Default low-priority-defect remErrXcon

Parameters **low-priority-defect** — Specifies the lowest priority defect using the following:

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>eth-ring>path>eth-cfm>mep
config>eth-tunnel>path>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 SDP).

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.

one-way-delay-threshold

Syntax **one-way-delay-threshold** *seconds*

Context config>eth-tunnel>path>eth-cfm>mep

Description This command enables one way delay threshold time limit.

Default 3 seconds

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

Values 0 to 600

one-way-delay-threshold

Syntax **one-way-delay-threshold** *seconds*

Context config>eth-ring>path>eth-cfm>mep

Description This command configures a one way delay threshold time limit.

Default one-way-delay-threshold 3

Parameters *seconds* — Specifies the value, in seconds, for the threshold.

Values 0 to 600

rpl-end

Syntax [**no**] **rpl-end**

Context config>eth-ring>path

Description This command configures the G.8032 path as a ring protection link end. The ring should be declared as either a RPL owner or RPL neighbor for this command to be allowed. Only path a or path b can be declared an RPL-end.

The **no** form of this command sets the rpl-end to default no rpl-end.

Default no rpl-end

revert-time

Syntax **revert-time** *time*
no revert-time

Context config>eth-ring

Description This command configures the revert time for an Eth-Ring. It ranges from 60 seconds to 720 second by 1 second intervals.

The **no** form of this command means non-revertive mode and revert time is essentially 0, and the revert timers are not set.

Default revert-time 300

Parameters *time* — Specifies the guard-time, in seconds.

Values 60 to 720

rpl-node

Syntax **rpl-node** [**owner** | **nbr**]
no rpl-node

Context config>eth-ring

Description This command configures the G.8032 ring protection link type as owner or neighbor. The **no** form of the command means this node is not connected to an RPL link. When RPL owner or neighbor is specified either the a or b path must be configured with the RPL end command. An owner is responsible for operation of the rpl link. Configuring the RPL as neighbor is optional (can be left as no rpl-node) but if the command is used the nbr is mandatory.

On a sub-ring without virtual channel it is mandatory to configure sub-ring non-virtual-link on all nodes on the sub-ring to propagation the R-APS messages around the sub-ring.

The **no** form of this command removes the RPL link.

Default no rpl-node

sub-ring

Syntax [**no**] **sub-ring** {**virtual-link** | **non-virtual-link**}

Context config>eth-ring

Description This command additionally specifies this ring-id to be sub-ring as defined in G.80312. By declaring this ring as a sub-ring object, this ring will only have one valid path and the sub-ring will be connected to a major ring or a VPLS instance. The virtual-link parameter declares that a sub-ring is connected to another ring and that control messages can be sent over the attached ring to the other side of the sub-ring. The non-virtual channel parameter declares that a sub-ring may be connected to a another ring or to a VPLS instance but that no control messages from the sub-ring use the attached ring or VPLS instance. The non-virtual channel behavior is standard G.8032 capability.

The **no** form of this command deletes the sub-ring and its virtual channel associations.

Default no sub-ring

Parameters **virtual-link** — Specifies that the interconnection is to a ring and a virtual link will be used.

non-virtual-link — Specifies that the interconnection is to a ring or a VPLS instance and a virtual link will not be used.

interconnect

Syntax **interconnect** {**ring-id** *ring-index* | **vpls**}
no interconnect

Context config>eth-ring>sub-ring

Description This command links the G.8032 sub-ring to a ring instance or to a VPLS instance. The ring instance must be a complete ring with two paths but may itself be a sub-ring or a major ring (declared by its configuration on another node). When the interconnection is to another node, the sub-ring may have a virtual link or a non-virtual-link. When the sub-ring has been configured with a non-virtual link, the sub ring may be alternatively be connected to a VPLS service. This command is only valid on the interconnection node where a single sub-ring port connects to a major ring or terminates on a VPLS service.

The **no** form of this command removes the interconnect node.

Default no interconnect

Parameters *ring-id* — Specifies the identifier for the ring instance of the connection ring for this sub-ring on this node.

Values 0 to 128

vpls — Specifies that the sub-ring is connected to the VPLS instance that contains the sub-ring SAP.

propagate-topology-change

Syntax [**no**] **propagate-topology-change**

Context config>eth-ring>sub-ring>interconnect

Description This command configures the G.8032 sub-ring to propagate topology changes. From the sub-ring to the major ring as specified in the G.8032 interconnection flush logic. This command is only valid on the sub-ring and on the interconnection node. Since this command is only valid on a Sub-ring, a virtual link or non-virtual link must be specified for this command to be configured. The command is blocked on major rings (when both path a and b are specified on a ring).

The **no** form of this command sets propagate to the default.

Default no propagate-topology-change

2.19.2.9 ETH CFM Configuration Commands

eth-cfm

Syntax **eth-cfm**

Context config

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

default-domain

Syntax **default-domain**

Context config>eth-cfm

Description This command enables the context to configure MIP creation parameters per index (**bridge-identifier** *bridge-id* **vlan** *vlan-id*) if the MIP creation statement exists as part of the service connection. The mip creation statement must be present on the connection before any configuration can occur for a MIP under this context. The determining factor for MIP creation is based on the authoritative properties of the **eth-cfm domain association** configuration. The individual indexes in this table are used for MIP creation only when the association context is not authoritative; this includes the lack of association for a matching index.

bridge-identifier

Syntax **bridge-identifier** *bridge-id* **vlan** *vlan-id*

Context config>eth-cfm>default-domain

Description This command configures the cross-reference required to link the CFM function with the service context. The link is created when the **bridge-id**, **service-id**, and **vlan-id** (for a primary VLAN) match.

Under the **default-domain** context, this command allows the entry of MIP-specific parameters for the index (**bridge-identifier** and **vlan**) in the default-domain table.

This command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**).

A **no** form of this command is only available under the **association** context. Negating the line will remove the **bridge-identifier** and the link between the ETH-CFM configuration and the matching **service-id**.

Parameters *bridge-id* — Specifies the ID for a link to a specific service. Note that there is no verification that a service has been created with a matching service ID.

Values 1 to 2147483647

vlan-id — Specifies the VLAN ID for the **default-domain** index. The complete index allows the user to reference specific MIP entries in the **default-domain** table. The *vlan-id* value must match the configured **primary-vlan-enable** *vlan-id* corresponding to the **bridge-identifier**. If the MIP does not have **primary-vlan-enable** configured, the *vlan-id* must be configured as “none”. When the *vlan-id* is configured as none, the MIP relies on the service delineation for extraction and installs no additional VLAN in that portion of the index.

Values 1 to 4094 | none

id-permission

Syntax **id-permission** {**chassis** | **defer**}
no id-permission

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

Description This command enables the inclusion of the Sender ID TLV information specified under the **config**>**eth**>**system**>**sender-id** command for installed MEPs and MIPs. The inclusion of the Sender ID TLV is based on the configured value. The Sender ID TLV is supported for ETH-CC, ETH-LB, and ETH-LB PDUs.

Note: LBR functions reflect back all TLVs received in the LBM, unchanged, including the Sender ID TLV. Transmission of the Management Domain and Management Address fields are not supported in this TLV.

The **no** form of this command disables the inclusion of the Sender ID TLV.

Default config>eth-cfm>default-domain>bridge-identifier>id-permission defer

config>eth-cfm>domain>association>bridge>no id-permission

Parameters **chassis** — Keyword to include the Sender ID TLV with a value equal to the *sender-id* configured under the eth-cfm>system context.

defer — Keyword to specify that **id-permission** will inherit the value from the global read-only system values.

mhf-creation

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

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

Description This command defines the MIP method of creation. MIP creation mode and other factors are part of the MIP creation authority (**association** or **default-domain**) logic. The MIP creation algorithm may result in multiple potential MIPs. Only the lowest-level valid MIP is installed. The **static** creation mode is the exception to the single MIP installation rule.

Under the association context, the **level** *level* parameter is not supported as part of this command. The level is derived from the level configuration of the domain.

The **no** form of this command is only available under the **association** context, and reverts the current mode of creation to the default **none**. In order to transition to and from the **static** mode of operation, the active **mhf-creation** mode must be **none**.

Default config>eth-cfm>default-domain>bridge-identifier>mhf-creation defer config>eth-cfm>domain>association>bridge-identifier>mhf-creation none

Parameters **none** — Specifies that no MHFs (MIPs) can be created for this SAP or spoke SDP.

default — Specifies MHFs (MIPs) can be created for this SAP or spoke SDP without the requirement for a MEP at some lower MA level. If a lower-level MEP exists, the creation method will behave as **explicit**.

explicit — Specifies that MHFs (MIPs) can be created for this SAP or spoke SDP only if a MEP is created at some lower MD Level. There must be at least one lower MD Level MEP provisioned on the same SAP or spoke SDP.

defer — Defers the MIP creation process to the system-wide read-only values. This parameter is only configurable under the **default-domain** context.

level — Specifies the requested level of the MIP. This is used by the MIP creation algorithm to determine its validity in comparison to other ETH-CFM MIPs in the same service. If *level* is configured as “defer”, the level value will be inherited from the global read-only system values, and “-1” will be stored as a MIB value in the table.

Values 0 to 7, **defer**

Default defer

mip-ltr-priority

Syntax **mip-ltr-priority** *priority*

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

Description This command allows the operator to set the priority of the Linktrace Response Message (ETH-LTR) from a MIP for this association.

Default config>eth-cfm>default-domain>bridge-identifier-vlan>mip-ltr-priority defer

config>eth-cfm>domain>association>bridge-identifier>mip-ltr-priority 7

Parameters *priority* — Specifies the priority of the Linktrace Response Message (ETH-LTR) from a MIP. The “defer” value is only supported under the default-domain context and causes **mip-ltr-priority** to inherit values from the global read-only-system values.

Values 0 to 7, **defer**

domain

Syntax **domain** *md-index* [**format** {*format*}] [**name** *md-name*] **level** *level* [**admin-name** *admin-name*]
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 to 4294967295

format *format* — Specifies a value that represents the type (format).

Values dns, mac, none, string

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 to 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 to 7

admin-name *admin-name* — Specifies a creation time required parameter that allows the operator to assign a name value to the domain container. This is used for information and migration purposes. This value cannot be modified without destroying the domain. If no **admin-name** exists, the configured *md-index* value will be converted into a character string to become the **admin-name** reference. When upgrading from a release that does not include the **admin-name** configuration option, the *md-index* will be converted into a character string. Once a value is assigned to this *admin-name* value it cannot be modified.

Values 1 to 64 characters

association

Syntax **association** *ma-index* [**format** {*format*}] **name** *ma-name* [**admin-name** *admin-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 to 4294967295

format — Specifies a value that represents the type (format).

Values icc-based, integer, string, vid, vpn-id

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 to 65535 (integer value 0 means the MA is not attached to a VID.)
string:	raw ascii
vid:	0 to 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

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

Values 1 to 45 characters

admin-name *admin-name* — Specifies a creation time required parameter that allows the operator to assign a name value to the domain container. This is used for information and migration purposes. This value cannot be modified without destroying the domain. If no **admin-name** exists, the configured *md-index* value will be converted into a character string to become the **admin-name** reference. When upgrading from a release that does not include the **admin-name** configuration option, the *md-index* will be converted into a character string. Once a value is assigned to this *admin-name* value it cannot be modified.

Values 1 to 64 characters

auto-mep-discovery

Syntax [**no**] **auto-mep-discovery**

Context config>eth-cfm>domain>association

Description Enable/disable the ability to auto-discover remote MEPs from a peer MEP sending ETH-CC.

Default no auto-mep-discovery

bridge-identifier

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

Context config>eth-cfm>domain>association

Description This command configures the cross-reference required to link the CFM function with the service context. The link is created when the **bridge-id**, or **bridge-name** matches the **service-id**, or **service-name**, respectively.

The **no** form of this command removes the **bridge-identifier** and the link between the ETH-CFM configuration and the matching service.

There is no verification that any service has been created with a matching value. An existing **bridge-identifier** configuration can be overwritten with the alternate type, as long as the new reference does not change the existing service linkage.

Any SNMP reference that performs this remaining function must copy all sub parameters to correct table, as is done with the interactive CLI command. The bridge-id based tables are dot1agCfmMaCompTable and tmnxDot1agCfmMaCompTable. The bridge-name table is tmnxDot1agCfmMaBrNameTable

Parameters *bridge-id* — Specifies the ID for a link to a specific service. Note that there is no verification that a service has been created with a matching service ID.

This *bridge-id* variant of the command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **bridge-identifier bridge-name** *bridge-name* variant can be used in all configuration modes.

Values 1 to 2147483647

bridge-name *bridge-name* — Specifies a link to a service by service-name, up to 64 characters.

id-permission

Syntax **id-permission** {**chassis**}
no id-permission

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

Description This command enables the inclusion of the Sender ID TLV information specified under the **config**>**eth**>**system**>**sender-id** command for installed MEPs and MIPs. The inclusion of the Sender ID TLV is based on the configured value. The Sender ID TLV is supported for ETH-CC, ETH-LB, and ETH-LB PDUs.

Note: LBR functions reflect back all TLVs received in the LBM, unchanged, including the Sender ID TLV. Transmission of the Management Domain and Management Address fields are not supported in this TLV.

The **no** form of this command disables the inclusion of the Sender ID TLV.

Default config>eth-cfm>default-domain>bridge-identifier>id-permission defer

config>eth-cfm>domain>association>bridge>no id-permission

Parameters **chassis** — Keyword to include the Sender ID TLV with a value equal to the *sender-id* configured under the eth-cfm>system context.

defer — Keyword to specify that **id-permission** will inherit the value from the global read-only system values.

mhf-creation

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

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

Description This command defines the MIP method of creation. MIP creation mode and other factors are part of the MIP creation authority (**association** or **default-domain**) logic. The MIP creation algorithm may result in multiple potential MIPs. Only the lowest-level valid MIP is installed. The **static** creation mode is the exception to the single MIP installation rule.

Under the association context, the **level** *level* parameter is not supported as part of this command. The level is derived from the level configuration of the domain.

The **no** form of this command is only available under the **association** context, and reverts the current mode of creation to the default **none**. In order to transition to and from the **static** mode of operation, the active **mhf-creation** mode must be **none**.

Default config>eth-cfm>domain>association>bridge-identifier>mhf-creation none

Parameters **none** — Specifies that no MHFs (MIPs) can be created for this SAP or spoke SDP.

default — Specifies MHFs (MIPs) can be created for this SAP or spoke SDP without the requirement for a MEP at some lower MA level. If a lower-level MEP exists, the creation method will behave as **explicit**.

explicit — Specifies that MHFs (MIPs) can be created for this SAP or spoke SDP only if a MEP is created at some lower MD Level. There must be at least one lower MD Level MEP provisioned on the same SAP or spoke SDP.

static — Specifies the exact level of the MHF (MIP) that will be created for this SAP. Multiple MHFs (MIPs) are allowed as long as the MD Level hierarchy is properly configured for the particular Primary VLAN. Ingress MHFs (MIPs) with primary VLAN are not supported on SDP Bindings.

vlan

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

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

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

Default no vlan

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

Values 0 to 4094

ccm-hold-time

Syntax **ccm-hold-time down** *timer*
no ccm-hold-time

Context config>eth-cfm>domain>association

Description This command allows a sub second CCM enabled MEP to delay a transition to a failed state if a configured remote CCM peer has timed out. The MEP will remain in the UP state for 3.5 times CCM interval + down-delay.

The **no** form of this command removes the additional delay

Default no ccm-hold-time

Parameters **down** *timer* — Specifies the amount of time to delay in 100ths of a second.

Values 0-1000

ccm-interval

Syntax **ccm-interval** *interval*
no ccm-interval

Context config>eth-cfm>domain>association

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

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

Default no ccm-interval

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

Values 10 milliseconds, 100 milliseconds, 1 second, 10 seconds, 60 seconds, 600 seconds, 100 milliseconds

Default 10 (seconds)

facility-id-permission

Syntax **facility-id-permission** {**chassis**}
no facility-id-permission

Context config>eth-cfm>domain>association

Description This command configures the id-permission for facility MEPs for the association.

Default no facility-id-permission

remote-mepid

Syntax **remote-mepid** *mep-id* **remote-mac** {*unicast-da* | **default**}
no remote-mepid *mep-id*

Context config>eth-cfm>domain>association

Description This command identifies remote maintenance association endpoint (MEP) the systems is expecting to receive packets form. Optionally, the operator may configure a unicast MAC address associated with the remote-mep. This unicast value will replace the default layer two class 1 multicast address that is typically associated with ETH-CC packets.

Note: This command is not supported with sub second CCM intervals. The **unicast-da** parameter may only be configured when a single remote MEP exists in the association.

Parameters *mep-id* — Specifies the remote MEP identifier.

Values 1 to 8191

remote-mac {*unicast-da* | **default**} — Specifies the remote MAC type.

Values unicast-da —The unicast layer two destination address in the form xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx.

default — Removes the unicast address and reverts back to class 1 multicast.

redundancy

Syntax **redundancy**

Context config>eth-cfm

Description This command enables the context under which the ETH-CFM redundancy parameters are to be configured.

Default none

mc-lag

Syntax **mc-lag**

Context config>eth-cfm>redundancy

Description This command enables the context under which the MC-LAG specific ETH-CFM redundancy parameters are to be configured

Default none

propagate-hold-time

Syntax **propagate-hold-time** *second*
no propagate-hold-time

Context config>eth-cfm>redundancy>mc-lag

Description This command configures the delay, in seconds, that fault propagation is delayed because of port or MC-LAG state changes. This provides the amount of time for system stabilization during a port state changes that may be protected by MC-LAG. This command requires the standby-mep-shutdown command in order to take effect.

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

Default propagate-hold-time 1

Parameters *seconds* — Specifies the amount of time in seconds. Zero means no delay.

Values 0 to 60

standby-mep-shutdown

Syntax [**no**] **standby-mep-shutdown**

Context config>eth-cfm>redundancy>mc-lag

Description This system wide command enables MEPs to track the state of MC-LAG. This allows MEPs on the standby MC-LAG to act administratively down.

The **no** form of command disables the MEP tracking.

Default no standby-mep-shutdown

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. Should the time between packets exceed this values within a test the responder will mark the previous test as complete. It will treat 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.

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

Default inactivity-timer 100

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

Values 10 100

system

Syntax **system**

Context config>eth-cfm

Description This command enables the context to configure Connectivity Fault Management General System parameters.

grace-tx-enable

Syntax [**no**] **grace-tx-enable**

Context config>eth-cfm>system

Description This command enables and disables the transmission of ETH-VSM messages to delay CCM timeout and AIS churn during ISSU and soft reset functions.

Default grace-tx-enable

sender-id

Syntax **sender-id local** *local-name*
sender-id system
no sender-id

Context config>eth-cfm>system

Description This command configures the ETH-CFM sender-id used in CFM PDUs.

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

Default sender-id system

Parameters **system** — Specifies to use the system name.

local-name — Specifies to use the local name up to 45 alphanumeric characters in length.

2.19.2.10 Port and LAG ETH CFM Commands

mep

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

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

Description This command provisions the maintenance endpoint (MEP).

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

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

Values 1 to 81921

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

Values 1 to 4294967295

ma-index — Specifies the MA index value.

Values 1 to 4294967295

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

Values 1 to 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 to 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** — Specifies 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* — Specifies the priority value of the AIS messages originated by the node.

Values 0 to 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-padding-size

Syntax **ccm-padding-size** *ccm-padding*

no ccm-padding-size

Context config>eth-tunnel>path>eth-cfm>mep

Description This command inserts additional padding in the CCM packets.

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

Parameters **ccm-padding** — Specifies the additional padding in the CCM packets.

Values 3 to 1500 octets

control-mep

Syntax [**no**] **control-mep**

Context config>eth-ring>path>eth-cfm>mep

Description This command enables the Ethernet ring control on the MEP. The use of control-mep command is mandatory for an Ethernet ring. MEP detection of failure using CCM may be enabled or disabled independently of the control mep.

The **no** form of this command disables Ethernet ring control.

Default no control-mep

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.

Default no mac-address

Parameters **mac-address** *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.

facility-fault

Syntax [**no**] **facility-fault**

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

Description Allows the facility MEP to move from alarming only to network actionable function. This means a facility MEP will not merely report the defect conditions but will be able to action based on the transition of the MEP state. Without this command the facility MEP will only monitor and report and conditions of the MEP do not affect related services.

Default no facility-fault

2.19.2.11 ETH-Tunnel Commands

eth-tunnel

Syntax [**no**] **eth-tunnel** *tunnel-index*

Context config

Description This command configures a unique Ethernet Tunnel Identifier for an Ethernet Tunnel Group.

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

Default none

Parameters *tunnel-index* — Specifies a tunnel index identifier.

Values 1 to 1024

ccm-hold-time

Syntax **ccm-hold-time** [**down** *down-timeout*] [**up** *up-timeout*]
no ccm-hold-time

Context config>eth-tunnel

Description This command allows a sub second CCM enabled MEP to delay a transition to a failed state if a configured remote CCM peer has timed out. The MEP will remain in the UP state for 3.5 times CCM interval + down-delay.

The **no** form of this command removes the additional delay

Parameters **down** *down-timeout* — Specifies the time, in centiseconds, used for the hold-timer for associated Continuity Check (CC) Session down event dampening. This guards against reporting excessive member operational state transitions.

This is implemented by not advertising subsequent transitions of the CC state to the Ethernet Tunnel Group until the configured timer has expired.

Values 0 to 1000

Default 0

up *up-timeout* — Specifies the time, in deciseconds, used for the hold-timer for associated Continuity Check (CC) Session up event dampening. This guards against reporting excessive member operational state transitions.

This is implemented by not advertising subsequent transitions of the CC state to the Ethernet Tunnel Group until the configured timer has expired.

Values 0 to 5000

Default 20

ethernet

Syntax **ethernet**

Context config>eth-tunnel

Description This command enables the context to configure Ethernet parameters for the Ethernet tunnel.

encap-type

Syntax **encap-type** {**dot1q|qinq**}
no encap-type

Context config>eth-tunnel>ethernet

Description This command configures the encapsulation method used to distinguish customer traffic on a LAG. The encapsulation type is configurable on a LAG port. The LAG port and the port member encapsulation types must match when adding a port member.

If the encapsulation type of the LAG port is changed, the encapsulation type on all the port members will also change. The encapsulation type can be changed on the LAG port only if there is no interface associated with it. If the MTU is set to a non-default value, it will be reset to the default value when the encap type is changed.

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

Default encap-type dot1q

Parameters **dot1q** — Specifies that frames carry 802.1Q tags where each tag signifies a different service.

qinq — Specifies the qinq encapsulation method.

mac

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

Context config>eth-tunnel>ethernet

Description This command assigns a specific MAC address to an Ethernet port, Link Aggregation Group (LAG), Ethernet tunnel, or BCP-enabled port or sub-port.

Only one MAC address can be assigned to a port. When multiple **mac** commands are entered, the last command overwrites the previous command. When the command is issued while the port is operational, IP will issue an ARP, if appropriate, and BPDUs are sent with the new MAC address.

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

Default A default MAC address is assigned by the system from the chassis MAC address pool.

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. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses6-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.

lag-emulation

Syntax **lag-emulation**

Context config>eth-tunnel

Description This command enables the context to configure eth-tunnel loadsharing parameters.

access

Syntax **access**

Context config>eth-tunnel>lag-emulation

Description This command enables the context to configure eth-tunnel loadsharing access parameters.

adapt-qos

Syntax **adapt-qos** {**distribute** | **link** | **port-fair**}
no adapt-qos

Context config>eth-tunnel>lag-emulation>access

Description This command specifies how the emulated LAG queue and virtual scheduler buffering and rate parameters are adapted over multiple active MDAs.

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

Parameters **distribute** — Creates an additional internal virtual scheduler per line card as parent of the configured SAP queues and virtual schedulers per member path on that line card. This internal virtual scheduler limits the total amount of egress bandwidth for all member paths on the line card to that line card’s share of the bandwidth specified in the egress qos policy. This mode is not supported together with an egress port scheduler or the use of egress queue groups.

link — Specifies that the emulated LAG will create the SAP queues and virtual schedulers with the bandwidth specified in the egress QoS policy on each member path.

port-fair — Specifies that the emulated LAG will create the SAP queues and virtual schedulers on each member path based on the bandwidth specified in the egress QoS policy divided by the number of active paths.

per-fp-ing-queuing

Syntax [**no**] **per-fp-ing-queuing**

Context config>eth-tunnel>lag-emulation>access

Description This command specifies whether a more efficient method of queue allocation for the LAG should be utilized.

The **no** form of the command disables the method of queue allocation.

path-threshold

Syntax **path-threshold** *num-paths*
no path-threshold

Context config>eth-tunnel>lag-emulation

Description This command configures whether a more efficient method of queue allocation for Ethernet Tunnel Group SAPs should be utilized.

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

Default no path-threshold

Parameters **num-paths** — Specifies the behavior for the eth-tunnel if the number of operational members is equal to or below a threshold level.

Values 0 to 15

path

Syntax [**no**] **path** *path-index*

Context config>eth-tunnel

Description This command configures one of the two paths supported under the Ethernet tunnel.

The **no** form of this command removes the path from under the Ethernet tunnel. If this is the last path, the associated SAP need to be un-configured before the path can be deleted.

Default no path

Parameters **path-index** — Specifies the identifier for the path.

Values 1 to 16

control-tag

Syntax **control-tag** *qtag*[.*qtag*]
no control-tag

Context config>eth-tunnel>path

Description This command specifies the VLAN-ID to be used for Ethernet CFM and G.8031 control plane exchanges. If the operator wants to replace an existing control-tag, the parent path needs to be in shutdown state, then deleted and recreated before a new control-tag can be specified.

The **no** form of this command is used just to indicate that a control-tag is not configured. The procedure described above, based on ‘no path’ command must be used to un-configure/ change the control-tag assigned to the path.

Default no control-tag

Parameters **vlan-id** — Specifies the value of the VLAN ID to be used for the control tag.

Values 0 to 4094

eth-cfm

Syntax **eth-cfm**

Context config>eth-tunnel>path

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

mep

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

Context config>eth-tunnel>path>eth-cfm

Description This command provisions an 802.1ag 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 to 81921

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

Values 1 to 4294967295

ma-index — Specifies the MA index value.

Values 1 to 4294967295

alarm-notification

Syntax **alarm-notification**

Context config>eth-tunnel>path>eth-cfm>mep

Description This command enables the context to configure the MEP alarm notification parameters.

ccm-enable

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

Context config>eth-tunnel>path>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>eth-tunnel>path>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 to 7

ccm-padding-size

Syntax **ccm-padding-size** *ccm-padding*
no ccm-padding-size

Context config>eth-tunnel>path>eth-cfm>mep

Description This command inserts additional padding in the CCM packets.

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

Parameters **ccm-padding** — Specifies the additional padding in the CCM packets.

Values 3 to 1500 octets

control-mep

Syntax [**no**] **control-mep**

Context config>eth-tunnel>path>eth-cfm>mep

Description This command enables the Ethernet tunnel control on the MEP. The use of control-mep command is mandatory for an Ethernet tunnel. MEP detection of failure using CCM may be enabled or disabled independently of the control mep.

The **no** form of this command disables Ethernet ring control.

Default no control-mep

eth-test-enable

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

Context config>eth-tunnel>path>eth-cfm>mep

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

oam eth-cfm eth-test *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.

bit-error-threshold

Syntax **bit-error-threshold** *bit-errors*

Context config>eth-tunnel>path>eth-cfm>mep>eth-test-enable

Description This command specifies the lowest priority defect that is allowed to generate a fault alarm.

Default bit-error-threshold 1

Parameters *bit-errors* — Specifies the lowest priority defect.

Values 0 to 11840

test-pattern

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

Context config>eth-tunnel>path>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.

Default test-pattern all-zeros

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

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

crc-enable — Generates a CRC checksum.

member

Syntax **member** *port-id*
no member

Context config>eth-tunnel>path

Description This command configures the path member.

The **no** form of the command removes the port-id from the configuration.

Default none

Parameters **port-id** — Specifies the path member.

Values slot/mda/port

port-id	*slot*/*mda*/*port*[.*channel*]		
	pxc-id	psc-id.sub-port	
		pxc psc-id.sub-port	
		pxc: keyword	
		id: 1 to 64	
		sub-port: a, b	
	aps-id	aps-*group-id*[.*channel*]	
		aps keyword	
		group-id	1 to 64
		group-id	1 to 16
	bundle-*type*-*slot/mda*.*bundle-num*		
		bundle	keyword
		type	ima, ppp
		bundle-num	1 to 256
	bpgrp-id:	**bpgrp**-*type*-*bpgrp-num*	
		bpgrp	keyword
		type	ima
		bpgrp-num	1 to 1280
	ccag-id - ccag-<id>.<path-id>[cc-type]		
		ccag	keyword
		id	1 to 8
		path-id	a, b
		cc-type[.sap-net | .net-sap]	
	lag-id	lag-*id*	
		lag	keyword
		id	1 to 800

precedence

Syntax **precedence** {**primary**|**secondary**}

Context config>eth-tunnel>path

Description This command specifies the precedence to be used for the path. Only two precedence options are supported: **primary** and **secondary**.

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

Default precedence secondary

Parameters **primary | secondary** — Specifies the path precedence as either primary or secondary.

protection-type

Syntax **protection-type** {**g8031-1to1 | loadsharing**}

Context config>eth-tunnel

Description This command configures the model used for determining which members are actively receiving and transmitting data.

When the value is set to “g8031-1to1 (1)”, as per the G.8031 specification, only two members are allowed, and only one of them can be active at one point in time.

When the value is set to “loadsharing (2)”, multiple members can be active at one point in time.

Default protection-type g8031-1to1

revert-time

Syntax **revert-time** *time*
no revert-time

Context config>eth-tunnel

Description This command configures the revert time for an Eth tunnel. It ranges from 60 seconds to 720 seconds by 1-second intervals.

The **no** form of this command means non-revertive mode and revert time is 0, so the revert timers are not set.

Default revert-time 300

Parameters *value* — Specifies the guard-time, in seconds.

Values 60 to 720

2.19.2.12 Connection Profile VLAN Commands

connection-profile-vlan

Syntax **connection-profile-vlan** *conn-prof-id* [**create**]
no connection-profile-vlan *conn-prof-id*

Context config

Description This command enables the context to configure the VLAN ranges that will be associated with a service SAP.

Default none

Each connection-profile-vlan must be explicitly configured.

Parameters *conn-prof-id* — Specifies the connection-profile identifier. This value will be configured in the service along with the SAP when the user associates a VLAN bundle to a single SAP. For example, a SAP defined in a dot1q port 1/1/1 that matches all the VLANs defined in the connection-profile-vlan 1 will be created as '**sap 1/1/1:cp-1 create**'.

Values 1 to 8000

vlan-range

Syntax **vlan-range** *from* [**to** *to*]
no vlan-range *from*

Context config>connection-profile-vlan

Description This command allows the user to configure different ranges in the connection-profile-vlan. The ranges have the following characteristics:

- Ranges can contain a single VID or start-and-end values. When the *to-vid* is not specified, the end vid value is the same as the start vid value.
- On the fly addition/removal of ranges is allowed.
- When removing an entry, the **no vlan-range** *vid* **to** *vid* must be configured by the user.
- Multiple ranges are allowed under the same connection-profile-vlan. No VLAN values should overlap within the same connection-profile-vlan.
- The index for connection-profile and connection-profile-vlan must be unique between the two. For example, if **connection-profile 100** is present, then **connection-profile-vlan 100** will be disallowed.

Default none

Each vlan-range must be explicitly configured.

Parameters *from* — Specifies the beginning of the **vlan-range** associated to the **connection-profile-vlan**.

Values 1 to 4094

to — Specifies the end of the **vlan-range** associated to the **connection-profile-vlan**. If not specified, the **vlan-range** is comprised of only the *from* VLAN ID.

Values 1 to 4094

2.19.2.13 Network Group Encryption (NGE) Commands

group-encryption

Syntax **group-encryption**

Context config

Description This command enables the context to configure group encryption parameters.

encryption-keygroup

Syntax **encryption-keygroup** *keygroup-id* [**create**]
no encryption-keygroup *keygroup-id*

Context config>grp-encryp

Description This command is used to create a key group. Once the key group is created, use the command to enter the key group context or delete a key group.

The **no** form of the command removes the key group. Before using the **no** form, the key group association must be deleted from all services that are using this key group.

Parameters *keygroup-id* — The number or name of the key group being referenced.

Values 1 to 15, or *keygroup-name* (up to 64 characters)

create — Creates a key group.

active-outbound-sa

Syntax **active-outbound-sa** *spi*
no active-outbound-sa

Context config>grp-encryp>encryp-keygrp

Description This command specifies the Security Association, referenced by the Security Parameter Index (SPI), to use when performing encryption and authentication on NGE packets egressing the node for all services configured using this key group.

The **no** form of the command returns the parameter to its default value and is the same as removing this key group from all outbound direction key groups in all services configured with this key group (that is, all packets of services using this key group will egress the node in without being encrypted).

Parameters *spi* — Specifies the SPI to use for packets of services using this key group when egressing the node.

Values 1 to 127

esp-auth-algorithm

Syntax **esp-auth-algorithm** {**sha256** | **sha512**}
no esp-auth-algorithm

Context config>grp-encryp>encryp-keygrp

Description This command specifies the hashing algorithm used to perform authentication on the Encapsulating Security Payload (ESP) within NGE packets for services configured using this key group. All SPI entries must be deleted before the **no** form of the command may be entered or the **esp-auth-algorithm** value changed from its current value.

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

Default sha256

Parameters **sha256** — Configures the ESP to use the HMAC-SHA-256 algorithm for authentication.

sha512 — Configures the ESP to use the HMAC-SHA-512 algorithm for authentication.

esp-encryption-algorithm

Syntax **esp-encryption-algorithm** {**aes128** | **aes256**}
no esp-encryption-algorithm

Context config>grp-encryp>encryp-keygrp

Description This command specifies the encryption algorithm used to perform encryption on the Encapsulating Security Payload (ESP) within NGE packets for services configured using this key group. All SPI entries must be deleted before the **no** form of the command may be entered or the **esp-encryption-algorithm** value changed from its current value.

The **no** form of the command resets the parameter to the default value.

Default aes128

Parameters **aes128** — Configures the AES algorithm with a block size of 128 bits—a very strong algorithm choice.

aes256 — Configures the AES algorithm with a block size of 256 bits—the strongest available version of AES.

keygroup-name

Syntax **keygroup-name** *keygroup-name*
no keygroup-name

Context config>grp-encryp>encryp-keygrp

Description This command is used to name the key group. The key group name can be used to reference a key group when configuring services or displaying information.

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

Parameters *keygroup-name* — The name of the key group, up to 64 characters.

security-association

Syntax **security-association spi** *spi* **authentication-key** *authentication-key* **encryption-key** *encryption-key* [**crypto**]
no security-association spi *spi*

Context config>grp-encryp>encryp-keygrp

Description This command is used to create a security association for a specific SPI value in a key group. The command is also used to enter the authentication and encryption key values for the security association, or to delete a security association.

The SPI value used for the security association is a node-wide unique value, meaning that no two security associations in any key group on the node may share the same SPI value.

Keys are entered in clear text. After configuration, they are never displayed in their original, clear text form. Keys are displayed in an encrypted form, which is indicated by the system-appended **crypto** keyword when an **info** or an **admin>save** command is run. For security reasons, keys encrypted on one node are not usable on other nodes (that is, keys are not exchangeable between nodes).

The **no** form of the command removes the security association and related key values from the list of security associations for the key group. If the **no** form of the command is attempted using the same SPI value that is configured for **active-outbound-sa**, then a warning is issued and the command is blocked. If the **no** form of the command is attempted on the last SPI in the key group and the key group is configured on a service, then the command is blocked.

Parameters *spi* — Specifies the SPI ID of the SPI being referenced for the security association.

Values 1 to 127

authentication-key — Specifies the authentication key for the SPI, in hexadecimal format. The number of characters in the hexadecimal string must be 64 or 128, depending on whether the authentication algorithm is set to sha256 or sha512, respectively.

encryption-key — Specifies the encryption key for the SPI, in hexadecimal format. The number of characters in the hexadecimal string must be 32 or 64, depending on whether the encryption algorithm is set to aes128 or aes256, respectively.

crypto — Displays the keys showing on the CLI **info** display in an encrypted form.

group-encryption-label

Syntax **group-encryption-label** *encryption-label*
no group-encryption-label

Context config>grp-encryp

Description This command configures the group encryption label used to identify when an MPLS payload is encrypted. This label must be unique network-wide and must be configured consistently on all nodes participating in a network group encryption domain. The label cannot be changed or deleted when there are any key groups configured on the node.

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

Parameters *encryption-label* — The network-wide, unique reserved MPLS label for group encryption.

Values 32 to 2047

2.19.2.14 NGE Services Commands

encryption-keygroup

Syntax **encryption-keygroup** *keygroup-id* **direction** {**inbound | outbound**}
no encryption-keygroup direction {**inbound | outbound**}

Context config>service>sdp
config>service>vprn

Description This command is used to bind a key group to an SDP or VPRN service for inbound or outbound packet processing. When configured in the outbound direction, packets egressing the node use the **active-outbound-sa** associated with the key group configured. When configured in the inbound direction, received packets must be encrypted using one of the valid security associations configured for the key group. Services using the SDP will be encrypted.

The encryption (enabled or disabled) configured on an SDP used to terminate a Layer 3 spoke SDP of a VPRN always overrides any VPRN-level configuration for encryption.

Encryption is enabled once the outbound direction is configured.

The **no** form of the command removes the key group from the SDP or service in the specified direction (inbound or outbound).

Parameters *keygroup-id* — The number of the key group being configured.

Values 1 to 15 or *keygroup-name* (up to 64 characters)

direction {**inbound | outbound**} — Specifies the direction of the service that the keygroup will be bound to.

2.19.2.15 Model-Driven Automatic ID Commands

md-auto-id

Syntax **md-auto-id**

Context config>service

Description This command automatically assigns numerical ID values for model-driven (MD) management interfaces.

Classic management interfaces use a numerical service ID, customer ID, and PW template ID as the primary key for services, customers, and PW templates. In model-driven interfaces, services, customers, and PW templates use string names as keys. The services, customers, and PW templates can optionally be created without having to explicitly select and specify a numerical ID. In this case, SR OS assigns an ID using the configured ID range.

customer-id-range

Syntax **customer-id-range start** *customer-id* **end** *customer-id*
no customer-id-range

Context config>service>md-auto-id

Description This command specifies the range of IDs used by SR OS to automatically assign an ID to customers that are created in model-driven interfaces without an ID explicitly specified by the user or client.

A customer created with an explicitly-specified ID cannot use an ID in this range. The ID range cannot be changed while customers exist inside the previous or new range.

The **no** form of this command removes the range values.

See the md-auto-id command for further details.

Default no customer-id-range

Parameters **start** *customer-id* — Specifies the lower value of the ID range. The value must be less than or equal to the **end** value.

Values 2 to 2147483647

end *customer-id* — Specifies the upper value of the ID range. The value must be greater than or equal to the **start** value.

Values 2 to 2147483647

pw-template-id-range

Syntax **pw-template-id-range start** *pw-template-id* **end** *pw-template-id*
no pw-template-id-range

Context config>service>md-auto-id

Description This command specifies the range of IDs used by SR OS to automatically assign an ID to PW templates that are created in model-driven interfaces without an ID explicitly specified by the user or client.

A PW template created with an explicitly-specified ID cannot use an ID in this range. The ID range cannot be changed while pw-templates exist inside the previous or new range.

The **no** form of this command removes the range values.

See the md-auto-id command for further details.

Default no pw-template-id-range

Parameters **start** *pw-template-id* — Specifies the lower value of the ID range. The value must be less than or equal to the **end** value.

Values 1 to 2147483647

end *pw-template-id* — Specifies the upper value of the ID range. The value must be greater than or equal to the **start** value.

Values 1 to 2147483647

service-id-range

Syntax **service-id-range start** *service-id* **end** *service-id*
no service-id-range

Context config>service>md-auto-id

Description This command specifies the range of IDs used by SR OS to automatically assign an ID to services that are created in model-driven interfaces without an ID explicitly specified by the user or client.

A service created with an explicitly-specified ID cannot use an ID in this range. The ID range cannot be changed while services exist inside the previous or new range.

The **no** form of this command removes the range values.

See the md-auto-id command for further details.

Default no service-id-range

Parameters **start** *service-id* — Specifies the lower value of the ID range. The value must be less than or equal to the **end** value.

Values 1 to 2147483647

end *service-id* — Specifies the upper value of the ID range. The value must be greater than or equal to the **start** value.

Values 1 to 2147483647

2.20 Show, Clear, and Tools Command Reference

This section provides an overview of the show, clear, debug and tools command reference.

Topics in this section include:

- Command Hierarchies
- Command Descriptions

2.20.1 Command Hierarchies

- Show Commands
- Clear Commands
- Tools Perform Commands
- Tools Dump Commands

Note: For information on egress multicast group commands, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN*.

2.20.1.1 Show Commands

```
show
    — service
          — customer [customer-id] [site customer-site-name]
          — fdb-mac [ieee-address] [expiry]
          — id service-id
          — id service-id base
          — id service-id bgp [bgp-instance]
          — id service-id mac-notification
          — id service-id macsec
          — id service-id mrp
          — id service-id mvrp vlan vlan-id
          — id service-id mvrp vlan detail
          — id service-id provider-tunnel
          — id service-id sap base
          — id service-id sap mrp
          — id service-id sdp
          — id service-id vpls-group [vpls-group-id]
          — id service-id vpls-group vpls-group-id non-template-saps
```

  - **isid-using** [*range-id*]
  - **l2-route-table** [**detail**] [**bgp-ad**] [**multi-homing**] [**bgp-vpls**] [**bgp-vpws**] [**all-routes**]
  - **md-auto-id**
  - **oper-group** [*group-name*]
  - **oper-group** [*group-name*] **detail**
  - **oper-group** [*group-name*] **members** [**sap**] [**sdp**] [**site**]
  - **oper-group** [*group-name*] **monitoring** [**sap**] [**sdp**] [**site**] [**mvrp**]
  - **pw-sap-using**
  - **pw-template** [*policy-id*]
  - **saii-type2-using** *global-id*[:*prefix*[:*ac-id*]]
  - **sap-using** [**msap**] [**dyn-script**] [**description**]
  - **sap-using** [**sap** *sap-id*] [**vlan-transaction** | **anti-spoof**] [**description**]
  - **sap-using** {*ingress* | *egress*} **atm-td-profile** *td-profile-id*
  - **sap-using** {*ingress* | *egress*} **filter** *any-filter-id*
  - **sap-using** {*ingress* | *egress*} **qos-policy** *qos-policy-id* [**msap**]
  - **sap-using** **etree**
  - **sdp** *sdp-id* **pw-port** [*pw-port-id*] [**statistics**]
  - **sdp** [**consistent** | **inconsistent** | **na**] **egressifs**
  - **sdp** *sdp-id* **keep-alive-history**
  - **sdp** **far-end** *ip-address* **keep-alive-history**
  - **sdp** [*sdp-id*] [**detail**]
  - **sdp** **far-end** *ip-address* [**detail**]
  - **sdp-group** [*group-name*]
  - **sdp-group-using** [*sdp-group*]
  - **sdp-using** [*sdp-id*[:*vc-id*] | **far-end** *ip-address*]
  - **service-using** [**epipe**] [**ies**] [**vpls**] [**vprn**] [**mirror**] [**b-vpls**] [**i-vpls**] [**m-vpls**] [**apipe**] [**fpipe**] [**ipipe**] [**sdp** *sdp-id*] [**customer** *customer-id*]
  - **system**
    - **bgp-auto-rd**
    - **bgp-route-distinguisher** [**vprn**] [**vpls**] [**epipe**]
    - **bgp-route-distinguisher** **svc**
    - **bgp-route-distinguisher** **ad-evi-rt-set**
    - **bgp-route-distinguisher** **system**
  - **taii-type2-using** *global-id*[:*prefix*[:*ac-id*]]
  - **template**
    - **vpls-sap-template**
    - **vpls-sap-template** *template-name*
    - **vpls-sap-template-using** *template-name*
    - **vpls-template**
    - **vpls-template** *template-name*
    - **vpls-template-using** *template-name*
- **connection-profile-vlan** [*con-prof-id*]
- **eth-tunnel** {**aps** | **status**}
- **eth-tunnel** *tunnel-index* [**path** *path-index*] [**detail**]
- **eth-tunnel**

2.20.1.1.1 ETH-CFM Show Commands

show

- **eth-cfm**
  - **association** [*ma-index*] [**detail**]

```
          — cfm-stack-table [port [port-id [vlan vlan-id]] | sdp sdp-id[:vc-id]] [level level]
               [direction up | down]
          — cfm-stack-table
          — cfm-stack-table port [{all-ports | all-sdps | all-virtuals}] [level level] [direction up |
               down]
          — cfm-stack-table port port-id [vlan qtag[.qtag]] [level level] [direction up | down]
          — cfm-stack-table sdp sdp-id[:vc-id] [level level] [direction up | down]
          — cfm-stack-table virtual service-id [level 0..7]
          — cfm-stack-table facility [{all-ports | all-lags | all-lag-ports | all-tunnel-meps | all-
               router-interfaces}] [level level] [direction up | down]
          — cfm-stack-table facility collect-lmm-stats
          — cfm-stack-table facility lag id [tunnel tunnel-id [level level] [direction up | down]
          — cfm-stack-table facility port id [level level] [direction up | down]
          — cfm-stack-table facility router-interface ip-int-name [level level] [direction up |
               down]
          — default-domain [bridge-identifier bridge-id vlan vlan-id]
          — domain [md-index] [association ma-index | all-associations] [detail]
          — lbm-svc-act-responder [domain md-index] [association ma-index] [mep mep-id]
          — mep mep-id domain md-index association ma-index [loopback] [linktrace] [eth-
               bandwidth-notification] [statistics]
          — mep mep-id domain md-index association ma-index remote-mepid mep-id | all-
               remote-mepids
          — mep mep-id domain md-index association ma-index eth-test-results [remote-peer
               mac-address]
          — mep mep-id domain md-index association ma-index one-way-delay-test [remote-
               peer mac-address]
          — mep mep-id domain md-index association ma-index two-way-delay-test [remote-
               peer mac-address]
          — mep mep-id domain md-index association ma-index two-way-slm-test [remote-
               peer mac-address]
          — mip
          — mip-instantiation [level level] [{sap sap-id | sdp sdp-id}]
          — system-config
     — eth-ring [status]
     — eth-ring [ring-index] hierarchy
     — eth-ring ring-index [path {a | b}]
```

2.20.1.1.2 PW-Port Show Commands

```
show
     — pw-port [pw-port-id] [detail]
     — pw-port sdp sdp-id
     — pw-port sdp none
     — pw-port sdp statistics
```

2.20.1.1.3 NGE Show Commands

```
show
     — group-encryption
```

```
          — encryption-keygroup keygroup-id [spi spi]
          — summary
```

2.20.1.2 Clear Commands

```
clear
     — group-encryption
          — encryption-keygroup keygroup-id [spi spi]
```

2.20.1.3 Tools Perform Commands

The following commands are applicable to the 7750 SR and 7450 ESS.

```
tools
     — perform
          — service
               — id service-id
                    — admin-lock
                         — pw
                              — sdp sdp-id:vc-id [test-service-id service-id] [start]
                    — loopback
                         — eth
                              — sap sap-id start mode [mac-swap] [mac ieee-address]
                                   [all]
                              — sap sap-id stop
                              — sdp sdp-id:vc-id start mode [mac-swap] [mac ieee-
                                   address] [all]
                              — sdp sdp-id:vc-id stop
                         — pw
                              — sdp sdp-id:vc-id {start | stop}
          — eth-ring
               — clear ring-index
               — force ring-index path {a | b}
               — manual ring-index path {a | b}
```

2.20.1.4 Tools Dump Commands

```
tools
     — dump
          — service
               — loopback
               — id service-id
                    — loopback sap sap-id
                    — loopback sdp sdp-id:vc-id
          — eth-ring ring-index [clear]
```

2.20.2 Command Descriptions

This section provides show command descriptions and output.

- Service Commands
- Connection Profile VLAN Commands
- ETH-CFM Show Commands
- NGE Show Commands
- Clear Commands

Note: For VLL, VPLS, and PBB show, clear, and debug commands, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN*.

For IES and VPRN show, clear, and debug commands, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN*.

2.20.2.1 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 to 2147483647

customer-site-name — Specifies the customer site name up to 32 characters in length which is an anchor point for an ingress and egress virtual scheduler hierarchy.

Output The following shows an example of customer information.

Table 15 describes the **show customer** command output fields:

Sample Output

```
*A:ALA-12# show service customer
==========================================================
Customers
```

```
==========================================================
Customer-ID : 1
Contact     : Manager
Description : Default customer
Phone       : (123) 555-1212

Customer-ID : 2
Contact     : Tech Support
Description : TiMetra Networks
Phone       : (234) 555-1212

Customer-ID : 3
Contact     : Test
Description : TiMetra Networks
Phone       : (345) 555-1212

Customer-ID : 6
Contact     : Test1
Description : Epipe Customer
Phone       : (456) 555-1212

Customer-ID : 7
Contact     : Test2
Description : VPLS Customer
Phone       : (567) 555-1212

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

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

Customer-ID : 94043
Contact     : Test Engineer on Duty
Description : TEST Customer
Phone       : (789) 555-1212
------------------------------------------------------
Total Customers : 8
-----------------------------------------------------------
*A:ALA-12#
*A:ALA-12# show service customer 274
==============================================================================
Customer  274
==============================================================================
Customer-ID : 274
Contact     : Mssrs. Beaucoup
Description : ABC Company
Phone       : 650 123-4567
------------------------------------------------------------------------------
Multi Service Site
------------------------------------------------------------------------------
Site        : west
Description : (Not Specified)
==============================================================================
*A:ALA-12#
```

```
*A:ALA-12# show service customer 274 site west
==============================================================================
Customer  274
==============================================================================
Customer-ID : 274
Contact     : Mssrs. Beaucoup
Description : ABC Company
Phone       : 650 123-4567
------------------------------------------------------------------------------
Multi Service Site
------------------------------------------------------------------------------
Site        : west
Description : (Not Specified)
Assignment  : Card 1
I. Sched Pol: SLA1
E. Sched Pol: (Not Specified)
------------------------------------------------------------------------------
Service Association
------------------------------------------------------------------------------
No Service Association Found.
==============================================================================
*A:ALA-12#
```

Table 15 **Service Commands Customer Field Descriptions**

Label	Description
Customer-ID	Displays the ID that uniquely identifies a customer.
Contact	Displays the name of the primary contact person.
Description	Displays generic information about the customer.
Phone	Displays the phone or pager number to reach the primary contact.
Total Customers	Displays the total number of customers configured.
Site	Displays the 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	Displays the port ID, MDA, or card number, where the SAP's that are members of this multi- service site are defined.
I. Sched Pol	Displays the ingress QoS scheduler policy assigned to this multi-service site.

Table 15 **Service Commands Customer Field Descriptions (Continued)**

Label	Description
E. Sched Pol	Displays the egress QoS scheduler policy assigned to this multi-service site.
Service-ID	Displays the ID that uniquely identifies a service.
SAP	Displays the SAP assigned to the service.

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 the MAC is aged out.

Output The following displays an example of service FDB MAC address information.

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#
```

id

Syntax **id** *service-id*
id *service-id* **base**
id *service-id* **bgp** [*bgp-instance*]
id *service-id* **mac-notification**
id *service-id* **macsec**
id *service-id* **mrp**
id *service-id* **mvrp vlan**
id *service-id* **mvrp vlan detail**
id *service-id* **sap base**
id *service-id* **sap mrp**
id *service-id* **sdp**
id *service-id* **vpls-group** [*vpls-group-id*]
id *service-id* **vpls-group** *vpls-group-id* **non-template-saps**

Context show>service

Description This command displays VPLS template information used to instantiate this service and m-VPLS that controls this service.

Output The following displays information about a specified service ID.

Sample Output

```
*A:SwSim100>config>service>epipe>sap$ show service id 10 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 10                  Vpn Id            : 0
Service Type      : Epipe
MACSec enabled: yes
Name              : 10
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 07/23/2018 17:55:04
Last Mgmt Change  : 07/23/2018 20:16:49
Test Service      : No
Admin State       : Down                Oper State        : Down
MTU               : 1514
Vc Switching      : False
SAP Count         : 2                   SDP Bind Count    : 0
Per Svc Hashing   : Disabled
Vxlan Src Tep Ip  : N/A
Force QTag Fwd    : Disabled
Oper Group        : <none>

-------------------------------------------------------------------------------
```

```
BGP Information
-------------------------------------------------------------------------------
No vxlan information for 10
-------------------------------------------------------------------------------
ETH-CFM service specifics
-------------------------------------------------------------------------------
Tunnel Faults     : ignore
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
No Matching Entries
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
SAP 1/1/c2/1:10
-------------------------------------------------------------------------------
Service Id         : 10
SAP                : 1/1/c2/1:10              Encap             : q-tag
MACSec Enabled: Yes, support 5, ecap-type dot1q:*, ca-1
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Down
Flags              : ServiceAdminDown
Multi Svc Site     : None
Last Status Change : 07/23/2018 17:55:04
Last Mgmt Change   : 07/23/2018 20:17:04
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU          : 1518                     Oper MTU          : 1518
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
qinq-pbit-marking  : both
Endpoint           : N/A
Egr Agg Rate Limit : max
Q Frame-Based Acct : Disabled                 Limit Unused BW   : Disabled
Vlan-translation   : None
Qinq-vlan-                                    Qinq-vlan-
translation        : None                     translation Ids   : None

Acct. Pol          : None                     Collect Stats     : Disabled

Oper Group         : (none)                   Monitor Oper Grp  : (none)
Host Lockout Plcy  : n/a
Ignore Oper Down   : Disabled
Lag Link Map Prof  : (none)
Cflowd             : Disabled
Bandwidth          : Not-Applicable
Oper DCpu Prot Pol*: _default-access-policy
-------------------------------------------------------------------------------
MACSec
-------------------------------------------------------------------------------
MACSec enabled: yes
==============================================================
PortId (SAP)                      MACSec Subport       encap-match    ca
-------------------------------------------------------------------------------
1/1/c2/
```

```
1:1                                        5                  dot1q:*             ca
-1
-------------------------------------------------------------------------------
Number of MACSec enabled SAPs : 1
-------------------------------------------------------------------------------
==============================================================


*A:Dut# show service id 1200 bgp (for VPLS service)
=======================================================================
BGP Information
=======================================================================
Vsi-Import           : None
Vsi-Export           : None
Route Dist           : auto-rd
Oper Route Dist      : 192.0.2.69:1200
Oper RD Type         : auto
Rte-Target Import    : 65000:1200         Rte-Target Export: 65000:1200
Oper RT Imp Origin   : configured         Oper RT Import   : 65000:1200
Oper RT Exp Origin   : configured         Oper RT Export   : 65000:1200
PW-Template Id       : None
-----------------------------------------------------------------------
======================================================================
*A:Dut#
*A:Dut# show service id 4096 bgp (for Epipe service)
=======================================================================
BGP Information
=======================================================================
Route Dist           : auto-rd
Oper Route Dist      : 192.0.2.69:1201
Oper RD Type         : auto
Rte-Target Import    : 65000:4096         Rte-Target Export: 65000:4096
PW-Template Id       : None
-----------------------------------------------------------------------
=======================================================================


*A:PE-6# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPLS
Name              : 1
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 05/08/2018 09:40:32
Last Mgmt Change  : 05/08/2018 09:40:24
Etree Mode        : Disabled
Admin State       : Up                  Oper State        : Up
MTU               : 1514
SAP Count         : 1                   SDP Bind Count    : 1
Snd Flush on Fail : Disabled            Host Conn Verify  : Disabled
SHCV pol IPv4     : None
Propagate MacFlush: Disabled            Per Svc Hashing   : Disabled
Allow IP Intf Bind: Disabled
Fwd-IPv4-Mcast-To*: Disabled            Fwd-IPv6-Mcast-To*: Disabled
Mcast IPv6 scope  : mac-based
Def. Gateway IP   : None
Def. Gateway MAC  : None
Temp Flood Time   : Disabled            Temp Flood        : Inactive
Temp Flood Chg Cnt: 0
SPI load-balance  : Disabled
TEID load-balance : Disabled
Src Tep IP        : N/A
Vxlan ECMP        : Disabled
VSD Domain        : <none>

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/c1/1:1                           q-tag        9000    9000    Up   Up
sdp:65:1 S(192.0.2.5)                    Spok         0       8974    Up   Down
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

provider-tunnel

Syntax **provider-tunnel**

Context show>service>id

Description This command displays provider tunnel information.

Output The following is an example of provider tunnel information.

Sample Output

```
A:PE-2# show service id 2000 provider-tunnel
===============================================================================
Service Provider Tunnel Information
===============================================================================
Type               : inclusive          Root and Leaf      : enabled
Admin State        : inService          Data Delay Intvl   : 15 secs
PMSI Type          : ldp                LSP Template       :
Remain Delay Intvl : 0 secs             LSP Name used      : 8193
PMSI Owner         : bgpEvpnMpls
===============================================================================
A:PE-2#


*A:Dut-B# /tools dump service id 1 provider-tunnels type originating
=======================================================================
VPLS 1 Inclusive Provider Tunnels Originating
=======================================================================
ipmsi (LDP)                                    P2MP-ID  Root-Addr
-----------------------------------------------------------------------
8193                                           8193    10.20.1.2
-----------------------------------------------------------------------


*A:Dut-B# /tools dump service id 1 provider-tunnels type terminating
```

```
=======================================================================
VPLS 1 Inclusive Provider Tunnels Terminating
=======================================================================
ipmsi (LDP)                                     P2MP-ID  Root-Addr
-----------------------------------------------------------------------
                                                8193    10.20.1.3
                                                8193    10.20.1.4
                                                8193    10.20.1.6
                                                8193    10.20.1.7
-----------------------------------------------------------------------


*A:Dut-B# /tools dump service id 1 provider-tunnels
=======================================================================
VPLS 1 Inclusive Provider Tunnels Originating
=======================================================================
ipmsi (LDP)                                     P2MP-ID  Root-Addr
-----------------------------------------------------------------------
8193                                            8193    10.20.1.2
-----------------------------------------------------------------------
=======================================================================
VPLS 1 Inclusive Provider Tunnels Terminating
=======================================================================
ipmsi (LDP)                                     P2MP-ID  Root-Addr
-----------------------------------------------------------------------
                                                8193    10.20.1.3
                                                8193    10.20.1.4
                                                8193    10.20.1.6
                                                8193    10.20.1.7
-----------------------------------------------------------------------
```

sdp

Syntax **sdp**

Context show>service>id

Description This command displays SDPs associated with this service.

Output The following is an example of SDP information associated with a service ID.

Sample Output

```
*A:Dut-C# show service id 1001 sdp 17407:4294967295 detail
========================================================================
Service Destination Point (Sdp Id : 17407:4294967295) Details
========================================================================
------------------------------------------------------------------------
Sdp Id 17407:4294967295  -(0.0.0.0)
------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 17407:4294967295         Type            : VplsPmsi
Split Horiz Grp    : (Not Specified)
VC Type            : Ether                    VC Tag            : n/a
```

```
Admin Path MTU     : 9194                     Oper Path MTU     : 9194
Far End            : not applicable           Delivery          : MPLS
Tunnel Far End     : n/a                    LSP Types         : None
Hash Label         : Disabled               Hash Lbl Sig Cap  : Disabled
Oper Hash Label    : Disabled
Admin State        : Up                     Oper State        : Up
Acct. Pol          : None                   Collect Stats     : Disabled
Ingress Label      : 0                        Egress Label      : 3
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Last Status Change : 01/31/2012 00:51:46      Signaling         : None
Last Mgmt Change   : 01/31/2012 00:49:58      Force Vlan-Vc     : Disabled
Endpoint           : N/A                      Precedence        : 4
PW Status Sig      : Enabled
Class Fwding State : Down
Flags              : None
Time to RetryReset : never                    Retries Left      : 3
Mac Move           : Blockable                Blockable Level   : Tertiary
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Application Profile: None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
MAC Learning       : Enabled                  Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled
Ignore Standby Sig : False                    Block On Mesh Fail: False
Oper Group         : (none)                   Monitor Oper Grp  : (none)
Rest Prot Src Mac  : Disabled
Auto Learn Mac Prot: Disabled                 RestProtSrcMacAct : Disable
Ingress Qos Policy : (none)                   Egress Qos Policy : (none)
Ingress FP QGrp    : (none)                   Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                   Egr Port QGrp Inst: (none)
------------------------------------------------------------------------
ETH-CFM SDP-Bind specifics
------------------------------------------------------------------------
V-MEP Filtering    : Disabled
KeepAlive Information :
Admin State        : Disabled                 Oper State        : Disabled
Hello Time         : 10                       Hello Msg Len     : 0
Max Drop Count     : 3                        Hold Down Time    : 10
Statistics            :
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                        I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 5937639                  E. Fwd. Octets    : 356258340
MCAC Policy Name   :
MCAC Max Unconst BW: no limit                 MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                        MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                        MCAC Avail Opnl BW: unlimited
-----------------------------------------------------------------------
RSVP/Static LSPs
------------------------------------------------------------------------
Associated LSP List :
No LSPs Associated
```

```
------------------------------------------------------------------------
Class-based forwarding :
------------------------------------------------------------------------
Class forwarding   : Disabled                EnforceDSTELspFc  : Disabled
Default LSP        : Uknwn                   Multicast LSP     : None
========================================================================
FC Mapping Table
========================================================================
FC Name             LSP Name
------------------------------------------------------------------------
No FC Mappings
------------------------------------------------------------------------
Stp Service Destination Point specifics
------------------------------------------------------------------------
Stp Admin State    : Down                    Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                     Port State        : Forwarding
Port Number        : 0                       Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : N/A
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                     Designated Port Id: N/A
Fwd Transitions    : 0                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
TC bit BPDUs rcvd  : 0                       TC bit BPDUs tx   : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
------------------------------------------------------------------------
Number of SDPs : 1
------------------------------------------------------------------------
========================================================================
*A:Dut-C#


A:Dut-B>config>service>vpls>bind# /show service id 1 sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
Sdp Id 120:1  -(10.20.1.1)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 120:1                   Type              : Spoke
Spoke Descr     : (Not Specified)
Split Horiz Grp    : (Not Specified)
Etree Root Leaf Tag: Disabled                Etree Leaf AC     : Disabled
VC Type            : Ether                   VC Tag            : n/a
Admin Path MTU     : 0                       Oper Path MTU     : 1570
Delivery           : MPLS
Far End            : 10.20.1.1
Tunnel Far End     : n/a                     LSP Types         : SR-TE
Hash Label         : Enabled                 Hash Lbl Sig Cap  : Enabled
Oper Hash Label    : Disabled
Entropy Label      : Disabled

Admin State        : Up                      Oper State        : Down
MinReqd SdpOperMTU : 1490
```

```
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 262130                   Egress Label      : 262135
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
BFD Template       : None
BFD-Enabled        : no                       BFD-Encap         : ipv4
Last Status Change : 07/15/2016 02:41:25      Signaling         : TLDP
Last Mgmt Change   : 07/15/2016 02:40:45
Endpoint           : N/A                      Precedence        : 4
PW Status Sig      : Enabled
Force Vlan-Vc      : Disabled                 Force Qinq-Vc     : Disabled
Class Fwding State : Down
Flags              : LabelStackLimitExceeded
Time to RetryReset : never                    Retries Left      : 3
Mac Move           : Blockable                Blockable Level   : Tertiary
Local Pw Bits      : pwNotForwarding
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing bfdFaultDet
Peer Vccv CC Bits  : mplsRouterAlertLabel

Application Profile: None
Transit Policy     : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
OAM MAC Addr       : 0                        DHCP MAC Addr     : 0
Host MAC Addr      : 0                        Intf MAC Addr     : 0
SPB MAC Addr       : 0                        Cond MAC Addr     : 0
BGP EVPN Addr      : 0                        EVPN Static Addr  : 0

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

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

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.      : 100                      I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 8800                     I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 112                      E. Fwd. Octets    : 9436


*A:Dut-C>config>router>mpls# /show service id 1 sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
Sdp Id 230:1  -(10.20.1.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 230:1                    Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type            : n/a                      VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 1578
Delivery           : MPLS
Far End            : 10.20.1.2
Tunnel Far End     : n/a                      LSP Types         : SR-TE
Hash Label         : Disabled                 Hash Lbl Sig Cap  : Disabled
Oper Hash Label    : Disabled
Entropy Label      : Disabled

Admin State        : Up                       Oper State        : Down
MinReqd SdpOperMTU : n/a
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 262134                   Egress Label      : 262138
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
BFD Template       : None
BFD-Enabled        : no                       BFD-Encap         : ipv4
Last Status Change : 07/21/2016 21:46:23      Signaling         : n/a
Last Mgmt Change   : 07/21/2016 21:42:38
Class Fwding State : Down
Flags              : LabelStackLimitExceeded
Local Pw Bits      : pwNotForwarding
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing bfdFaultDet
Peer Vccv CC Bits  : mplsRouterAlertLabel
Application Profile: None
Transit Policy     : None
AARP Id            : None
Ingress Qos Policy : (none)                   Egress Qos Policy : (none)
Ingress FP QGrp    : (none)                   Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                   Egr Port QGrp Inst: (none)
KeepAlive Information :
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.      : 22                       E. Fwd. Octets    : 1662
-------------------------------------------------------------------------------
Control Channel Status
-------------------------------------------------------------------------------
PW Status          : disabled                 Refresh Timer     : <none>
Peer Status Expire : false
Request Timer      : <none>
Acknowledgement    : false
-------------------------------------------------------------------------------
ETH-CFM SDP-Bind specifics
-------------------------------------------------------------------------------
```

```
Squelch Levels     : None

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
No LSPs Associated
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : Disabled                 EnforceDSTELspFc  : Disabled
Default LSP        : Uknwn                    Multicast LSP     : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings
-------------------------------------------------------------------------------
Segment Routing
-------------------------------------------------------------------------------
ISIS               : disabled
OSPF               : disabled
TE-LSP             : enabled
===============================================================================
SR-TE LSPs
===============================================================================
Lsp                             Admin   Oper    Time Since
                                                Last Trans
-------------------------------------------------------------------------------
LSP_CToB_1                      Up      Up      00h00m14s
===============================================================================
-------------------------------------------------------------------------------
Number of SDPs : 1
===============================================================================

A:Dut-B>config>service>vpls>bind# /show service id 1 sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
Sdp Id 120:1  -(10.20.1.1)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 120:1                    Type              : Spoke
Spoke Descr     : (Not Specified)
Split Horiz Grp    : (Not Specified)
Etree Root Leaf Tag: Disabled                 Etree Leaf AC     : Disabled
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 1570
Delivery           : MPLS
Far End            : 10.20.1.1
Tunnel Far End     : n/a                      LSP Types         : SR-TE
Hash Label         : Enabled                  Hash Lbl Sig Cap  : Enabled
Oper Hash Label    : Disabled
Entropy Label      : Disabled

Admin State        : Up                       Oper State        : Down
MinReqd SdpOperMTU : 1490
```

```
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 262130                   Egress Label      : 262135
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
BFD Template       : None
BFD-Enabled        : no                       BFD-Encap         : ipv4
Last Status Change : 07/15/2016 02:41:25      Signaling         : TLDP
Last Mgmt Change   : 07/15/2016 02:40:45
Endpoint           : N/A                      Precedence        : 4
PW Status Sig      : Enabled
Force Vlan-Vc      : Disabled                 Force Qinq-Vc     : Disabled
Class Fwding State : Down
Flags              : LabelStackLimitExceeded
Time to RetryReset : never                    Retries Left      : 3
Mac Move           : Blockable                Blockable Level   : Tertiary
Local Pw Bits      : pwNotForwarding
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing bfdFaultDet
Peer Vccv CC Bits  : mplsRouterAlertLabel
Application Profile: None
Transit Policy     : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
OAM MAC Addr       : 0                        DHCP MAC Addr     : 0
Host MAC Addr      : 0                        Intf MAC Addr     : 0
SPB MAC Addr       : 0                        Cond MAC Addr     : 0
BGP EVPN Addr      : 0                        EVPN Static Addr  : 0
MAC Learning       : Enabled                  Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled
Ignore Standby Sig : False                    Block On Mesh Fail: False
Oper Group         : (none)                   Monitor Oper Grp  : (none)
Auto Learn Mac Prot: Disabled
RestMacProtSrc Act : none
SendBvplsEvpnFlush : Disabled
Ingress Qos Policy : (none)                   Egress Qos Policy : (none)
Ingress FP QGrp    : (none)                   Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                   Egr Port QGrp Inst: (none)
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.      : 100                      I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 8800                     I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 112                      E. Fwd. Octets    : 9436
```

isid-using

Syntax **isid-using** [*range-id*]

Context show>service

Description This command displays services using the range ID.

Parameters *range-id* — Displays the service using the specified I-component Service ID (ISID).

Values 1 to 4294967295

Output The following is an example of services ISID information.

Sample Output

```
*A:SetupCLI# show service isid-using
=================================================================
Services
=================================================================
SvcId      ISID     Type   b-Vpls     Adm  Opr  SvcMtu CustId
-----------------------------------------------------------------
2001       122      i-VPLS 2002       Up   Down 1514   1
2005       2005     i-mVP* 2004       Down Down 1500   1
-----------------------------------------------------------------
Matching Services : 2
-----------------------------------------------------------------
*A:SetupCLI#
```

l2-route-table

Syntax **l2-route-table** [**detail**] [**bgp-ad**] [**multi-homing**] [**bgp-vpls**] [**bgp-vpws**] [**all-routes**]

Context show>service

Description This command displays Layer 2 route table information.

Parameters **all-routes** — Displays active or inactive routes.

detail — Displays detailed information.

Output The following is an example of Layer 2 route table information.

Sample Output

```
*A:PE-2# show service l2-route-table
========================================================================
Services: L2 Route Information - Summary
========================================================================
Svc Id L2-Routes (RD-Prefix) Next Hop Origin
Sdp Bind Id PW Temp Id
------------------------------------------------------------------------
1000 *192.0.2.3:60002-192.0.2.3 192.0.2.3 BGP-L2
32766:4294967293 1
------------------------------------------------------------------------
No. of L2 Route Entries: 1
========================================================================
===============================================================================
Services: L2 Multi-Homing Route Information - Summary
===============================================================================
Svc Id L2-Routes (RD-Prefix) Next Hop SiteId State DF
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of L2 Multi-Homing Route Entries: 0
===============================================================================
========================================================
Services: L2 Bgp-Vpls Route Information - Summary
========================================================
Svc Id L2-Routes (RD) Next Hop Ve-Id
Sdp Bind Id PW Temp Id
--------------------------------------------------------
1001 *192.0.2.3:60003 192.0.2.3 3
32765:4294967292 1
--------------------------------------------------------
No. of L2 Bgp-Vpls Route Entries: 1
========================================================
========================================================
Services: L2 Bgp-Vpws Route Information - Summary
========================================================
Svc Id L2-Routes (RD) Next Hop Ve-Id
Sdp Bind Id PW Temp Id
--------------------------------------------------------
1002 *192.0.2.3:60004 192.0.2.3 3
32764:4294967291 1
--------------------------------------------------------
No. of L2 Bgp-Vpws Route Entries: 1
========================================================
*A:PE-2# show service l2-route-table bgp-vpls
========================================================
Services: L2 Bgp-Vpls Route Information - Summary
========================================================
Svc Id L2-Routes (RD) Next Hop Ve-Id
Sdp Bind Id PW Temp Id
--------------------------------------------------------
1001 *192.0.2.3:60003 192.0.2.3 3
32765:4294967292 1
--------------------------------------------------------
No. of L2 Bgp-Vpls Route Entries: 1
========================================================
*A:PE-2# show service l2-route-table detail
===============================================================================
Services: L2 Route Information - Summary
===============================================================================
Svc Id : 1000
Origin : BGP-L2
PW Temp Id : 1
RD-Prefix : *192.0.2.3:60002-192.0.2.3
Next Hop : 192.0.2.3
Status : active
Sdp Bind Id : 32766:4294967293
===============================================================================
===============================================================================
Services: L2 Multi-Homing Route Information - Summary
===============================================================================
===============================================================================
===============================================================================
Services: L2 Bgp-Vpls Route Information - Summary
===============================================================================
Svc Id : 1001
```

```
VeId : 3
PW Temp Id : 1
RD : *192.0.2.3:60003
Next Hop : 192.0.2.3
State (D-Bit) : up(0)
Path MTU : 1514
Control Word : 0
Seq Delivery : 0
DF Bit : clear
Status : active
Sdp Bind Id : 32765:4294967292
===============================================================================
===============================================================================
Services: L2 Bgp-Vpws Route Information - Summary
===============================================================================
Svc Id : 1002
VeId : 3
PW Temp Id : 1
RD : *192.0.2.3:60004
Next Hop : 192.0.2.3
State (D-Bit) : up(0)
Path MTU : 1514
Control Word : 0
Seq Delivery : 0
Status : active
Tx Status : active
CSV : 0
Preference : 0
Sdp Bind Id : 32764:4294967291
===============================================================================
```

md-auto-id

Syntax **md-auto-id**

Context show>service

Description This command displays MD automatically-assigned ID information.

Output The following output is an example of MD automatically-assigned information, and Table 16 describes the MD automatic ID service output fields.

Sample Output

```
*A:node-6# show service md-auto-id
===============================================================================
MD Auto-Id Information
===============================================================================
Service-Id Range
Start                : 1073741823                  End        : 2147483647
Count                : 12
Customer-Id Range
Start                : 1073741823                  End        : 2147483647
```

```
Count              : 10
Pw-Template-Id Range
Start              : 1073741823                  End        : 2147483647
Count              : 5
===============================================================================
```

Table 16 **MD Auto-Id Service Output Fields**

Label	Description
Start	Specifies the start range for a service ID, customer ID, or PW template ID.
End	Specifies the end range for a service ID, customer ID, or PW template ID.
Count	Specifies the number of service IDs, customer IDs, or PW template IDs with automatically-assigned IDs.

oper-group

Syntax **oper-group** [*group-name*]
oper-group [*group-name*] **detail**
oper-group [*group-name*] **members** [**sap**] [**sdp**] [**site**]
oper-group [*group-name*] **monitoring** [**sap**] [**sdp**] [**site**] [**mvrp**]

Context show>service

Description This command displays oper-group information, member count, monitor-client count, and status in a single line for each of the configured oper-groups.

Parameters *group-name* — Displays oper-group information.

detail — Displays detailed information for each of the configured oper-groups.

members — Displays the members of the specified oper-group, or all oper-groups. A filter can be applied on the output to display only required member type, by specifying an optional parameter.

Values sap, sdp, site

monitoring — displays the clients that are monitoring the specified oper-group, or all oper-groups. A filter can be applied on the output to display only required client type, by specifying an optional parameter.

Values sap, sdp, site, mvrp

Output The following displays server oper group information.

Sample Output

```
*A:Dut-B#  show service oper-group
```

```
===============================================================================
Service Oper Group Information
===============================================================================
Name                             Oper   Creation Hold   Hold   Members Monitor
                                 Status Origin   UpTime DnTime
                                                 (secs) (secs)
-------------------------------------------------------------------------------
og-test                          up     manual     4      0       4       4
-------------------------------------------------------------------------------
Entries found: 1
=====================================================================
*A:Dut-B#


*A:Dut-B#  show service  oper-group  detail
=====================================================================
Service Oper Group Information
=====================================================================
Oper Group           : og-test
Creation Origin      : manual           Oper Status    : up
Hold DownTime        : 0 secs           Hold UpTime    : 4 secs
Members              : 4                Monitoring     : 4
=====================================================================
===================================================================
Member SDP-Binds for OperGroup: og-test
===================================================================
SdpId            SvcId      Type IP address      Adm     Opr
-------------------------------------------------------------------
201:1            1          Spok 10.20.1.1       Up      Up
201:2            1          Spok 10.20.1.1       Up      Up
-------------------------------------------------------------------
SDP Entries found: 4
===================================================================
===================================================================
Monitoring SDP-Binds for OperGroup: og-test
===================================================================
SdpId            SvcId      Type IP address      Adm     Opr
-------------------------------------------------------------------
205:1            1          Spok 10.20.1.5       Up      Up
205:2            1          Spok 10.20.1.5       Up      Up
-------------------------------------------------------------------
SDP Entries found: 4
===================================================================
*A:Dut-B#
```

pw-sap-using

Syntax **pw-sap-using**

Context show>service

Description This command displays service SAP PW port information.

Output The following example shows PW SAP port information.

Sample Output

```
===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
pw-1:0                          1          1     none    1     none   Up   Up
pw-1:1                          1          1     none    1     none   Up   Up
pw-2:2.1                        2          1     none    1     none   Up   Up
pw-2:0.*                        2          1     none    1     none   Up   Up
pw-2:1.*                        2          1     none    1     none   Up   Up
pw-3:3                          3          1     none    1     none   Up   Up
pw-4:4.*                        4          1     none    1     none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 7
-------------------------------------------------------------------------------
===============================================================================
```

pw-template

Syntax **pw-template** [*policy-id*]

Context show>service

Description This command displays PW template information.

Output The following example shows PW template information.

Sample Output

```
*A:Dut-B#    show service  pw-template 1
=======================================================================
PW Template Information
=======================================================================
PW Tmpl Id           : 1
Use Provisioned Sdp  : enabled              VcType           : vlan
Acctg Policy         : default              Collect Stats    : disabled
Mac-Learning         : enabled              Mac-Ageing       : enabled
Discard Unkn Src     : disabled             Limit MacMove    : blockable
Mac-Pinning          : disabled             Vlan VcTag       : 4095
MAC Address Limit    : no limit             Rest Prot Src Mac: disabled
Auto Learn Mac Prot  : disabled             RestProtSrcMacAct: disable
Block On Peer Fault  : disabled

SHG
Name                 :
Description          : (Not Specified)
Rest Prot Src Mac    : disabled             Rest Unprot Dst  : disabled
Auto Learn Mac Prot  : disabled             RestProtSrcMacAct: disable

Egress
Mac FilterId         : none                 Ip FilterId      : none
Ipv6 FilterId        : none                 QoS NetPlcyId    : none
```

```
Port RedirectQGrp    : none                 Instance Id      : none

Ingress
Mac FilterId         : none                 Ip FilterId      : none
Ipv6 FilterId        : none                 QoS NetPlcyId    : none
Fp RedirectQGrp      : none                 Instance Id      : none

IGMP
Fast Leave           : disabled             Import Plcy      : none
Last Memb Intvl      : 10 deci-secs         Max Nbr Grps     : 0
Send Queries         : disabled
Version              : 3

Force VlanVc Fwd     : disabled             Control Word     : disabled
Hash Label           : disabled             Hash Lbl Sig Cap : disabled
Last Changed         : 02/12/2013 22:11:49
-----------------------------------------------------------------------
Included SDP-Groups
-----------------------------------------------------------------------
red
-----------------------------------------------------------------------
```

saii-type2-using

Syntax **saii-type2-using** *global-id*[:*prefix*[:*ac-id*]]

Context show>service

Description This command displays the SDP used by a spoke SDP FEC with a specified FEC129 Type 2 SAII.

Parameters *global-id[:prefix[:ac-id]]* — Specifies the switch-point information using SAII-Type2.

Values <global-id[:prefix*> : <global-id>[:<prefix>[:<ac-id>]]

global-id	1 to 4294967295
prefix	a.b.c.d | 1 to 4294967295
ac-id	1 to 4294967295

Output The following example shows SAII information.

Sample Output

```
*A:Dut-E# show service saii-type2-using 3:10.20.1.3:1
===================================================================
Service Switch-Point Information
===================================================================
SvcId      Oper-SdpBind      SAII-Type2
-------------------------------------------------------------------
2147483598 17407:4294967195  3:10.20.1.3:1
-------------------------------------------------------------------
Entries found: 1
```

```
===================================================================
```

sap-using

Syntax **sap-using** [**msap**] [**dyn-script**] [**description**]
sap-using [**sap** *sap-id*] [**vlan-translation** | **anti-spoof**] [**description**]
sap-using {*ingress* | *egress*} **atm-td-profile** *td-profile-id*
sap-using {*ingress* | *egress*} **filter** *any-filter-id*
sap-using {*ingress* | *egress*} **qos-policy** *qos-policy-id* [**msap**]
sap-using etree

Context show>service

Description This command displays SAP information.

Parameters *sap-id* — The ID of the SAP to display.

td-profile-id — Displays SAPs using this traffic descriptor.

Values 1 to 1000

any-filter-id — The filter ID to display.

Values 1 to 65535

msap — Keyword to display MSAPs

vlan-translation — Keyword to display the SAPs where vlan-translation is enabled.

anti-spoof — Keyword to display the SAPs where anti-spoof can be enabled.

qos-policy-id — Keyword to display the SAPs that have a specific qos-policy applied.

Values 1 to 65535

dyn-script — Displays SAPs created by dynamic scripts.

Output The following is an example of SAP information.

Sample Output

```
show service sap-using eth-tunnel [tunnel-id ##]

*A:Dut-C># show service sap-using eth-tunnel
===============================================================================
Service Access Points (Ethernet Tunnel)
===============================================================================
SapId                                       SvcId    Path     Port          Tag
-------------------------------------------------------------------------------
eth-tunnel-1                                    50      1    1/1/2         4030
                                                        2    3/1/3         4031
eth-tunnel-2                                    51      1    3/1/1          100
                                                        2    3/1/3         4032
eth-tunnel-67                                   52      2    3/1/3          672
                                                        8    1/1/2          678
```

```
eth-tunnel-1:3                               3133       1    1/1/2           4
                                                        2    3/1/3           4
eth-tunnel-2:3                               3233       1    3/1/1           7
                                                        2    3/1/3           7
eth-tunnel-65:4094                           4094       2        -      4094.*
                                                        3    2/1/4      4094.*
                                                        8    1/1/3      4094.*
                                                       16    2/1/3      4094.*
eth-tunnel-1024:4094                         4094       1    2/1/1           -
                                                        2    3/1/2           -
eth-tunnel-1:4                               5154       1    1/1/2           5
                                                        2    3/1/3           5
eth-tunnel-2:4                               5254       1    3/1/1           8
                                                        2    3/1/3           8
eth-tunnel-1:5                               6165       1    1/1/2           6
                                                        2    3/1/3           6
eth-tunnel-2:5                               6265       1    3/1/1           9
                                                        2    3/1/3           9
eth-tunnel-65:3                             36533       3    2/1/4       65.10
                                                        8    1/1/3       65.10
                                                       16    2/1/3       65.10
eth-tunnel-66:3                             36633       2    2/1/4       66.13
                                                        4    1/1/3       66.13
eth-tunnel-67:3                             36733       2    3/1/3          16
                                                        8    1/1/2          16
eth-tunnel-68:3                             36833       2    3/1/3          19
                                                        3    3/1/1          19
eth-tunnel-65:4                             56554       3    2/1/4       65.11
                                                        8    1/1/3       65.11
                                                       16    2/1/3       65.11
eth-tunnel-66:4                             56654       2    2/1/4       66.14
                                                        4    1/1/3       66.14
eth-tunnel-67:4                             56754       2    3/1/3          17
                                                        8    1/1/2          17
eth-tunnel-68:4                             56854       2    3/1/3          20
                                                        3    3/1/1          20
eth-tunnel-65:5                             66565       3    2/1/4       65.12
                                                        8    1/1/3       65.12
                                                       16    2/1/3       65.12
eth-tunnel-66:5                             66665       2    2/1/4       66.15
                                                        4    1/1/3       66.15
eth-tunnel-67:5                             66765       2    3/1/3          18
                                                        8    1/1/2          18
eth-tunnel-68:5                             66865       2    3/1/3          21
                                                        3    3/1/1          21
-------------------------------------------------------------------------------
Number of SAPs : 23
```

This command can also be used to identify SAPs with the “EthTunTagMismatch” flag and can be used to prevent the flag from occurring before activating paths through the following CLI example:

```
*A:Dut-C> show service sap-using eth-tunnel | match "-"
-------------------------------------------------------------------------------
eth-tunnel-1                                   50       1    1/1/2        4030
eth-tunnel-2                                   51       1    3/1/1         100
eth-tunnel-67                                  52       2    3/1/3         672
```

```
eth-tunnel-1:3                                3133       1    1/1/2           4
eth-tunnel-2:3                                3233       1    3/1/1           7
eth-tunnel-65:4094                            4094       2        -      4094.*
eth-tunnel-1024:4094                          4094       1    2/1/1           -
                                                         2    3/1/2           -
...
eth-tunnel-65:3                              36533       3    2/1/4       65.10
eth-tunnel-66:3                              36633       2    2/1/4       66.13
...
```

SAP eth-tunnel-1024:4094 does not have the eth-tunnel tags configured for the corresponding paths which causes the SAP to be oper down. Ethernet tunnel 65 does not have path 2 configured. However, SAP eth-tunnel-65:4094 has a tag configured for path 2. This is acceptable and allows the operator to pre-provision tags under the same-fate SAPs before the corresponding path is configured under the Ethernet tunnel. This is the recommended configuration order so that there is no traffic disruption on the same-fate SAPs.

SAP eth-tunnel-65:5 has tags configured for paths 3, 8 and 16 and is operationally up.

If path 2 of Ethernet tunnel 65 was properly configured and active, SAP eth-tunnel-65:5 would be operationally down since it does not have a corresponding tag for path 2.

Any other tunnel is fine because it has no dash present in the port or tag location.

The **show eth-tunnel status** command summarizes the MEP status in one screen and also identifies the ports and tags associated in summary format for all loadsharing tunnels (similar to show eth-tunnel aps for g8031-1to1 mode).

```
show service sap-using eth-cfm squelch-ingress-levels [sap sap-id]
 <sap-id>           : null           - <port-id|lag-id>
                      dot1q          - <port-id|lag-id>:[qtag1|cp-conn-prof-id]
                      qinq           - <port-id|lag-id>:[qtag1|cp-conn-prof-
id].[qtag2]
                                       cp-conn-prof-id]
                        cp             - keyword
                        conn-prof-id   - [1..8000]
                      port-id        - slot/mda/port[.channel]
                                    eth-sat-id esat-id/slot/port
                                       esat: keyword
                                       id: 1 to 20
                                    pxc-idpxc-id.sub-port
                                       pxc pxc-id.sub-port
                                       pxc: keyword
                                       id: 1 to 64
                                       sub-port: a, b
                      eth-tunnel     - eth-tunnel-<id>[:<eth-tun-sap-id>]
                        id             - [1..128]
                        eth-tun-sap-id - [0..4094]
                      lag-id         - lag-<id>
                        lag            - keyword
                        id             - [1..200]
                      qtag1          - [0..4094]
                      qtag2          - [*|null|0..4094]
```

```
show service sap-using squelch-ingress-levels
```

```
===============================================================================
ETH-CFM Squelching
===============================================================================
SapId              SvcId      Squelch Level
-------------------------------------------------------------------------------
6/1/1:100.*        1          0 1 2 3 4 5 6 7
lag-1:100.*        1          0 1 2 3 4
6/1/1:200.*        2          0 1 2
lag-1:200.*        2          0 1 2 3 4 5
-------------------------------------------------------------------------------
Number of SAPs: 4
-------------------------------------------------------------------------------
===============================================================================

show service sdp-using eth-cfm squelch-ingress-levels [<sdp-id[:vc-id]>]
   <sdp-id[:vc-id]>     : sdp-id - [1..17407]
                          vc-id  - [1..4294967295]

show service sdp-using squelch-ingress-levels
================================================================================
ETH-CFM Squelching
================================================================================
SdpId              SvcId        Type Far End                 Squelch Level
-------------------------------------------------------------------------------
12345:4000000000   2147483650   Spok 1.1.1.1                0 1 2 3 4 5 6 7
===============================================================================

show service sap-using eth-cfm collect-lmm-stats
=========================================================
ETH-CFM SAPs Configured to Collect LMM Statistics
=========================================================
SapId                                        SvcId
---------------------------------------------------------
1/1/10:1000.*                                1000
---------------------------------------------------------
No. of SAPs: 1
=========================================================
```

sdp

Syntax **sdp** *sdp-id* **pw-port** [*pw-port-id*] [**statistics**]
sdp [**consistent** | **inconsistent** | **na**] **egressifs**
sdp *sdp-id* **keep-alive-history**
sdp far-end *ip-address* **keep-alive-history**
sdp [*sdp-id*] [**detail**]
sdp far-end *ip-address* [**detail**]

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 to 17407

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

detail — Displays detailed SDP information.

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

pw-port *pw-port-id* — Displays the SAP identifier for PW-SAPs.

Output The following example displays SDP information.

Table 17 describes the **show service SDP** output fields.

Sample Output

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

Admin State        : Up                       Oper State        : Up
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 0                        Egress Label      : 3
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Last Status Change : 12/14/2012 12:42:22      Signaling         : None
Last Mgmt Change   : 12/14/2012 12:42:19      Force Vlan-Vc     : Disabled
Endpoint           : N/A                      Precedence        : 4
PW Status Sig      : Enabled
Class Fwding State : Down
Flags              : None
Time to RetryReset : never                    Retries Left      : 3
Mac Move           : Blockable                Blockable Level   : Tertiary
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
```

```
Application Profile: None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0

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

Ingress Qos Policy : (none)                   Egress Qos Policy : (none)
Ingress FP QGrp    : (none)                   Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                   Egr Port QGrp Inst: (none)
-----------------------------------------------------------------------
ETH-CFM SDP-Bind specifics
-----------------------------------------------------------------------
V-MEP Filtering    : Disabled

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

Statistics            :
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                        I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 2979761                  E. Fwd. Octets    : 476761760
-----------------------------------------------------------------------
Control Channel Status
-----------------------------------------------------------------------
PW Status          : disabled                 Refresh Timer     : <none>
Peer Status Expire : false                    Clear On Timeout  : true

MCAC Policy Name   :
MCAC Max Unconst BW: no limit                 MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                        MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                        MCAC Avail Opnl BW: unlimited
-----------------------------------------------------------------------
RSVP/Static LSPs
-----------------------------------------------------------------------
Associated LSP List :
No LSPs Associated
-----------------------------------------------------------------------
Class-based forwarding :
-----------------------------------------------------------------------
Class forwarding   : Disabled                 EnforceDSTELspFc  : Disabled
Default LSP        : Uknwn                    Multicast LSP     : None
=======================================================================
FC Mapping Table
=======================================================================
FC Name             LSP Name
-----------------------------------------------------------------------
No FC Mappings
-----------------------------------------------------------------------
Stp Service Destination Point specifics
```

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

Fwd Transitions    : 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
TC bit BPDUs rcvd  : 0                        TC bit BPDUs tx   : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
-----------------------------------------------------------------------
Number of SDPs : 1
-----------------------------------------------------------------------
=======================================================================


A:Dut-F# show service sdp 1600 detail
===============================================================================
Service Destination Point (Sdp Id : 1600) Details
===============================================================================
-------------------------------------------------------------------------------
Sdp Id 1600  -2.2.2.2
-------------------------------------------------------------------------------
Description           : (Not Specified)
SDP Id               : 1600                  SDP Source         : manual
Admin Path MTU       : 0                     Oper Path MTU      : 1532
Delivery             : GRE
Far End              : 2.2.2.2               Tunnel Far End     : n/a
Oper Tunnel Far End  : 2.2.2.2
Local End            : 6.6.6.6
LSP Types            : n/a

Admin State          : Up                    Oper State         : Up
Signaling            : TLDP                  Metric             : 0
Acct. Pol            : None                  Collect Stats      : Disabled
Last Status Change   : 10/04/2018 21:18:06   Adv. MTU Over.     : No
Last Mgmt Change     : 10/04/2018 21:17:59   VLAN VC Etype      : 0x8100
Bw BookingFactor     : 100
Oper Max BW(Kbps)    : 0                     Avail BW(Kbps)     : 0
Net-Domain           : default               Egr Interfaces     : Consistent
Allow Fragmentation  : No
FPE LSP Id           : 0
Weighted ECMP        : Disabled
Flags                : None
Mixed LSP Mode Information :
Mixed LSP Mode       : n/a                   Active LSP Type    : 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
```

```
-------------------------------------------------------------------------------
MPLS-TP LSPs
-------------------------------------------------------------------------------
Associated LSP List :
SDP Delivery Mechanism is not MPLS
-------------------------------------------------------------------------------
Segment Routing
-------------------------------------------------------------------------------
ISIS                 : disabled
OSPF                 : disabled
TE-LSP               : disabled
===============================================================================


*A:Dut-B# show service sdp
============================================================================
Services: Service Destination Points
============================================================================
SdpId  AdmMTU  OprMTU  Far End          Adm  Opr         Del     LSP   Sig
----------------------------------------------------------------------------
230    0       1582    10.20.1.3        Up   Up          MPLS    I     TLDP
----------------------------------------------------------------------------
Number of SDPs : 1
----------------------------------------------------------------------------
Legend: R = RSVP, L = LDP, B = BGP, M = MPLS-TP, n/a = Not Applicable
============================================================================
*A:Dut-B#


*A:Dut-B# show service sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 230  -10.20.1.3
-------------------------------------------------------------------------------
Description           : (Not Specified)
SDP Id               : 230                   SDP Source         : manual
Admin Path MTU       : 0                     Oper Path MTU      : 1582
Delivery             : MPLS
Far End              : 10.20.1.3
Tunnel Far End       : n/a                   LSP Types          : SR-ISIS

Admin State          : Up                    Oper State         : Up
Signaling            : TLDP                  Metric             : 0
Acct. Pol            : None                  Collect Stats      : Disabled
Last Status Change   : 01/28/2015 22:00:07   Adv. MTU Over.     : No
Last Mgmt Change     : 01/28/2015 21:59:53   VLAN VC Etype      : 0x8100
Bw BookingFactor     : 100                   PBB Etype          : 0x88e7
Oper Max BW(Kbps)    : 0                     Avail BW(Kbps)     : 0
Net-Domain           : default               Egr Interfaces     : Consistent
Flags                : None

Mixed LSP Mode Information :
Mixed LSP Mode       : Disabled              Active LSP Type    : SR-ISIS

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

Src B-MAC LSB        : <none>                  Ctrl PW VC ID      : <none>
Ctrl PW Active       : n/a
-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
No LSPs Associated
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding     : Disabled                EnforceDSTELspFc   : Disabled
Default LSP          : Uknwn                   Multicast LSP      : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings
-------------------------------------------------------------------------------
Segment Routing
-------------------------------------------------------------------------------
ISIS                 : enabled                 LSP Id             : 524289
Oper Instance Id     : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:Dut-B#


*A:Dut-B> show service sdp
============================================================================
Services: Service Destination Points
============================================================================
SdpId  AdmMTU  OprMTU  Far End          Adm  Opr         Del     LSP   Sig
----------------------------------------------------------------------------
230    0       1582    10.20.1.3        Up   Up          MPLS    O     TLDP
----------------------------------------------------------------------------
Number of SDPs : 1
----------------------------------------------------------------------------
Legend: R = RSVP, L = LDP, B = BGP, M = MPLS-TP, n/a = Not Applicable
        I = SR-ISIS, O = SR-OSPF
============================================================================


*A:Dut-B> show service sdp 230 detail
===============================================================================
Service Destination Point (Sdp Id : 230) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 230  -10.20.1.3
-------------------------------------------------------------------------------
Description           : (Not Specified)
SDP Id               : 230                     SDP Source         : manual
Admin Path MTU       : 0                       Oper Path MTU      : 1582
```

```
Delivery             : MPLS
Far End              : 10.20.1.3
Tunnel Far End       : n/a                  LSP Types          : SR-OSPF

Admin State          : Up                   Oper State         : Up
Signaling            : TLDP                 Metric             : 0
Acct. Pol            : None                 Collect Stats      : Disabled
Last Status Change   : 05/27/2015 03:08:37  Adv. MTU Over.     : No
Last Mgmt Change     : 05/27/2015 03:05:36  VLAN VC Etype      : 0x8100
Bw BookingFactor     : 100                  PBB Etype          : 0x88e7
Oper Max BW(Kbps)    : 0                    Avail BW(Kbps)     : 0
Net-Domain           : default              Egr Interfaces     : Consistent
Flags                : None

Mixed LSP Mode Information :
Mixed LSP Mode       : Disabled             Active LSP Type    : SR-OSPF

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

Src B-MAC LSB        : <none>               Ctrl PW VC ID      : <none>
Ctrl PW Active       : n/a

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
No LSPs Associated

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

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

-------------------------------------------------------------------------------
Segment Routing
-------------------------------------------------------------------------------
OSPF                 : enabled              LSP Id             : 524289
Oper Instance Id     : 0
===============================================================================
*A:Dut-B>config>service>sdp#


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

*A:ALA-12#
===============================================================================
Service Destination Point (Sdp Id : 8) Details
===============================================================================
Sdp Id 8  -(10.10.10.104)
-------------------------------------------------------------------------------
Description           : MPLS-10.10.10.104
SDP Id               : 8
Admin Path MTU       : 0                     Oper Path MTU      : 0
Far End              : 10.10.10.104          Delivery           : MPLS
Admin State          : Up                    Oper State         : Down
Flags                : SignalingSessDown TransportTunnDown
Signaling            : TLDP                  VLAN VC Etype      : 0x8100
Last Status Change   : 02/01/2007 09:11:39   Adv. MTU Over.     : No
Last Mgmt Change     : 02/01/2007 09:11:46
KeepAlive Information :
Admin State          : Disabled              Oper State         : Disabled
Hello Time           : 10                    Hello Msg Len      : 0
Hello Timeout        : 5                     Unmatched Replies  : 0
Max Drop Count       : 3                     Hold Down Time     : 10
Tx Hello Msgs        : 0                     Rx Hello Msgs      : 0
Associated LSP LIST :
Lsp Name             : to-104
Admin State          : Up                    Oper State         : Down
Time Since Last Tran*: 01d07h36m
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:ALA-12#


*A:MV-SR12> show service sdp 10 detail
===============================================================================
Service Destination Point (Sdp Id : 10) Details
===============================================================================
Sdp Id 10  -(203.20.1.201)
-------------------------------------------------------------------------------
Description           : (Not Specified)
SDP Id               : 10                    SDP Source         : manual
Admin Path MTU       : 0                     Oper Path MTU      : 9182
Far End              : 203.20.1.201          Delivery           : MPLS/LDP
Admin State          : Up                    Oper State         : Up
Signaling            : TLDP                  Metric             : 0
Acct. Pol            : None                  Collect Stats      : Disabled
Last Status Change   : 02/12/2010 22:37:08   Adv. MTU Over.     : No
Last Mgmt Change     : 02/12/2010 22:37:03   VLAN VC Etype      : 0x8100
Bw BookingFactor     : 100                   PBB Etype          : 0x88e7
Oper Max BW(Kbps)    : 0                     Avail BW(Kbps)     : 0
Net-Domain           : default               Egr Interfaces     : Consistent
Mixed LSP Mode       : Enabled
Revert Time          : 0                     Revert Count Down  : n/a
Flags                : None

KeepAlive Information :
```

```
Admin State          : Disabled              Oper State         : Disabled
Hello Time           : 10                    Hello Msg Len      : 0
Hello Timeout        : 5                     Unmatched Replies  : 0
Max Drop Count       : 3                     Hold Down Time     : 10
Tx Hello Msgs        : 0                     Rx Hello Msgs      : 0
-------------------------------------------------------------------------------
LDP Information :
-------------------------------------------------------------------------------
LDP LSP Id           : 65539                 LDP Active         : No
-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP LIST :
Lsp Name             : To_7710
Admin State          : Up                    Oper State         : Up
Time Since Last Tran*: 01h20m56s
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding     : Disabled              EnforceDSTELspFc   : Disabled
Default LSP          : Uknwn                 Multicast LSP      : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:MV-SR12>config>service>vprn#


*B:Dut-B>config>router>mpls>lsp# /show service sdp
============================================================================
Services: Service Destination Points
============================================================================
SdpId  AdmMTU  OprMTU  Far End          Adm  Opr         Del     LSP   Sig
----------------------------------------------------------------------------
230    0       1578                     Up   Up          MPLS    I     TLDP
                       2001:db8::
----------------------------------------------------------------------------
Number of SDPs : 1
----------------------------------------------------------------------------
Legend: R = RSVP, L = LDP, B = BGP, M = MPLS-TP, n/a = Not Applicable
        I = SR-ISIS, O = SR-OSPF, T = SR-TE, F = FPE
============================================================================


*B:Dut-B>config>router>mpls>lsp# /show service sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
Sdp Id 230  2001:db8::
-------------------------------------------------------------------------------
Description           : Default sdp description
SDP Id               : 230                   SDP Source         : manual
Admin Path MTU       : 0                     Oper Path MTU      : 1578
Delivery             : MPLS
```

```
Far End              : 2001:db8::
Tunnel Far End       : n/a                   LSP Types          : SR-ISIS
Admin State          : Up                    Oper State         : Up
Signaling            : TLDP                  Metric             : 0
Acct. Pol            : None                  Collect Stats      : Disabled
Last Status Change   : 07/12/2016 19:40:17   Adv. MTU Over.     : No
Last Mgmt Change     : 07/12/2016 19:40:04   VLAN VC Etype      : 0x8100
Bw BookingFactor     : 100                   PBB Etype          : 0x88e7
Oper Max BW(Kbps)    : 0                     Avail BW(Kbps)     : 0
Net-Domain           : default               Egr Interfaces     : Consistent
FPE LSP Id           : 0
Flags                : None
Mixed LSP Mode Information :
Mixed LSP Mode       : Disabled              Active LSP Type    : SR-ISIS
KeepAlive Information :
Admin State          : Disabled              Oper State         : Disabled
Hello Time           : 10                    Hello Msg Len      : 0
Hello Timeout        : 5                     Unmatched Replies  : 0
Max Drop Count       : 3                     Hold Down Time     : 10
Tx Hello Msgs        : 0                     Rx Hello Msgs      : 0
Src B-MAC LSB        : <none>                Ctrl PW VC ID      : <none>
Ctrl PW Active       : n/a

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
No LSPs Associated
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding     : Disabled              EnforceDSTELspFc   : Disabled
Default LSP          : Uknwn                 Multicast LSP      : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings
-------------------------------------------------------------------------------
Segment Routing
-------------------------------------------------------------------------------
ISIS                 : enabled               LSP Id             : 524355
Oper Instance Id     : 0
OSPF                 : disabled
TE-LSP               : disabled
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================


*B:Dut-B>config>router>mpls>lsp# /show service id 1 sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
Sdp Id 230:1  -(2001:db8::)
-------------------------------------------------------------------------------
```

```
Description     : Default sdp description
SDP Id             : 230:1                    Type              : Spoke
Spoke Descr     : Description for Sdp Bind 230 for Svc ID 1
VC Type            : VLAN                     VC Tag            : 0
Admin Path MTU     : 0                        Oper Path MTU     : 1578
Delivery           : MPLS
Far End            : 2001:db8::
Tunnel Far End     : n/a                      LSP Types         : SR-ISIS
Hash Label         : Disabled                 Hash Lbl Sig Cap  : Disabled
Oper Hash Label    : Disabled
Entropy Label      : Disabled

Admin State        : Up                       Oper State        : Up
MinReqd SdpOperMTU : 1514
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 262134                   Egress Label      : 262134
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Admin BW(Kbps)     : 0                        Oper BW(Kbps)     : 0
BFD Template       : None
BFD-Enabled        : no                       BFD-Encap         : ipv4
Last Status Change : 07/12/2016 19:40:18      Signaling         : TLDP
Last Mgmt Change   : 07/12/2016 19:40:04
Endpoint           : N/A                      Precedence        : 4
PW Status Sig      : Enabled
Force Vlan-Vc      : Disabled                 Force Qinq-Vc     : Disabled
Class Fwding State : Down
Flags              : None
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing bfdFaultDet
Peer Vccv CC Bits  : mplsRouterAlertLabel
Application Profile: None
Transit Policy     : None
Standby Sig Slave  : False
Block On Peer Fault: False
Use SDP B-MAC      : False
Ingress Qos Policy : (none)                   Egress Qos Policy : (none)
Ingress FP QGrp    : (none)                   Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                   Egr Port QGrp Inst: (none)
KeepAlive Information :
Admin State        : Disabled                 Oper State        : Disabled
Hello Time         : 10                       Hello Msg Len     : 0
Max Drop Count     : 3                        Hold Down Time    : 10
Statistics            :
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                        I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                        E. Fwd. Octets    : 0
-------------------------------------------------------------------------------
Control Channel Status
-------------------------------------------------------------------------------
PW Status          : disabled                 Refresh Timer     : <none>
Peer Status Expire : false
Request Timer      : <none>
Acknowledgement    : false
-------------------------------------------------------------------------------
```

```
ETH-CFM SDP-Bind specifics
-------------------------------------------------------------------------------
Squelch Levels     : None

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
No LSPs Associated
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : Disabled                 EnforceDSTELspFc  : Disabled
Default LSP        : Uknwn                    Multicast LSP     : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings
-------------------------------------------------------------------------------
Segment Routing
-------------------------------------------------------------------------------
ISIS               : enabled                  LSP Id            : 524355
Oper Instance Id   : 0
OSPF               : disabled
TE-LSP             : disabled
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

When network domains are configured, the SDP egress interface state can be verified by using the following command:

```
*A:Dut-T# show service sdp egressifs
===============================================================================
SDP Egress Ifs State Table
===============================================================================
SDP Id              Network Domain                   State
-------------------------------------------------------------------------------
100                 net1                             consistent
-------------------------------------------------------------------------------
SDPs : 1
===============================================================================
*A:Dut-Tr#
*A:Dut-C># show service sdp 1 pw-port
============================================================================
Service Destination Point (Sdp Id 1 Pw-Port )
=============================================================================
SDP Binding port     : 1/1/3

SDP: 1 Pw-port: 11
----------------------------------------------------------------------------
VC-Id                : 11                   Admin Status       : up
Encap                : dot1q                Oper Status        : up
VC Type              : vlan                 Vlan VC Tag        : 0
Oper Flags           : (Not Specified)
```

```
SDP: 1 Pw-port: 44
-------------------------------------------------------------------------------
VC-Id                : 2                      Admin Status       : up
Encap                : dot1q                  Oper Status        : up
VC Type              : ether
Oper Flags           : (Not Specified)

----------------------------------------------------------------------------
Entries found: 2
-----------------------------------------------------------------------------
*A:Dut-C> #


*A:Dut-C> # show service sdp 1 pw-port 44
============================================================================
Service Destination Point (Sdp Id 1 Pw-Port 44)
============================================================================
SDP Binding port     : 1/1/3
VC-Id                : 2                      Admin Status       : up
Encap                : dot1q                  Oper Status        : up
VC Type              : ether
Oper Flags           : (Not Specified)
============================================================================
*A:Dut-C> #
```

The following show output gives the source-bmac-lsb and control PW used for a given SDP.

```
A:bksim1613# show service sdp 1 detail
===============================================================================
Service Destination Point (Sdp Id : 1) Details
===============================================================================
-------------------------------------------------------------------------------
Sdp Id 1  -2.2.2.2
-------------------------------------------------------------------------------
Description           : (Not Specified)
SDP Id               : 1                      SDP Source         : manual
Admin Path MTU       : 0                      Oper Path MTU      : 1556
Delivery             : MPLS
Far End              : 2.2.2.2
Tunnel Far End       : n/a                    LSP Types          : RSVP

Admin State          : Up                     Oper State         : Up
Signaling            : TLDP                   Metric             : 0
Acct. Pol            : None                   Collect Stats      : Disabled
Last Status Change   : 08/12/2013 06:33:57    Adv. MTU Over.     : No
Last Mgmt Change     : 08/12/2013 06:32:47    VLAN VC Etype      : 0x8100
Bw BookingFactor     : 100                    PBB Etype          : 0x88e7
Oper Max BW(Kbps)    : 0                      Avail BW(Kbps)     : 0
Net-Domain           : default                Egr Interfaces     : Consistent
Flags                : None

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

KeepAlive Information :
Admin State          : Disabled               Oper State         : Disabled
Hello Time           : 10                     Hello Msg Len      : 0
Hello Timeout        : 3                     Unmatched Replies  : 0
Max Drop Count       : 3                     Hold Down Time     : 10
Tx Hello Msgs        : 0                     Rx Hello Msgs      : 0
Src B-MAC LSB        : 00-13                 Ctrl PW VC ID      : 550
```

The following show output indicates whether use-sdp-bmac is applied to a given PW.

```
A:bksim1613# show service id 550 sdp 1:550 detail
===============================================================================
Service Destination Point (Sdp Id : 1:550) Details
===============================================================================
-------------------------------------------------------------------------------
Sdp Id 1:550  -(2.2.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 1:550                    Type               : Spoke
Spoke Descr     : (Not Specified)
VC Type            : Ether                    VC Tag             : n/a
Admin Path MTU     : 0                        Oper Path MTU      : 1556
Delivery           : MPLS
Far End            : 2.2.2.2
Tunnel Far End     : n/a                      LSP Types          : RSVP
Hash Label         : Disabled                 Hash Lbl Sig Cap   : Disabled
Oper Hash Label    : Disabled

Admin State        : Up                       Oper State         : Up
Acct. Pol          : None                     Collect Stats      : Disabled
Ingress Label      : 131048                   Egress Label       : 131063
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id    : n/a
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id     : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id   : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord   : False
Admin BW(Kbps)     : 0                        Oper BW(Kbps)      : 0
Last Status Change : 08/12/2013 06:33:57      Signaling          : TLDP
Last Mgmt Change   : 08/12/2013 06:32:47      Force Vlan-Vc      : Disabled
Endpoint           : N/A                      Precedence         : 4
PW Status Sig      : Enabled
Class Fwding State : Down
Flags              : None
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : mplsRouterAlertLabel

Application Profile: None
Transit Policy     : None
Standby Sig Slave  : False
Block On Peer Fault: False
Use sdp B-MAC      : True

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

KeepAlive Information :
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

-------------------------------------------------------------------------------
Control Channel Status
-------------------------------------------------------------------------------
PW Status          : disabled                 Refresh Timer     : <none>
Peer Status Expire : false
Request Timer      : <none>
Acknowledgement    : false

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
Lsp Name           : to-bksim1611-1
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 05h44m54s

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

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

-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

Table 17 **Service Commands SDP Field Descriptions**

Label	Description
SDP Id	Displays 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.

Table 17 **Service Commands SDP Field Descriptions (Continued)**

Label	Description (Continued)
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.
Local End	Displays the local end IP address.
Delivery	Displays the type of delivery used by the SDP: GRE or MPLS.
IP address	Displays the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Adm Admin State	Displays the desired state of the SDP.
Opr Oper State	Displays the operating state of the SDP.
Flags	Displays all the conditions that affect the operating status of this SDP.
Signal Signaling	Displays the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.
Last Status Change	Displays the time of the most recent operating status change to this SDP.
Adv. MTU Over	Specifies whether the advertised MTU of a VLL spoke SDP bind includes the 14-byte Layer 2 header.
Last Mgmt Change	Displays the time of the most recent management-initiated change to this SDP.
KeepAlive Information	Displays Keepalive information.
Hello Time	Displays how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	Displays the length of the SDP echo request messages transmitted on this SDP.
Hello Timeout	Displays the number of seconds to wait for an SDP echo response message before declaring a timeout.
Unmatched Replies	Displays 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.

Table 17 Service Commands SDP Field Descriptions (Continued)

Label	Description (Continued)
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. If the SDP type is GRE, the following message displays: SDP Delivery Mechanism is not MPLS.
Lsp Name	Displays the LSP name.
Time Since Last Transaction	Displays the time of the last transaction.
Signaling	Displays the signaling type.
Collect Stats	Specifies whether the agent collects accounting statistics for this SDP. When the value is true the agent collects accounting statistics on this SDP.
VLAN VC Etype	Displays the VLAN VC type.
BW Booking Factor	Displays 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 specified as **vlan** for Provider Backbone Bridging frames.
Oper Max BW (kb/s)	Displays the operational bandwidth in kilobits per seconds (kb/s) 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 (kb/s)	Displays the bandwidth that is still free for booking by the SDP bindings on the SDP.
Net-Domain	Displays 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.

Table 17 **Service Commands SDP Field Descriptions (Continued)**

Label	Description (Continued)
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	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.

sdp-group

Syntax **sdp-group** [*group-name*]

Context show>service

Description This show command will display the SDPs and the PW templates that are associated with the group-name.

Output The following is an example of SDP group information.

Sample Output

```
*A:Dut-B# show service sdp-group
==================================================
SDP Group Information
==================================================
Group                               Value
--------------------------------------------------
red                                 1
blue                                2
--------------------------------------------------
```

```
Entries found: 2
==================================================
*A:Dut-B#


*A:Dut-B#  show service sdp-group "red"
=======================================================================
SDP-Group Information
=======================================================================
Name               : red                Value              : 1

Associated SDPs
=======================================================================
SdpId              : 204                Sdp-Group          : red
SdpId              : 205                Sdp-Group          : red
-----------------------------------------------------------------------
Number of Entries: 2
=======================================================================
Associated pw-template included
=======================================================================
Pw-Template        : 1                  Sdp-Group          : red
-----------------------------------------------------------------------
Number of Entries: 1
=======================================================================
Associated pw-template excluded
=======================================================================
No Entries found
=======================================================================
*A:Dut-B#
```

sdp-group-using

Syntax **sdp-group-using**

Context show>service

Description This command displays groups using SDP.

Output The following is an example of information pertaining to objects using SDP groups.

Sample Output

```
*A:Dut-D#  show service sdp-group-using
=======================================================================
SDP-Group Information
=======================================================================
SdpId              : 402                Sdp-Group          : red
SdpId              : 405                Sdp-Group          : red
SdpId              : 4021               Sdp-Group          : blue
SdpId              : 4051               Sdp-Group          : blue

Associated pw-template included
=======================================================================
Pw-Template        : 1                  Sdp-Group          : red
Pw-Template        : 2                  Sdp-Group          : blue
```

```
Associated pw-template excluded
=======================================================================
No Entries found
=======================================================================
*A:Dut-D#
```

sdp-using

Syntax **sdp-using** [*sdp-id*[:*vc-id*] **| far-end** *ip-address*]
sdp-using *sdp-id*[:*vc-id*] **eth-cfm collect-lmm-stats**

Context show>service

Description This command displays services using SDP or far-end address options.

Parameters *sdp-id* — Displays only services bound to the specified SDP ID.

Values 1 to 17407

vc-id — The virtual circuit identifier.

Values 1 to 4294967295

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

Default Services with any far-end IP address.

eth-cfm collect-lmm-stats — Displays the LMM statistics for the specified MPLS SDP binding

Output The following example displays information about services using certain options.

Table 18 describes show service sdp-using output fields.

Sample Output

```
*A:Dut-A# show service sdp-using
===============================================================================
SDP Using
===============================================================================
SvcId       SdpId               Type Far End        Opr State I.Label  E.Label
-------------------------------------------------------------------------------
1           13:1                Spok 2001:db8::  Up        262130   262130
-------------------------------------------------------------------------------
Number of SDPs : 1
===============================================================================

*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
-------------------------------------------------------------------------------
Number of SDPs : 4
===============================================================================
*A:ALA-1#

show service sap-using eth-cfm squelch-ingress-levels [sap <sap-id>]
 <sap-id>           : null           - <port-id|lag-id>
                      dot1q          - <port-id|lag-id>:[qtag1|cp-conn-prof-id]
                      qinq           - <port-id|lag-id>:[qtag1|cp-conn-prof-id].
[qtag2|
                                       cp-conn-prof-id]
                        cp             - keyword
                        conn-prof-id   - [1..8000]
                      port-id        - slot/mda/port[.channel]
                                    eth-sat-id esat-id/slot/port
                                       esat: keyword
                                       id: 1 to 20
                                    pxc-idpxc-id.sub-port
                                       pxc pxc-id.sub-port
                                       pxc: keyword
                                       id: 1 to 64
                                       sub-port: a, b
                      eth-tunnel     - eth-tunnel-<id>[:<eth-tun-sap-id>]
                        id             - [1..128]
                        eth-tun-sap-id - [0..4094]
                      lag-id         - lag-<id>
                        lag            - keyword
                        id             - [1..200]
                      qtag1          - [0..4094]
                      qtag2          - [*|null|0..4094]


show service sap-using squelch-ingress-levels
===============================================================================
ETH-CFM Squelching
===============================================================================
SapId              SvcId      Squelch Level
-------------------------------------------------------------------------------
6/1/1:100.*        1          0 1 2 3 4 5 6 7
lag-1:100.*        1          0 1 2 3 4
6/1/1:200.*        2          0 1 2
lag-1:200.*        2          0 1 2 3 4 5
-------------------------------------------------------------------------------
Number of SAPs: 4
-------------------------------------------------------------------------------
===============================================================================

show service sdp-using eth-cfm squelch-ingress-levels [<sdp-id[:vc-id]>]
   <sdp-id[:vc-id]>     : sdp-id - [1..17407]
                          vc-id  - [1..4294967295]

show service sdp-using squelch-ingress-levels
================================================================================
ETH-CFM Squelching
================================================================================
SdpId              SvcId        Type Far End                 Squelch Level
```

```
-------------------------------------------------------------------------------
12345:4000000000   2147483650   Spok 1.1.1.1               0 1 2 3 4 5 6 7
===============================================================================

show service sdp-using eth-cfm collect-lmm-stats
===============================================================================
ETH-CFM SDPs Configured to Collect LMM Statistics
===============================================================================
SdpId               SvcId       Type    Far End
-------------------------------------------------------------------------------
1:1000              1000        spoke   1.1.1.31
-------------------------------------------------------------------------------
No. of SDPs: 1
===============================================================================
```

Table 18 **Service Commands SDP-Using Field Descriptions**

Label	Description
Svc ID	Displays the service identifier.
Sdp ID	Displays the SDP identifier.
Type	Displays the type of SDP: spoke or mesh.
Far End	Displays the far end address of the SDP.
Oper State	Displays the operational state of the service.
Ingress Label	Displays the label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	Displays the label used by this device to send packets to the far-end device in this service by this SDP.

service-using

Syntax **service-using** [**epipe**] [**ies**] [**vpls**] [**vprn**] [**mirror**] [**b-vpls**] [**i-vpls**] [**m-vpls**] [**apipe**] [**fpipe**] [**ipipe**] [**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.

The following parameters are applicable to the 7750 SR and 7450 ESS:

ies — Displays matching IES instances.

i-vpls — Displays matching I-VPLS instances.

b-vpls — Displays matching B-VPLS instances.

m-vpls — Displays matching M-VPLS instances.

mirror — Displays matching mirror services.

ipipe — Displays matching Ipipe services.

The following parameters are applicable to the 7750 SR:

apipe — Displays matching Apipe services.

fpipe — Displays matching Fpipe services.

vprn — Displays matching VPRN services.

sdp-id — Displays only services bound to the specified SDP ID.

Default Services bound to any SDP ID.

Values 1 to 17407

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

Default Services associated with a customer.

Values 1 to 2147483647

Output The following example shows service information.

Table 19 describes the show command output fields.

Sample Output

```
A:PE6# show service service-using
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm  Opr  CustomerId Service Name
-------------------------------------------------------------------------------
1            VPLS      Up   Up   1
11           VPRN      Up   Up   1
12           VPRN      Up   Up   1
70           Epipe     Up   Up   1
80           VPRN      Up   Up   1
100          Epipe     Up   Up   1
113          VPRN      Up   Up   1
600          VPLS      Down Down 1
601          VPRN      Down Down 1
4000         VPLS      Up   Up   1
4001         VPRN      Up   Up   1
5000         VPLS      Up   Up   1
5001         VPRN      Up   Up   1
6000         Epipe     Up   Up   1
6001         VPRN      Up   Up   1
2147483648   IES       Up   Down 1          _tmnx_InternalIesService
2147483649   intVpls   Up   Down 1          _tmnx_InternalVplsService
-------------------------------------------------------------------------------
Matching Services : 17
-------------------------------------------------------------------------------
===============================================================================
A:PE6#


*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#
```

The following output is applicable to the 7750 SR and 7450 ESS:

```
*A:SetupCLI# show service service-using
  - service-using [epipe] [ies] [vpls] [mirror] [ipipe] [b-vpls] [i-vpls]
[m-vpls] [sdp <sdp-id>] [customer <customer-id>]

 <epipe>              : keyword - displays epipe services
 <ies>                : keyword - displays ies services
 <vpls>               : keyword - displays vpls services
 <mirror>             : keyword - displays mirror services
 <ipipe>              : keyword - displays ipipe services
 <sdp-id>             : [1..17407] - display services using this sdp
 <customer-id>        : [1..2147483647] - display services using this customer
 <b-vpls>             : keyword - displays b-vpls services
 <i-vpls>             : keyword - displays i-vpls services
 <m-vpls>             : keyword - displays m-vpls services


*A:SetupCLI# show service service-using
===========================================================================
Services
===========================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
23           mVPLS     Up     Down       2                 09/25/2007 21:45:58
100          Epipe     Up     Down       2                 09/25/2007 21:45:58
101          Epipe     Up     Down       2                 09/25/2007 21:45:58
102          Epipe     Up     Down       2                 09/25/2007 21:45:58
105          Epipe     Up     Down       2                 09/25/2007 21:45:58
110          Epipe     Up     Down       1                 09/25/2007 21:45:58
990          IES       Up     Down       1                 09/25/2007 21:45:58
1000         Mirror    Up     Down       1                 09/25/2007 21:45:59
1001         Epipe     Up     Down       1                 09/25/2007 21:45:58
1002         Epipe     Up     Down       1                 09/25/2007 21:45:58
1003         Epipe     Up     Down       1                 09/25/2007 21:45:58
1004         Epipe     Up     Down       1                 09/25/2007 21:45:58
2000         Mirror    Up     Down       1                 09/25/2007 21:45:59
2001         i-VPLS    Up     Down       1                 09/25/2007 21:45:59
2002         b-VPLS    Up     Down       1                 09/25/2007 21:45:59
2003         i-VPLS    Down   Down       1                 09/25/2007 21:45:59
2004         b-mVPLS   Down   Down       1                 09/25/2007 21:45:59
2005         i-mVPLS   Down   Down       1                 09/25/2007 21:45:59
8787         IES       Up     Down       2                 09/25/2007 21:45:58
8888         IES       Up     Down       1                 09/25/2007 21:45:58
10000        IES       Down   Down       1                 09/25/2007 21:45:59
10001        VPLS      Up     Down       1                 09/25/2007 21:45:58
483000       Ipipe     Down   Down       2                 09/25/2007 21:45:59
483001       Ipipe     Up     Down       2                 09/25/2007 21:45:59
483004       Ipipe     Down   Down       2                 09/25/2007 21:45:59
483007       VPLS      Down   Down       2                 09/25/2007 21:45:59
483010       Ipipe     Down   Down       1                 09/25/2007 21:45:59
-------------------------------------------------------------------------------
Matching Services : 27
-------------------------------------------------------------------------------
*A:ALA-14#
```

Table 19 **Service Commands Service-Using Field Descriptions**

Label	Description
Service Id	Displays the service identifier.
Type	Displays the service type configured for the service ID.
Adm	Displays the desired state of the service.
Opr	Displays the operating state of the service.

Table 19 **Service Commands Service-Using Field Descriptions**

Label	Description (Continued)
CustomerID	Displays the ID of the customer who owns this service.
Last Mgmt Change	Displays the date and time of the most recent management-initiated change to this service.

system

Syntax **system**

Context show>service

Description This command enables the context to display service system information.

bgp-auto-rd

Syntax **bgp-auto-rd**

Context show>service>system

Description This command displays BGP auto route distinguisher (RD) information.

Output The following shows an example of BGP auto route distinguisher (RD) information

Sample Output

```
*A:Dut#show service system bgp-auto-rd
=======================================================================
Service BGP Auto Route Distinguisher Information
=======================================================================
IP address         : 192.0.2.69
Comm Val Start     : 1200                        End           : 1300
In Use             : 1
=======================================================================
```

bgp-route-distinguisher

Syntax **bgp-route-distinguisher** [**vprn**] [**vpls**] [**epipe**]
bgp-route-distinguisher svc
bgp-route-distinguisher ad-evi-rt-set
bgp-route-distinguisher system

Context show>service>system

Description This command displays the BGP operational route-distinguishers used by all the bgp-enabled services in the system. The information can be filtered by service: VPRN, VPLS. or Epipe. The output can also be filtered to show only the relevant route-distinguisher information related to services (**svc**), or the EVPN Auto-Discovery routes (**ad-evi-rt-set**), or the system route-distinguishers (**system**).

Output The following is an example of service BGP auto route distinguisher (RD) information.

Sample Output

```
*A:PE-2# show service system bgp-route-distinguisher
===============================================================================
Service Route Distinguishers
===============================================================================
Svc Id     Type  Oper Route-Distinguisher         Route-Distinguisher
-------------------------------------------------------------------------------
501        vprn  192.0.2.2:60000                  auto
800        vprn  192.0.2.2:60001                  auto
1          vpls  192.0.2.2:1                      configured
1          vpls  192.0.2.2:2                      configured
101        vpls  192.0.2.2:101                    configured
101        vpls  192.0.2.2:102                    configured
500        vpls  192.0.2.2:500                    derivedEvi
600        vpls  192.0.2.2:600                    derivedEvi
804        vpls  192.0.2.2:804                    derivedEvi
701        epipe 192.0.2.2:701                    derivedEvi
702        epipe 192.0.2.2:702                    derivedEvi
-------------------------------------------------------------------------------
Number of RD Entries: 11
===============================================================================
===============================================================================
Service System BGP Route Distinguisher Information
===============================================================================
                    Oper Route Distinguisher                        Type
-------------------------------------------------------------------------------
Auto-rd             192.0.2.2:60000-192.0.2.2:65000                 configured
Ethernet-segment    192.0.2.2:0                                     default
EVI RT Set RD Range <none>
===============================================================================
===============================================================================
BGP EVPN Ethernet Segment AD EVI RT Set Route Distinguishers
===============================================================================
Eth Seg                          EVI     Svc ID     Route Distinguisher
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of Entries: 0
===============================================================================
*A:PE-2# show service system bgp-route-distinguisher vpls
===============================================================================
Service Route Distinguishers
===============================================================================
Svc Id     Type  Oper Route-Distinguisher         Route-Distinguisher
-------------------------------------------------------------------------------
1          vpls  192.0.2.2:1                      configured
1          vpls  192.0.2.2:2                      configured
101        vpls  192.0.2.2:101                    configured
101        vpls  192.0.2.2:102                    configured
```

```
500        vpls  192.0.2.2:500                    derivedEvi
600        vpls  192.0.2.2:600                    derivedEvi
804        vpls  192.0.2.2:804                    derivedEvi
-------------------------------------------------------------------------------
Number of RD Entries: 7
===============================================================================
===============================================================================
Service System BGP Route Distinguisher Information
===============================================================================
                    Oper Route Distinguisher                        Type
-------------------------------------------------------------------------------
Auto-rd             192.0.2.2:60000-192.0.2.2:65000                 configured
Ethernet-segment    192.0.2.2:0                                     default
EVI RT Set RD Range <none>
===============================================================================
===============================================================================
BGP EVPN Ethernet Segment AD EVI RT Set Route Distinguishers
===============================================================================
Eth Seg                          EVI     Svc ID     Route Distinguisher
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of Entries: 0
===============================================================================
```

taii-type2-using

Syntax **taii-type2-using** *global-id*[:*prefix*[:*ac-id*]]

Context show>service

Description This command displays switch-point information using TAII.

Parameters *global-id[:prefix[:ac-id]]* — Specifies the switch-point information using SAII-Type2.

Values

<global-id[:prefix*> : <global-id>[:<prefix>[:<ac-id>]]

global-id	1 to 4294967295
prefix	a.b.c.d | 1 to 4294967295
ac-id	1 to 4294967295

Output The following is an example of service switch-point information using TAII information.

Sample Output

```
*A:Dut-E# show service taii-type2-using 6:10.20.1.6:1
===================================================================
Service Switch-Point Information
===================================================================
SvcId      Oper-SdpBind      TAII-Type2
-------------------------------------------------------------------
2147483598 17407:4294967195  6:10.20.1.6:1
-------------------------------------------------------------------
```

```
Entries found: 1
=====================================================
```

template

Syntax **template**

Context show>service

Description This command enables the context to display service template information.

vpls-sap-template

Syntax **vpls-sap-template**
vpls-sap-template *template-name*

Context show>service>template

Description This command displays basic information such as summary, template name, and so on, for all SAP VPLS-templates.

Output The following example displays VPLS SAP template information.

Sample Output

```
A:Dut-C# show service template vpls-sap-template squelch
===============================================================================
SAP template
===============================================================================
Template                           Saps           Last Update
-------------------------------------------------------------------------------
saptemplate                        30             07/26/2010 08:39:51
-------------------------------------------------------------------------------
Entries found: 1
===============================================================================
===============================================================================
SAP Template Information
===============================================================================
Template             : saptemplate          Discard Unkn Src : disabled
MAC Aging            : enabled              MAC Learning     : enabled
BPDU Translation     : disabled             MAC Address Limit: no limit
L2pt Termination     : disabled

STP
Admin Status         : up                   Port Priority    : 128
Port Path Cost       : 10                   Admin Edge       : disabled
Link Type            : Pt-pt
Auto Edge            : enabled              Root Guard       : disabled

MAC Move
Limit                : blockable            Limit Level      : tertiary
```

```
Ingress
QoS Policy           : 1                     MAC Fltr         : n/a
IP Fltr              : n/a                   QoS Sched Pol    : n/a
Match QinQ Dot1p Bits: default               Shared Q Pol     : n/a
IPv6 Fltr            : n/a
Use Multi-Pt Shared  : disabled              Agg Rate Limit   : Max
Policer Pol          : n/a

Egress
QoS Policy           : 1                     MAC Fltr         : n/a
IP Fltr              : n/a                   QoS Sched Pol    : n/a
IPv6 Fltr            : n/a                   QinQ Mark Top    : disabled
Agg Rate Limit       : Max                   Policer Pol      : n/a
Frame Based Acctg    : disabled

CPM Prot Plcy        : def                   CPM Monitor MAC  : disabled
Coll Acctg Stats     : disabled

ETH-CFM MIP          : disabled
ETH-CFM Squelch Level: 0 1 2 3 4 5
===============================================================================
```

vpls-sap-template-using

Syntax **vpls-sap-template-using** *template-name*

Context show>service>template

Description This command displays services instantiated using vpls-sap-template.

Output The following example displays information about services instantiated using this VPLS template.

Sample Output

```
A:Dut-C# show service template vpls-sap-template-using "saptemplate"
===============================================================================
SAP template 'saptemplate' created SAPs
===============================================================================
SvcId          Sap                                Creator Svc    Vpls Group
-------------------------------------------------------------------------------
1-10           2/1/2:1-2/1/2:10                   5000           1
               2/2/8:1-2/2/8:10
               lag-1:1.*-lag-1:10.*
-------------------------------------------------------------------------------
Entries found: 30
===============================================================================
```

vpls-template

Syntax **vpls-template**

vpls-template *template-name*

Context show>service>template>vpls-template

Description This command displays basic information/summary, template name, etc. for all VPLS templates. When a template name is specified, detailed information for the specified template, VPLS parameters, and so on, are displayed.

Output The following example displays VPLS template information.

Sample Output

```
A:Dut-C# show service template vpls-template
===============================================================================
Service template
===============================================================================
Template                           Services       Last Update
-------------------------------------------------------------------------------
test                              0               07/26/2010 08:40:01
svctemplate                       10              07/26/2010 08:39:51
-------------------------------------------------------------------------------
Entries found: 2
===============================================================================
A:Dut-C# show service template vpls-template "svctemplate"
===============================================================================
Service template Information
===============================================================================
Template             : svctemplate
MTU Size             : 1514                 Customer         : 10
MAC Aging            : enabled              MAC Learning     : enabled
Discard Unkn Dest    : disabled             Temp Flood Time  : Disabled
Per Svc Hashing      : disabled

FDB
Local Age Time       : 300 secs             Remote Age Time  : 900 secs
High Watermark       : 95%                  Low Watermark    : 90%
Table Size           : 250

STP
Admin State          : disabled             Priority         : 32768
Bridge Max Age       : 20 secs              Bridge Hello Time: 2 secs
Bridge Fwd Delay     : 15 secs              Mode             : rstp
Hold Cnt             : 6

MAC Move
Rate                 : 2/sec                Retry Timeout    : 10 secs
Admin State          : disabled             Num Retries      : 3
Pri-Ports Cumu Factor: 3                    Sec Cumu Factor  : 2
=================================================================================
```

vpls-template-using

Syntax **vpls-template-using** *template-name*

Context show>service>template

Description This command displays services instantiated using the VPLS-template.

Output The following example displays service template information.

Sample Output

```
A:Dut-C# show service template vpls-template-using "svctemplate"
===========================================================================
Service template 'svctemplate' created Services
===========================================================================
SvcId                    Creator Svc              Vpls Group
---------------------------------------------------------------------------
1-10                     5000                     1
---------------------------------------------------------------------------
Entries found: 10
===========================================================================
```

2.20.2.2 Connection Profile VLAN Commands

connection-profile-vlan

Syntax **connection-profile-vlan** [*conn-prof-id*]

Context show

Description This command displays information about the connection-profiles (VLAN) in the system. When a specific connection profile is shown, the vlan-ranges that it contains are displayed.

Parameters *conn-prof-id* — Specifies the VLAN connection profile ID.

Values 1 to 8000

Output The following is an example of connection profile VLAN information.

Sample Output

```
*A:Dut# show connection-profile-vlan
=======================================================================
Connection Profile Vlan Summary Information
=======================================================================
CP Index                                Number of Members
-----------------------------------------------------------------------
1                                       2
=======================================================================
*A:Dut# show connection-profile-vlan 1
=======================================================================
Connection Profile 1 Information
=======================================================================
Description : (Not Specified)
Last Change : 12/01/2015 16:50:34
=======================================================================
Connection Profile Vlan Eth Information
=======================================================================
Range Start         Range End                     Last Change
-----------------------------------------------------------------------
5                   100                           12/01/2015 16:50:34
150                 300                           12/01/2015 16:50:34
=======================================================================
```

2.20.2.3 ETH-CFM Show Commands

eth-cfm

Syntax **eth-cfm**

Context show

Description This command enables the context to display eth-cfm information.

eth-tunnel

Syntax **eth-tunnel**
eth-tunnel {**aps | status**}
eth-tunnel *tunnel-index* [**path** *path-index*] [**detail**]

Context show

Description This command displays Ethernet tunnel information. Any data SAP missing a tag for a defined path has the EthTunTagMismatch flag generated. In the example provided below, SAP eth-tunnel-1:1 does not have the tag for path 2 configured. Therefore, it is operationally down with the reason indicated by the EthTunTagMismatch flag.

Parameters *tunnel-index* — Specifies the tunnel index.

Values 1 to 1024

path-index — Specifies the path index.

Values 1 to 16

detail — Displays detailed information.

status — Displays Ethernet tunnel status information.

aps — Displays APS Ethernet tunnel information.

Output The following example of Ethernet status information.

Sample Output

```
*A:Dut-C>show>service>id# show eth-tunnel status
=============================================================================
Ethernet Tunnel Groups (Status information)
=============================================================================
Tunnel Admin  Oper       Member Information            MEP Information
ID     State  State  Path          Tag       State  Ctrl-MEP CC-Intvl Defects
-----------------------------------------------------------------------------
1      Up     Up      1 - 1/1/2    4030      Up     Yes      1         -----
                      2 - 3/1/3    4031      Up     Yes      1         -----
2      Up     Up      1 - 3/1/1    100       Up     Yes      1         -----
                      2 - 3/1/3    4032      Up     Yes      1         -----
65     Up     Up      3 - 2/1/4    65.4003   Up     -        -         -----
                      8 - 1/1/3    65.4008   Up     -        -         -----
                     16 - 2/1/3    65.4016   Up     -        -         -----
66     Up     Up      2 - 2/1/4    66.4002   Up     -        -         -----
                      4 - 1/1/3    66.4004   Up     -        -         -----
67     Up     Up      2 - 3/1/3    672       Up     Yes      1         -----
                      8 - 1/1/2    678       Up     Yes      1         -----
68     Up     Up      2 - 3/1/3    682       Up     -        -         -----
                      3 - 3/1/1    683       Up     -        -         -----
1024   Up     Up      1 - 2/1/1    1024      Up     -        -         -----
                      2 - 3/1/2    1024      Up     -        -         -----
=============================================================================
Ethernet Tunnel MEP Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM


*A:TOP_NODE# show eth-tunnel aps
===============================================================================
Ethernet Tunnel APS Groups
===============================================================================
Tunnel Admin  Oper   Working Path            Path   Active Rx PDU
ID     State  State  Protecting Path         State  Path   Tx PDU
-------------------------------------------------------------------------------
1      Up     Up      1 - 5/1/14   3070      Up     Yes    0F000000 (  NR)
                      2 - 2/1/9    3070      Up     No     0F000000 (  NR)
2      Up     Up      1 - 5/1/6    3071      Up     Yes    0F000000 (  NR)
                      2 - 2/1/13   3071      Up     No     0F000000 (  NR)
3      Up     Up      1 - 5/1/6    3072      Up     Yes    0F000000 (  NR)
                      2 - 2/1/13   3072      Up     No     0F000000 (  NR)
4      Up     Up      1 - 2/1/10   4.3073    Up     Yes    0F000000 (  NR)
                      2 - 2/1/4    4.3073    Up     No     0F000000 (  NR)
5      Up     Up      1 - 2/1/16   5.3074    Up     Yes    0F000000 (  NR)

show service id 3131 sap eth-tunnel-1:1

Flags              : EthTunTagMismatch
-------------------------------------------------------------------------------
SAP eth-tunnel-1:1
-------------------------------------------------------------------------------
Service Id         : 3131
SAP                : eth-tunnel-1:1           Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Down
Flags              : EthTunTagMismatch
Multi Svc Site     : None
Last Status Change : 01/13/2010 19:05:05
Last Mgmt Change   : 01/13/2010 17:01:33
Sub Type           : regular
```

```
Split Horizon Group: (Not Specified)

Admin MTU          : 2023                     Oper MTU          : 2023
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
                                              qinq-pbit-marking : both
Ing Agg Rate Limit : max                      Egr Agg Rate Limit: max
Endpoint           : N/A
Vlan-translation   : None

Acct. Pol          : None                     Collect Stats     : Disabled
Application Profile: None
-------------------------------------------------------------------------------
Eth-Tunnel Data Information
-------------------------------------------------------------------------------
Path               : 2                        Tag               : 1
```

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 to 4294967295

detail — Displays detailed information for the eth-cfm association.

Output The following example displays ETH CFM association information.

Table 20 describes show eth-cfm association command output fields:

Sample Output

```
*A:node-1# show eth-cfm association
======================================================================
eth-cfm CFM Association Table
======================================================================
Md-index   Ma-index   Name                     CCM-interval Bridge-id
----------------------------------------------------------------------
1          1          test-ma-1                10           2
1          2          2                        10           20
======================================================================
*A:node-1#


*A:node-1# show eth-cfm association 1 detail
-------------------------------------------------------------------------------
Domain 1 Associations:
-------------------------------------------------------------------------------
Md-index          : 1                        Ma-index          : 1
```

```
Name Format       : charString              CCM-interval      : 10
Name              : test-ma-1
Bridge-id         : 2                       MHF Creation      : defMHFnone
PrimaryVlan       : 0                       Num Vids          : 0
Remote Mep Id     : 1
Remote Mep Id     : 4
Remote Mep Id     : 5
-------------------------------------------------------------------------------
*A:node-1#
```

Table 20 **ETH-CFM Association Field Descriptions**

Label	Description
Md-index	Displays the maintenance domain (MD) index.
Ma-index	Displays the maintenance association (MA) index.
Name	Displays the part of the maintenance association identifier 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

cfm-stack-table

Syntax **cfm-stack-table**
cfm-stack-table [{**all-ports** | **all-sdps** | **all-virtuals**}] [**level** *level*] [**direction** {**up** | **down**}]
cfm-stack-table port *port-id* [**vlan** *qtag*[.*qtag*]] [**level** *level*] [**direction** {**up** | **down**}]
cfm-stack-table sdp *sdp-id*[:*vc-id*] [**level** *level*] [**direction** {**up** | **down**}]
cfm-stack-table virtual *service-id* [**level** *level*]
cfm-stack-table facility [{**all-ports** | **all-lags** | **all-lag-ports** | **all-tunnel-meps** | **all-router-interfaces**}] [**level** *level*] [**direction** {**up** | **down}]**
cfm-stack-table facility collect-lmm-stats
cfm-stack-table facility lag *id* [**tunnel** *tunnel*-id] [**level** *level*] [**direction** {**up** | **down**}]
cfm-stack-table facility port *id* [**level level**] [**direction** {**up** | **down**}]
cfm-stack-table facility router-interface *ip-int-name* [**level** *level*] [**direction** {**up** | **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 options allow the operator to be specific. If no parameters are include then the entire stack-table will be displayed.

Parameters *port-id* — Specifies a bridge port or aggregated port on which MEPs or MHFs are configured.

vlan-id — Specifies an associated VLAN ID to be displayed.

sdp-id[:*vc-id*] — Specifies an SDP for which CFM stack table information will be displayed.

level — Specifies the MD level of the maintenance point.

Values 0 to 7

direction {**up** | **down**} — Specifies the direction in which the MP faces on the bridge port.

facility — Keyword to display the CFM stack table information for facility MEPs. The base command will display all facility MEPs. Options may be included in order to further parse the table for specific facility MEP information.

service-id — Specifies an SDP for which CFM stack table information will be displayed.

tunnel-id — Specifies the tunnel ID.

Values 1 to 4094

Output The following example displays eth-cfm CFM stack table information.

Table 21 describes the show eth-cfm CFM stack table command output fields:

Sample Output

```
show eth-cfm cfm-stack-table
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM
A = AisRx, L = CSF LOS Rx, F = CSF AIS/FDI rx, r = CSF RDI rx
===============================================================================
CFM SAP Stack Table
===============================================================================
Sap                Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect
-------------------------------------------------------------------------------
1/1/6:20.0           4 B         14        803  MIP d8:1c:01:01:00:06   -------
1/1/6:3000.1001      4 B         14        800  MIP 00:00:00:00:00:28   -------
1/1/6:2000.1002      4 B         14        802  MIP d8:1c:01:01:00:06   -------
1/1/6:0.*            4 B         14        805  MIP d8:1c:01:01:00:06   -------
1/1/9:300            2 U         12        300   28 00:00:00:00:00:28   -------
1 to 4094
1/1/9:401            2 U         12        401   28 00:00:00:00:00:28   -------
1/1/9:600            2 U         12        600   28 00:00:00:00:00:28   -------
1/1/10:4.*           2 U         12          4   28 00:00:00:00:00:28   --C----
1/1/10:1000.*        5 U         15       1000   28 00:00:00:00:00:28   -------
1/1/10:1001.*        5 U         15       1001   28 00:00:00:00:00:28   -------
1/2/1:2000.2000      4 B         14       2000  MIP 00:00:00:00:01:28   -------
```

```
1/2/1:3000.3000      4 B          0          0  MIP d8:1c:01:02:00:01   -------
===============================================================================

===============================================================================
CFM Ethernet Tunnel Stack Table
===============================================================================
Eth-tunnel         Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
===============================================================================
CFM Ethernet Ring Stack Table
===============================================================================
Eth-ring           Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
===============================================================================
CFM Facility Port Stack Table
===============================================================================
Port     Tunnel    Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect
-------------------------------------------------------------------------------
1/2/4     0          0 D         10          1   28 00:00:00:00:00:28   -------
===============================================================================
===============================================================================
CFM Facility LAG Stack Table
===============================================================================
Lag      Tunnel    Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
===============================================================================
CFM Facility Tunnel Stack Table
===============================================================================
Port/Lag Tunnel    Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
===============================================================================
CFM Facility Interface Stack Table
===============================================================================
Interface          Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect
-------------------------------------------------------------------------------
v28-v33              1 D         11          1   28 00:00:00:00:00:28   -------
===============================================================================
===============================================================================
CFM SAP Primary VLAN Stack Table
===============================================================================
Sap
  Primary VlanId   Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect
-------------------------------------------------------------------------------
1/1/6:20.*
      21             4 B         14        804  MIP d8:1c:01:01:00:06   -------
===============================================================================
===============================================================================
CFM SDP Stack Table
===============================================================================
Sdp                Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect
-------------------------------------------------------------------------------
```

```
1:1000               4 D         14       1000   28 00:00:00:00:00:28   -------
2:777                4 D         14        777   28 d8:1c:ff:00:00:00   -------
400:800              4 B         14        800  MIP 00:00:00:00:01:28   -------
===============================================================================
===============================================================================
CFM Virtual Stack Table
===============================================================================
Service            Lvl Dir Md-index   Ma-index  MepId  Mac-address     Defect
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

Table 21 ETH-CFM CFM Stack Table Field Descriptions

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 maintenance domain (MD) index.
Ma-index	Displays the maintenance association (MA) index.
Mep-id	Displays the integer that is unique among all the MEPs in the same MA.
Mac-address	Displays the MAC address of the MP.

default-domain

Syntax **default-domain** [**bridge-identifier** *bridge-id* **vlan** *vlan-id*]

Context show>eth-cfm

Description This command displays per-MIP index (**bridge-identifier** and **vlan**) configuration as entered under the **default-domain** entries.

Parameters *bridge-id* — The bridge identifier related to the MIP. This is equivalent to the *service-id*.

vlan-id — The VLAN ID matching the primary VLAN, or “none” if **primary-vlan-enable** is not configured.

Output The following is an example of default domain information.

Table 22 describes the show default domain command output fields.

Sample Output

```
show eth-cfm default-domain
===============================================================================
Default Domain Information
===============================================================================
System Settings
MHF Creation  : none                    Level           : 0
Id Permission : none                    MIP Ltr Priority : 7
===============================================================================
BridgeId     VLAN   Valid   Level   MhfCreation   IdPermission  LtrPriority
-------------------------------------------------------------------------------
2000         none   true      3       default       none          defer
===============================================================================
```

Table 22 ETH-CFM Default Domain Field Descriptions

Label	Description
Valid	Indicates whether the row is valid and can be used for MIP creation. It does not indicate whether the row is being used to create the specific MIP. The show command **eth-cfm mip-instantiation** shows the authoritative creation routine.
Level	Displays the configured level value
MhfCreation	Displays the configured **mhf-creation** mode
IdPermission	Displays the configured ID permission action
LtrPriority	Displays the configured MIP LTR priority

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.

ma-index — Displays the index to which the MP is associated, or 0, if none.

all-associations — Displays all associations to the MD.

detail — Displays detailed domain information.

Output The following is an example of ETH CFM domain information.

Table 23 describes the show eth-cfm domain command output fields:

Sample Output

```
*A:node-1# show eth-cfm domain
==============================================================================
eth-cfm CFM Domain Table
```

```
==============================================================================
Md-index   Level Name                                       Format
------------------------------------------------------------------------------
1          4     test-1                                     charString
7          4     AA:BB:CC:DD:EE:FF-0                        macAddressAndUint
==============================================================================
*A:node-1#

*A:node-1# show eth-cfm domain 1 detail
===============================================================================
Domain 1
Md-index          : 1                        Level            : 4
Permission        : sendIdNone               MHF Creation     : defMHFnone
Name Format       : charString               Next Ma Index    : 3
Name              : test-1
===============================================================================
*A:node-1#
```

Table 23 **ETH-CFM Domain Field Descriptions**

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*.

lbm-svc-act-responder

Syntax **lbm-svc-act-responder** [**domain** *md-index*] [**association** *ma-index*] [**mep** *mep-id*]

Context show>eth-cfm

Description This command displays all the MEPs that have been created with this optional parameter, which allocates additional resources to facilitate high-speed LBM-to-LBR processing typically used during service activation testing. The optional filters are cumulative. These filters can be used to narrow the focus of the display to a specific area.

Parameters *md-index* — Displays the MD index.

Values 1 to 4294967295

ma-index — Displays the MA index.

Values 1 to 4294967295

mep-id — Displays the local MEP ID.

Values 1 to 8191

Output The following is an example of ETH CFM **lbm-svc-act-responder** information.

Table 24 describes the output fields:

Sample Output

```
show eth-cfm lbm-svc-act-responder
==============================================================================
Eth-CFM Local MEP LBM Service Activation Responder Enabled
==============================================================================
MdIndex    MaIndex    MepId  SrcMacAddress
------------------------------------------------------------------------------
14         1000       28     d8:1c:ff:00:00:00
------------------------------------------------------------------------------
No. of MEPs: 1
==============================================================================
```

Table 24 **ETH-CFM lbm-service-act-responder Field Descriptions**

Label	Description
MdIndex	Displays the Maintenance Domain (MD) index value
MaIndex	Displays the MA index value
MepId	Displays the maintenance association endpoint identifier
SrcMacAddress	Displays the source MAC address

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**loopback**] [**linktrace**] [**eth-bandwidth-notification**] [**statistics**]

mep *mep-id* **domain** *md-index* **association** *ma-index* **remote-mepid** *mep-id* | **all-remote-mepids**

mep *mep-id* **domain** *md-index* **association** *ma-index* **eth-test-results** [**remote-peer** *mac-address*]

mep *mep-id* **domain** *md-index* **association** *ma-index* **one-way-delay-test** [**remote-peer** *mac-address*]

mep *mep-id* **domain** *md-index* **association** *ma-index* **two-way-delay-test** [**remote-peer** *mac-address*]

mep *mep-id* **domain** *md-index* **association** *ma-index* **two-way-slm-test** [**remote-peer** *mac-*

address]

Context show>eth-cfm

Description This command displays Maintenance Endpoint (MEP) information.

Parameters *mep-id* — Displays the integer that is unique among all the MEPs in the same MA.

md-index — Displays the index of the MD to which the MP is associated, or 0, if none.

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.

eth-bandwidth-notification — Displays the active eth-bn notification parameters received form the peer and reported to the rate function on the associated port.

statistics — Includes specified statistic counter information for the specified MEP.

remote-mepid *mep-id* — Includes specified remote MEP ID information for specified the MEP.

all-remote-mepids — Includes all remote MEP ID information for the specified MEP.

eth-test-results — Includes eth-test-result information for the specified MEP.

one-way-delay-test — Includes one-way-delay-test information for the specified MEP.

two-way-delay-test — Includes two-way-delay-test information for the specified MEP.

two-way-slm-test — Includes two-way-slm-test information for the specified MEP.

remote-peer *mac-address* — Includes specified remote MEP ID information for the specified MEP.

Output The following is an example of ETH CFM MEP information.

Sample Output

```
# show eth-cfm mep 101 domain 3 association 1
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index           : 3                        Direction         : Down
Ma-index           : 1                        Admin             : Enabled
MepId              : 101                      CCM-Enable        : Enabled
IfIndex            : 1342177281               PrimaryVid        : 6553700
Description        : (Not Specified)
FngState           : fngReset                 ControlMep        : False
LowestDefectPri    : macRemErrXcon            HighestDefect     : none
Defect Flags       : None
Mac Address        : d0:0d:1e:00:01:01        ControlMep        : False
CcmLtmPriority     : 7
CcmTx              : 19886                    CcmSequenceErr    : 0
Fault Propagation  : disabled                 FacilityFault     : n/a
MA-CcmInterval     : 1                        MA-CcmHoldTime    : 0ms
Eth-1Dm Threshold  : 3(sec)                   MD-Level          : 3
Eth-Ais:           : Enabled                  Eth-Ais Rx Ais:   : No
```

```
Eth-Ais Tx Priorit*: 7                        Eth-Ais Rx Interv*: 1
Eth-Ais Tx Interva*: 1                        Eth-Ais Tx Counte*: 388
Eth-Ais Tx Levels  : 5
Eth-Tst:           : Disabled

Redundancy:
    MC-LAG State   : active

CcmLastFailure Frame:
    None

XconCcmFailure Frame:
    None
===============================================================================

show eth-cfm mep 607 domain 6 association 607
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index           : 6                        Direction         : Down
Ma-index           : 607                      Admin             : Enabled
MepId              : 607                      CCM-Enable        : Enabled
IfIndex            : 1342177283               PrimaryVid        : 268369927
Description        : (Not Specified)
FngState           : fngReset                 ControlMep        : False
LowestDefectPri    : macRemErrXcon            HighestDefect     : none
Defect Flags       : None
Mac Address        : 8c:d3:ff:00:01:43        ControlMep        : False
CcmLtmPriority     : 7
CcmTx              : 78122                    CcmSequenceErr    : 0
Fault Propagation  : useIfStatusTLV           FacilityFault     : n/a
MA-CcmInterval     : 1                        MA-CcmHoldTime    : 0ms
Eth-1Dm Threshold  : 3(sec)                   MD-Level          : 6
Eth-Ais:           : Disabled
Eth-Tst:           : Disabled

Redundancy:
    MC-LAG State   : n/a

CcmLastFailure Frame:
    None

XconCcmFailure Frame:
    None
===============================================================================
show eth-cfm association
===============================================================================
CFM Association Table
===============================================================================
Md-index   Ma-index   Name                     CCM-intrvl Hold-time Bridge-id
-------------------------------------------------------------------------------
2          106        MA-0000000106            1          n/a       none
2          207        MA-0000000207            1          n/a       none
2          308        MA-0000000308            1          n/a       none
3          1          ma-0000000001            1          n/a       none
3          2          ma-0000000002            1          n/a       none
3          3          ma-0000000003            1          n/a       none
3          4          ma-0000000004            1          n/a       none
3          5          ma-0000000005            1          n/a       none
```

```
5          555        MA-0000000555            10         n/a       47
6          607        MA-0000000607            1          n/a       207
7          707        MA-0000000707            1          n/a       207
===============================================================================

# show eth-cfm mep 28 domain 10 association 1 eth-bandwidth-notification
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index         : 10                       Direction         : Down
Ma-index         : 1                        Admin             : Disabled
MepId            : 28                       CCM-Enable        : Disabled
Port             : 1/1/1                    VLAN              : 0
Description      : (Not Specified)
FngAlarmTime     : 0                        FngResetTime      : 0
FngState         : fngReset                 ControlMep        : False
LowestDefectPri  : macRemErrXcon            HighestDefect     : none
Defect Flags     : None
Mac Address      : d8:1c:01:01:00:01        Collect LMM Stats : disabled
LMM FC Stats     : None
LMM FC In Prof   : None
TxAis            : noTransmit               TxGrace           : noTransmit
Facility Fault   : disabled
CcmLtmPriority   : 7                        CcmPaddingSize    : 0 octets
CcmTx            : 0                        CcmSequenceErr    : 0
CcmTxIfStatus    : Absent                   CcmTxPortStatus   : Absent
CcmTxRdi         : False                    CcmTxCcmStatus    : noTransmit
CcmIgnoreTLVs    : (Not Specified)
Fault Propagation: disabled                 FacilityFault     : Ignore
MA-CcmInterval   : 10                       MA-CcmHoldTime    : 0ms
MA-Primary-Vid   : Disabled
Eth-1Dm Threshold: 3(sec)                   MD-Level          : 0
Eth-1Dm Last Dest: 00:00:00:00:00:00
Eth-Dmm Last Dest: 00:00:00:00:00:00
Eth-Ais          : Disabled
Eth-Ais Tx defCCM: allDef
Eth-Tst          : Disabled
Eth-CSF          : Disabled
Eth-Cfm Grace Tx : Enabled                  Eth-Cfm Grace Rx  : Enabled
Eth-Cfm ED Tx    : Disabled                 Eth-Cfm ED Rx     : Enabled
Eth-Cfm ED Rx Max: 0
Eth-Cfm ED Tx Pri: CcmLtmPri (7)
Eth-BNM Receive  : Enabled                 Eth-BNM Rx Pacing : 5
Redundancy:
    MC-LAG State : n/a
CcmLastFailure Frame:
    None
XconCcmFailure Frame:
    None
-------------------------------------------------------------------------------
MEP Received Bandwidth Notification Message Information
-------------------------------------------------------------------------------
PortID                 : 0x0000000F
Received Period (s)    : N/A
Nominal BW (Mbps)      : 10000   Current BW (Mbps)  : 1000
Reported BW (Mbps)     : 1000    Last Reported      : 2017/12/13 20:56:57 UTC
Update Pacing Timer (s): 4.23
-------------------------------------------------------------------------------
===============================================================================
When no ETH-GNM PDU is received or ETH-BNM info has been purged by CFM.
-------------------------------------------------------------------------------
MEP Received Bandwidth Notification Message Information
-------------------------------------------------------------------------------
PortID                 : N/A
Received Period (s)    : N/A
Nominal BW (Mbps)      : N/A       Current BW (Mbps)  : N/A
Reported BW (Mbps)     : N/A       Last Reported      : N/A

Update Pacing Timer (s): N/A
-------------------------------------------------------------------------------

*A:sr7_A# show eth-cfm mep 1 domain 103 association 99  all-remote-mepids
===========================================================================
Eth-CFM Remote-Mep Table
===========================================================================
R-mepId Rx CC  Rx Rdi Port-Tlv If-Tlv peer Mac Addr     CCM status since
---------------------------------------------------------------------------
2       True   False  Up       Up     8a:d9:ff:00:00:00 02/17/2009 16:27:48
3       True   False  Up       Up     8a:da:01:01:00:02 02/17/2009 16:27:48
===========================================================================
*A:sr7_A#


*A:sr7_A# show eth-cfm mep 1 domain 103 association 99  remote-mepid 3
===========================================================================
Eth-CFM Remote-Mep Table
===========================================================================
R-mepId Rx CC  Rx Rdi Port-Tlv If-Tlv peer Mac Addr     CCM status since
---------------------------------------------------------------------------
3       True   False  Up       Up     8a:da:01:01:00:02 02/17/2009 16:27:48
===========================================================================
*A:sr7_A#


*A:7710_C# show eth-cfm mep 1 domain 103 association 99 eth-test-results
==============================================================
Eth CFM ETH-Test Result Table
==============================================================
                                 Current        Accumulate
                  FrameCount     ErrBits        ErrBits
Peer Mac Addr     ByteCount      CrcErrs        CrcErrs
--------------------------------------------------------------
22:34:56:78:9a:bc 1              0              0
                  100            0              0
32:34:56:78:9a:bc 1              0              0
                  100            0              0
42:34:56:78:9a:bc 1              0              0
                  100            0              0
52:34:56:78:9a:bc 1              0              0
                  100            0              0
62:34:56:78:9a:bc 1              0              0
                  100            0              0
72:34:56:78:9a:bc 1              0              0
                  100            0              0
82:34:56:78:9a:bc 1              0              0
                  100            0              0
92:34:56:78:9a:bc 1              0              0
                  100            0              0
```

```
c2:34:56:78:9a:bc 1              0              0
                  100            0              0
d2:34:56:78:9a:bc 1              0              0
                  100            0              0
==============================================================
*A:7710_C#


*A:7710_C# show eth-cfm mep 1 domain 103 association 99 eth-test-results remote-
peer
22:34:56:78:9a:bc
==============================================================
Eth CFM ETH-Test Result Table
==============================================================
                                 Current        Accumulate
                  FrameCount     ErrBits        ErrBits
Peer Mac Addr     ByteCount      CrcErrs        CrcErrs
--------------------------------------------------------------
22:34:56:78:9a:bc 1              0              0
                  100            0              0
==============================================================
*A:7710_C#


*A:7710_C# show eth-cfm mep 1 domain 103 association 99 one-way-delay-test
==================================================================
Eth CFM One-way Delay Test Result Table
==================================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
------------------------------------------------------------------
8a:d8:01:01:00:01     759606              2840
aa:bb:cc:dd:ee:ff     760256              760256
==================================================================
*A:7710_C#


*A:7710_C# show eth-cfm mep 1 domain 103 association 99 one-way-delay-test  remote-
peer 8a:d8:01:01:00:01
==================================================================
Eth CFM One-way Delay Test Result Table
==================================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
------------------------------------------------------------------
8a:d8:01:01:00:01     759606              2840
==================================================================
*A:7710_C#


*A:sim_B# show eth-cfm mep 2 domain 103 association 99 two-way-delay-test
==================================================================
Eth CFM Two-way Delay Test Result Table
==================================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
------------------------------------------------------------------
00:16:4d:54:49:db     10190               13710
==================================================================
*A:sim_B#
```

```
*A:sim_B# show eth-cfm mep 2 domain 103 association 99 two-way-delay-test remote-
peer
00:16:4D:54:49:DB
==================================================================
Eth CFM Two-way Delay Test Result Table
==================================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
------------------------------------------------------------------
00:16:4d:54:49:db     10190               13710
==================================================================
*A:sim_B#

domain 14 format none level 4
            association 1 format icc-based name "test000000001"
                bridge-identifier 3
                exit
                auto-mep-discovery
                ccm-interval 1
                remote-mepid 409
            exit
        exit

show eth-cfm mep 28 domain 14 association 2 all-remote-mepids
=============================================================================
Eth-CFM Remote-Mep Table
=============================================================================
R-mepId AD Rx CC RxRdi Port-Tlv If-Tlv Peer Mac Addr     CCM status since
-----------------------------------------------------------------------------
30      T  True  False Up       Up     00:00:00:00:00:30 02/03/2014 21:05:01
32         True  False Up       Up     00:00:00:00:00:32 02/03/2014 21:04:32
=============================================================================
Entries marked with a 'T' under the 'AD' column have been auto-discovered.


show eth-cfm domain 14 association 2 detail
===============================================================================
Domain 14
Md-index          : 14                     Level            : 4
                                           MHF Creation     : defMHFnone
Name Format       : none                   Next Ma Index    : 1
Name              : (Not Specified)
Creation Origin   : manual
-------------------------------------------------------------------------------
Domain 14 Associations:

Md-index          : 14                     Ma-index         : 2
Name Format       : icc-based              CCM-interval     : 1
Auto Discover     : True                   CCM-hold-time    : n/a
Name              : epipe00000005
Permission        : sendIdNone
Bridge-id         : 5                      MHF Creation     : defMHFnone
PrimaryVlan       : 0                      Num Vids         : 0
MIP LTR Priority  : 7
Total MEP Count   : 3
Remote Mep Id     : 30   (AutoDiscovered)    Remote MAC Addr   : default
Remote Mep Id     : 32                     Remote MAC Addr   : default

===============================================================================
```

mip

Syntax **mip**

Context show>eth-cfm

Description This command displays SAPs/bindings provisioned for allowing the default MIP creation.

mip-instantiation

Syntax **mip-instantiation** [**level** *level*] [{**sap** *sap-id* | **sdp** *sdp-id*}]

Context show>eth-cfm

Description This command displays the active MIPs created on the node, their related object values, and the SAP or SDP binding. The attributes include a column that indicates which MIP table was responsible and authoritative for the specific active attribute. Authorities can be the association (asn), default-domain (def), or the global read-only values (sys).

Parameters *level* — The level for which all created MIPs will be displayed

Values 0 to 7

sap-id — The SAP for which created MIPs will be displayed

sdp-id — The SDP binding for which created MIPs will be displayed

Output The following is an example of ETH CFM MIP instantiation information.

Table 25 describes the show MIP instantiation command output fields.

Sample Output

```
show eth-cfm mip-instantiation
===============================================================================
CFM SAP MIP Instantiation Information
===============================================================================
SAP                          Lvl  LA   Creation    CA   IdPerm     IdA  Pri  PA
-------------------------------------------------------------------------------
1/2/1:2000.2000              4    asn  default     asn  chassis    asn  7    asn
1/2/1:3000.3000              4    def  default     def  none       sys  7    sys
-------------------------------------------------------------------------------
No. of SAP MIPs: 2
===============================================================================

===============================================================================
CFM SAP Primary VLAN MIP Instantiation Information
===============================================================================
SAP                    VLAN  Lvl  LA   Creation    CA   IdPerm     IdA  Pri  PA
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
CFM SDP MIP Instantiation Information
===============================================================================
SDP                        Lvl  LA   Creation    CA   IdPerm     IdA  Pri  PA
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
```

Table 25 **ETH-CFM MIP Instantiation Field Descriptions**

Label	Description
VLAN	Displays the primary *vlan-id* associated with the MIP, or “none” if primary-vlan-enable is not configured
L	Displays the numerical value indicating the CFM level of the MIP
LA	Displays the level authority indicating the creation routine responsible for the level
Creation	Displays the MHF creation mode that was used to create the MIP
CA	Displays the creation authority
IdPerm	Indicates whether the SenderID TLV is being included (chassis) or not (none)
IdA	Displays the IdPermission authority
Pri	Displays the numerical value that indicates the **mip-ltr-priority**
PA	Displays the **mip-ltr-priority** authority

system-config

Syntax **system-config**

Context show>eth-cfm

Description This command shows various system level configuration parameters. These global ETH CFM commands are those which are configured directly under the **config>eth-cfm** context.

Output The following is an example of ETH CFM system configuration information.

Sample Output

```
# show eth-cfm system-config
===============================================================================
CFM System Configuration
===============================================================================
Redundancy
```

```
    MC-LAG Standby MEP Shutdown: true
    MC-LAG Hold-Timer          :   1 second(s)

Synthetic Loss Measurement
    Inactivity Timer           : 100 second(s)
===============================================================================
```

eth-ring

Syntax **eth-ring** [**status**]
eth-ring [*ring-index*] **hierarchy**
eth-ring *ring-index* [**path** {**a** | **b**}]

Context show

Description This command displays Ethernet Ring information.

Parameters **status** — Specifies to display an Ethernet Ring status summary

ring-index — Specifies an Ethernet Ring index

Values 1 to 128

hierarchy — Specifies to display Ethernet Ring hierarchical relationships

path — Specifies to show information for a specific path

2.20.2.4 PW-Port Show Commands

pw-port

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

Context show

Description This command displays FPE-based PW-port configuration information, state information and forwarding statistics.

Parameters *pw-port-id* — Specifies the PW-port ID.

Values 1 to 10239

sdp-id — Displays PW port information based on the known internal SDP ID

sdp none — Displays information about FPE-based PW-ports that are not associated with any internal SDPs

statistics — Displays forwarding statistics, such as the number or forwarded or dropped frames (Ethernet, VLANs, payload)

Output The following is an example of PW port information.

Table 26 describes the **show pw-port** command output fields.

Sample Output

```
*A:vSIM# show pw-port 1
===================================================================
PW Port Information
===================================================================
PW Port   Encap          SDP        IfIndex            VC-Id
-------------------------------------------------------------------
1         dot1q          17406      1526726657         100001
*A:vSIM# show pw-port 1 detail
===============================================================================
PW Port Information
===============================================================================
PW Port            : 1
Encap              : dot1q
SDP                : 17406
IfIndex            : 1526726657
VC-Id              : 100001
Description        : test
===============================================================================
===============================================================================
Service Destination Point (Sdp Id 17406 Pw-Port 1)
===============================================================================
SDP Binding port     : pxc-1.b
VC-Id                : 100001                Admin Status       : up
Encap                : dot1q                 Oper Status        : up
VC Type              : ether

Admin Ingress label  : 262142                Admin Egress label : 262143
Oper Flags           : (Not Specified)
Monitor Oper-Group   : (Not Specified)
*A:vSIM# show pw-port 1 statistics
===============================================================================
Service Destination Point (Sdp Id 17406 Pw-Port 1)
===============================================================================
SDP Binding port     : pxc-1.b
VC-Id                : 100001                Admin Status       : up
Encap                : dot1q                 Oper Status        : up
VC Type              : ether

Admin Ingress label  : 262142                Admin Egress label : 262143
Oper Flags           : (Not Specified)
Monitor Oper-Group   : (Not Specified)

Statistics            :
I. Fwd. Pkts.        : 0                     I. Dro. Pkts.      : 0
I. Fwd. Octs.        : 0                     I. Dro. Octs.      : 0
E. Fwd. Pkts.        : 0                     E. Fwd. Octets     : 0
*A:vSIM# show pw-port sdp 17406
===================================================================
PW Port Information
```

```
===================================================================
PW Port   Encap          SDP        IfIndex            VC-Id
-------------------------------------------------------------------
1         dot1q          17406      1526726657         100001
*A:vSIM# show pw-port sdp none
===================================================================
PW Port Information
===================================================================
PW Port   Encap          SDP        IfIndex            VC-Id
-------------------------------------------------------------------
2         dot1q                     1526726658
===========================================================
```

Table 26 **PW-Port Field Descriptions**

Label	Description
PW-Port	Displays the PW port ID.
Encap	Displays the PW port encapsulation (dot1q or qinq).
SDP	Displays the Internal SDP to which this PW port is bound.
IfIndex	Displays the Internal interface index.
VC-Id	Displays the VC-id of the internal spoke SDP that interconnects external PW to this PW port.
Description	Displays the description of this PW port.
SDP Binding Port	Displays the PXC sub-port to which this PW port is bound. This is termination side of PXC, always denoted as .b side.
VC Type	Displays the VC type of the PW port.
Admin Status	Displays the admin status of the internal SDP.
Oper Status	Displays the operational status of the internal SDP.
Admin Ingress Label	Displays the ingress VC-label associated with this PW port.
Admin Egress Label	Displays the egress VC-label associated with this PW port.
Oper Flags	Displays the operational flags on the internal SDP.
Monitor Oper-Group	Displays the operational group that is being monitored by this PW port.
I. Fwd. Pkts.	Displays the number of forwarded packets ingressing this PW port.
I. Fwd. Octs.	Displays the number of forwarded octets ingressing this PW port.
E. Fwd. Pkts.	Displays the number of forwarded packets egressing this PW port.

***Table 26* PW-Port Field Descriptions (Continued)**

Label	Description
I. Dro. Pkts.	Displays the number of dropped packets on ingress.
I. Dro. Octs.	Displays the number of dropped octets on ingress.
E. Fwd. Octets.	Displays the number of forwarded octets egressing this PW port.

2.20.2.5 NGE Show Commands

group-encryption

Syntax **group-encryption**

Context show

Description This command accesses the **show>group encryption** context.

encryption-keygroup

Syntax **encryption-keygroup** *keygroup-id* [**spi** *spi*]

Context show>grp-encryp

Description This command displays NGE information for a key group.

Parameters *keygroup-id* — Specifies the key group identifier to use for the output display.

Values 1 to 15 or *keygroup-name* (up to 64 characters)

spi — Specifies the SPI to use for the output display.

Output The following output is an example of encryption key group information, and Table 27 describes the fields.

Output Example

```
domain1>show>grp-encryp#  encryption-keygroup 2
===============================================================================
Encryption Keygroup Configuration Detail
===============================================================================
Keygroup Id        : 2
Keygroup Name      : KG1_secure
Description        : Most_secure_KG
Authentication Algo: sha256
Encryption Algo    : aes128
Active Outbound SA : 6
Activation Time    : 04/20/2015 20:07:31
-------------------------------------------------------------------------------
Security Associations
-------------------------------------------------------------------------------
Spi                : 2
Install Time       : 04/20/2015 20:08:17
Key CRC            : 0x806fb970
Spi                : 6
Install Time       : 04/20/2015 19:43:40
Key CRC            : 0xa4f2d262
-------------------------------------------------------------------------------
Encryption Keygroup Forwarded Statistics
-------------------------------------------------------------------------------
Encrypted Pkts          : 0             Encrypted Bytes         : 0
Decrypted Pkts          : 0             Decrypted Bytes         : 0
-------------------------------------------------------------------------------
Encryption Keygroup Outbound Discarded Statistics (Pkts)
-------------------------------------------------------------------------------
Total Discard           : 0             Unsupported Uplink      : 0
Enqueue Error           : 0             Other                   : 0
-------------------------------------------------------------------------------
Encryption Keygroup Inbound Discarded Statistics (Pkts)
-------------------------------------------------------------------------------
Total Discard           : 0             Invalid Spi             : 0
Authentication Failure *: 0             Control Word Mismatch   : 0
Padding Error           : 0             Enqueue Error           : 0
Other                   : 0
-------------------------------------------------------------------------------

---------------------------------------------
SDP Keygroup Association Table
---------------------------------------------
SDP ID         Direction
---------------------------------------------
61             Inbound   Outbound
---------------------------------------------
Inbound Keygroup SDP Association Count:  1
Outbound Keygroup SDP Association Count: 1

---------------------------------------------
VPRN Keygroup Association Table
---------------------------------------------
VPRN SVC ID    Direction
---------------------------------------------
12             Inbound   Outbound
---------------------------------------------
Inbound Keygroup VPRN Association Count:  1
Outbound Keygroup VPRN Association Count: 1
---------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.
domain1>show>grp-encryp#

domain1# show group-encryption encryption-keygroup 1 spi 1
```

```
===============================================================================
Encryption Keygroup Security Association Detail
===============================================================================
Keygroup Id      : 1                   SPI Id           : 1
Install Time     : 06/16/2015 11:28:49
Key CRC          : 0x36e5af55
-------------------------------------------------------------------------------
Encryption Keygroup Security Association Forwarded Statistics
-------------------------------------------------------------------------------
Encrypted Pkts          : 1662534       Encrypted Bytes         : 837917136
Decrypted Pkts          : 1662333       Decrypted Bytes         : 837815832
-------------------------------------------------------------------------------
Encryption Keygroup Security Association Outbound Discarded Statistics (Pkts)
-------------------------------------------------------------------------------
Total Discard           : 0             Enqueue Error           : 0
Other                   : 0
-------------------------------------------------------------------------------
Encryption Keygroup Security Association Inbound Discarded Statistics (Pkts)
-------------------------------------------------------------------------------
Total Discard           : 0             Authentication Failure  : 0
Control Word Mismatch   : 0             Padding Error           : 0
Enqueue Error           : 0             Other                   : 0
===============================================================================
```

Table 27 **Show Encryption Key Group Output Fields**

Label	Description
Encryption Keygroup Configuration Detail	
Keygroup Id	The key group identifier
Keygroup Name	The key group name
Description	The key group description
Authentication Algo	The authentication algorithm used for the key group
Encryption Algo	The encryption algorithm used for the key group
Active Outbound SA	The active outbound SA for the key group
Activation Time	The date and time that the key group was activated
Security Associations	
Spi	The security parameter index for the SA in the key group
Install Time	The date and time that the SA was installed in the key group
Key CRC	The CRC for the key belonging to the SA
Encryption Keygroup Forwarded Statistics	
Encrypted Pkts	The number of encrypted packets forwarded by the key group

Table 27 **Show Encryption Key Group Output Fields (Continued)**

Label	Description
Encrypted Bytes	The number of encrypted bytes forwarded by the key group
Decrypted Pkts	The number of decrypted packets forwarded by the key group
Decrypted Bytes	The number of decrypted bytes forwarded by the key group
Encryption Keygroup Outbound Discarded Statistics (Pkts)	
Total Discard	The total number of outbound packets discarded by the key group
Unsupported Uplink	The total number of outbound packets discarded by the key group due to an unsupported uplink
Enqueue Error	The total number of outbound packets discarded by the key group due to an enqueuing error
Other	The total number of outbound packets discarded by the key group due to some other reason, such as an internal configuration error (for example, a key group that points to an SA, but the SA is not valid)
Encryption Keygroup Inbound Discarded Statistics (Pkts)	
Total Discard	The total number of inbound packets discarded by the key group
Invalid Spi	The total number of inbound packets discarded by the key group due to an invalid SPI
Authentication Failure *	The total number of inbound packets discarded by the key group due to an authorization failure
Control Word Mismatch	The total number of inbound packets discarded by the key group due to a control word (CW) mismatch between the encrypted (protected) CW in the ESP payload and the CW that is not encrypted
Padding Error	The total number of inbound packets discarded by the key group due to a padding error
Enqueue Error	The total number of inbound packets discarded by the key group due to an enqueuing error
Other	The total number of inbound packets discarded by the key group due to some other reason (for example, an incoming packet length is incorrect)
SDP Keygroup Association Table	

Table 27 **Show Encryption Key Group Output Fields (Continued)**

Label	Description
SDP ID	The SDP ID
Direction	The direction in which key group authentication and encryption occurs for traffic on the SDP
Inbound Keygroup SDP Association Count	The number of SDPs configured to use inbound SA
Outbound Keygroup SDP Association Count	The number of SDPs configured to use outbound SA
VPRN Keygroup Association Table	
VPRN SVC ID	The VPRN service identifier
Direction	The direction in which key group authentication and encryption occurs for traffic on the VPRN
Inbound Keygroup VPRN Association Count	The number of VPRNs configured to use inbound SA
Outbound Keygroup VPRN Association Count	The number of VPRNs configured to use outbound SA

summary

Syntax **summary**

Context show>grp-encryp

Description This command shows NGE summary information.

Output The following output is an example of NGE summary information, and Table 28 describes the fields.

Output Example

```
domain1>show>grp-encryp# summary
============================
Group Encryption
============================
Encryption Label : 34
============================
=======================================================
Encryption Keygroup
=======================================================
Id Name         Auth Algo    Encr Algo    Active OutSA
-------------------------------------------------------
2  KG1_secure   sha256       aes128                  6
```

```
4              sha256       aes128                0
-------------------------------------------------------
No. of Encryption Keygroup: 2
=======================================================
domain1>show>grp-encryp#
```

Table 28 **Show Group Encryption Summary Output Fields**

Label	Description
Group Encryption	
Encryption Label	The unique network-wide group encryption label
Encryption Keygroup	
Id	The key group identifier value
Name	The key group name
Auth Algo	The authentication algorithm used by the key group
Encr Algo	The encryption algorithm used by the key group
Active OutSA	The active outbound SA for the key group
No. of Encryption Keygroup	The number of encryption key groups currently configured on the node

2.20.2.6 Clear Commands

group-encryption

Syntax **group-encryption**

Context clear

Description This command accesses the context to clear group encryption parameters.

encryption-keygroup

Syntax **encryption-keygroup** *keygroup-id*
encryption-keygroup *keygroup-id* **spi** *spi*

Context clear>grp-encryp

Description This command clears NGE information for a key group.

Parameters *keygroup-id* — Specifies the key group identifier.

Values 1 to 127, *keygroup-name* (up to 64 characters)

spi — Specifies the SPI ID.

Values 1 to 127

2.20.2.7 Tools Perform Commands

tools

Syntax **tools**

Context root

Description This command enables the context to enable useful tools for debugging purposes.

Default none

Parameters **dump** — Enables dump tools for the various protocols.

perform — Enables tools to perform specific tasks.

perform

Syntax perform

Context tools

Description This command enables the context to enable tools to perform specific tasks.

Default none

service

Syntax **service**

Context tools>perform

Description This command enables the context to configure tools for services.

id

Syntax **id** *service-id*

Context tools>perform>service

Description This command enables the context to configure tools for a specific service.

Parameters *service-id* — Specifies an existing service ID.

Values 1 to 2147483647

admin-lock

Syntax **admin-lock**

Context tools>perform>service>id

Description This command enters the context for applying an administrative lock for a spoke-sdp that is bound to a VLL SAP, another spoke-\ sdp or a VPLS interface for an MPLS-TP PW. Once the PW is locked it may be put into loopback mode. The command must be executed at both ends of the PW or MS-PW represented by the spoke-\ SDP. Test traffic can then be injected using a test SAP.

loopback

Syntax **loopback**

Context tools>perform>service>id

Description Tools for placing and removing SAPs and SDP bindings in data loopback. Overwrite will occur for any SAP or SDP binding when issuing a subsequent loopback command on the same SAP or SDP binding.

Interactions: Loopback functions are only applicable to Epipe, PBB Epipe, VPLS, I-VPLS and PBB core service contexts.

eth

Syntax **eth**

Context tools>perform>service>id>loopback

Description This command enables the context to configure a loopback on Ethernet SAPs or MPLS SDP bindings.

pw

Syntax **pw**

Context tools>perform>service>id>admin-lock
tools>perform>service>id>loopback

Description In the admin-lock context, this command administratively locks the specified spoke-sdp by locking the host service. The command must be executed at both ends of the PW or MS-PW represented by the spoke-SDP. Test traffic can then be injected using a test SAP.

In the loopback context, this command enters the MPLS-TP PW context for starting or stopping a loopback on a specified spoke-SDP. An administrative lock should first be applied to both ends of the PW or MS-PW represented by the spoke-SDP prior to configuring the loopback.

Interactions: Loopback functions for MPLS-TP pseudowire can be specified for either a T-PE or S-PE.

sdp

Syntax **sdp** *sdp-id:vc-id* [**test-service-id** *service-id*] **start**
sdp *sdp-id:vc-id* **stop**

Context tools>perform>service>admin-lock>pw

Description This command specifies the spoke SDP binding to which an administrative lock will be applied for the MPLS-TP pseudowire. The administrative lock can be placed on a spoke SDP that is bound to a VLL SAP, another spoke SDP or a VPLS interface. Once the pseudowire is locked it may be put into loopback mode. The command must be executed at both ends of the pseudowire or MS-PW represented by the spoke-SDP. Test traffic can then be injected using a configured test SAP on an Epipe, Apipe or Cpipe.

Parameters *sdp-id:vc-id* — Specifies the SDP-ID and VC-ID.

Values sdp-id 1 to 17407]
vc-id1 to 4294967295]

test-service-id — Keyword that specifies the ID of a test service (SAP) to which the SDP is bound.

sap

Syntax **sap** *sap-id* **start** *mode* [**mac-swap**] [**mac** *ieee-address*] [**all**]
sap *sap-id* **stop**

Context tools>perform>service>loopback>eth

Description This command places and removes the specific SAP in loopback mode for reflecting Ethernet traffic back in the direction of the received stream. This is only applicable to Ethernet-based SAPs.

Parameters *sap-id* — Specifies the SAP ID.

Values

null	*port-id* | lag-*id*	
dot1q	{*port-id* | lag-*id*}:{*qtag1* | cp-*conn-prof-id*	
qinq	{*port-id* | lag-*id*}:{*qtag1* | cp-*conn-prof-id*}.{*qtag2* | cp-*conn-prof-id*}	
	cp: keyword	
	conn-prof-id: 1..8000	
port-id	slot/mda/port [.channel]	
	eth-sat-id	esat-id/slot/port
		esat: keyword
		id: 1 to 20
	pxc-id	pxc-id.sub-port
		pxc pxc-id.sub-port
		pxc: keyword
		id: 1 to 64
		sub-port: a, b
lag-id	lag-*id*	
	lag*:* keyword	
	id: 1..800	
qtag1	0..4094	
qtag2	* | null | 0..4094	

start — Keyword that places the sap in loopback mode.

mode — Keywords that specify the location on the loopback in relation to the SAP.

Values **ingress** — Traffic arriving at the sap-ingress will be reflected back out the same SAP.

egress — Traffic arriving at the sap-egress will be reflected back into the service in the direction of the original source.

stop — Removes the SAP from loopback mode.

mac-swap — Enable source address and destination address swapping for the reflected packets when the arriving packet is unicast. Any broadcast and multicast packets arriving on a looped point will be dropped.

mac *ieee-address* — Optionally configures the source MAC address used in the reflected packet when the arriving packet is a broadcast or multicast. This does not apply to arriving unicast packets.

6-byte unicast mac-address in the form xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx.

all — Configured *ieee-address* is used as the source address for all reflected packets regardless of the arriving destination.

sdp

Syntax **sdp** *sdp-id:vc-id* **start** *mode* [**mac-swap**] [**mac** *ieee-address*] [**all**]
sdp *sdp-id:vc-id* **stop**

Context tools>perform>service>loopback>eth

Description This command places the specific MPLS SDP binding in loopback mode for reflecting Ethernet traffic back in the direction of the received stream. This is only applicable to MPLS SDP Bindings.

Parameters *sdp-id:vc-id* — Specifies the SDP ID and VC-ID.

Values sdp-id 1 to 17407
vc-id1 to 4294967295

start *mode* — Specifies the loopback in relation to the MPLS SDP Binding.

Values **ingress** — Traffic arriving at the sap-ingress will be reflected back out the same **sap**.

egress — Traffic arriving at the sap-egress will be reflected back into the service in the direction of the original source.

stop — Keyword that removes the MPLS SD-binding from loopback mode.

mac-swap — Enable source address and destination address swapping for the reflected packets when the arriving packet is unicast. Any broadcast and multicast packets arriving on a looped point will be dropped.

mac *ieee-address* — Optionally configure the source MAC address used in the reflected packet when the arriving packet is a broadcast or multicast. This does not apply to arriving unicast packets.

Values 6-byte unicast mac-address in the form
xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

all — Configured ieee-address is used as the source address for all reflected packets regardless of the arriving destination.

mac-swap — No swapping of MAC addresses are performed without specifying this option and any non-unicast destined packets will not be reflected back to the source.

sdp

Syntax **sdp** *sdp-id:vc-id* {**start | stop**}

Context tools>perform>service>loopback>pw

Description This command places or removes the specified MPLS-TP SDP binding in loopback mode for the purpose of an MPLS-TP pseudowire test service.

Note: The loopback is created at the PW level so everything under the PW label is looped back. It is recommended to configure an administrative lock for the MPLS-TP pseudowire for the specified test service prior to configuring the loopback.

Parameters *sdp-id:vc-id* — Specifies the SDP-ID and VC-ID.

Values sdp-id 1 to 17407
vc-id1 to 4294967295

start — Keyword that places the specified MPLS-TP PW in loopback mode for the purpose of an MPLS_TP PW test service.

stop — Keyword that removes the SDP binding from the loopback mode for the MPLS-TP pseudowire test service.

clear

Syntax **clear** *ring-index*

Context tools>perform>eth-ring

Description The clear command, at the Ethernet Ring Node, is used for the following operations:

- Clearing an active local administrative command, such as a Forced Switch or Manual Switch
- Triggering reversion before the WTR or WTB timer expires in case of revertive operation
- Triggering reversion in case of non-reactive operation

Parameters *ring-index* — Specifies an Ethernet Ring index.

Values 1 to 128

force

Syntax **force** *ring-index* **path** {**a** | **b**}

Context tools>perform>eth-ring

Description This command forces a block on the ring port where the command is issued.

Parameters *ring-index* — Specifies an Ethernet Ring index.

Values 1 to 128

manual

Syntax **manual** *ring-index* **path** {**a** | **b**}

Context tools>perform>eth-ring

Description This command forces a block on the ring port where the command is issued, in the absence of a failure or FS.

Parameters *ring-index* — Specifies an Ethernet Ring index.

Values 1 to 128

2.20.2.8 Tools Dump Commands

dump

Syntax **dump**

Context tools

Description This command enables the context to display output for tools-related tasks.

service

Syntax **service**

Context tools>dump

Description This command enables the context to display service dump information.

loopback

Syntax **loopback**

Context tools>dump>service

Description This command displays all configured Ethernet loopbacks.

id

Syntax **id** *service-id*

Context tools>dump>service

Description This command enables the context to display information for a specific service.

Parameters *service-id* — Specifies the service ID.

Values 1 to 2148007980 | *svc-name*: 64 characters max.

loopback

Syntax **loopback sap** *sap-id*
loopback sdp *sdp-id***:***vc-id*

Context tools>dump>service>id

Description This command displays configured service-specific Ethernet loopbacks.

Parameters *sap-id* — Specifies the SAP ID.

Values

null	*port-id* | lag-*id*	
dot1q	{*port-id* | lag-*id*}:{*qtag1* | cp-*conn-prof-id*	
qinq	{*port-id* | lag-*id*}:{*qtag1* | cp-*conn-prof-id*}.{*qtag2* | cp-*conn-prof-id*} cp: keyword *conn-prof-id*: 1..8000	
port-id	slot/mda/port [.channel]	
	eth-sat-id	esat-id/slot/port esat: keyword id: 1 to 20
	pxc-id	pxc-id.sub-port pxc pxc-id.sub-port pxc: keyword id: 1 to 64 sub-port: a, b
lag-id	lag-*id* lag*:* keyword *id*: 1..800	
qtag1	0..4094	
qtag2	* | null | 0..4094	

sdp-id:vc-id — Specifies the SDP ID and VC-ID.

Values *sdp-id*: 1 to 17407
vc-id: 1 to 4294967295

eth-ring

Syntax **eth-ring** *ring-index* [**clear**]

Context tools>dump

Description This command displays Ethernet Ring information.

Parameters *ring-index* — Specifies an Ethernet Ring index.

Values 1 to 128

clear — Keyword to clear stored information for the specified Ethernet Ring.

3 Common CLI Command Descriptions

3.1 In This Chapter

This chapter provides information about common Command Line Interface (CLI) syntax and command usage.

3.1.1 Common Service Commands

The section describes the common Service CLI command syntax.

3.1.1.1 SAP Commands

sap

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

Context config

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:

Table 29 **sap-id Formats**

Type	Syntax	Example
port-id	*slot*/*mda*/*port*[.*channel*]	1/1/5
null	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id* | *aps-id*]	*port-id*: 1/1/3 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*: lag-3 *aps-id*: aps-1

Table 29 **sap-id Formats (Continued)**

Type	Syntax	Example
dot1q	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id* | *aps-id*]:qtag1	*port-id*:qtag1: 1/1/3:100 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*:qtag1:lag-3:102 *aps-id*:qtag1: aps-1:27
qinq	[*port-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	*port-id*:qtag1.qtag2: 1/1/3:100.10 *bpgrp-id*: bpgrp-ima-1 *lag-id*:qtag1.qtag2: lag-10:
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:vpi/vci |vpi |vpi1.vpi2] [port-id | aps-id [:vpi/vci |vpi | vpi1.vpi2 | cp.conn-prof-id]	*port-id:* 1/1/1 *aps-id:* aps-1 *vpi/vci:* 16/26 *vpi:* 16 *vpi1.vpi2*: 16.200 *cp.conn-prof-id:* 1/2/1:cp.2
frame-relay	[*port-id* | *aps-id*]:*dlci*	*port-id*: 1/1/1:100 *bundle-id*: bundle-fr-3/1.1:100 *aps-id*: aps-1 *dlci*: 16
cisco-hdlc	slot/mda/port.channel	*port-id*: 1/1/3.1

The following values apply to the 7750 SR:

Values

sap-id	null	{port-id | bundle-id | bpgrp-id | lag-id | aps-id}	
	dot1q	{port-id | bundle-id | bpgrp-id | lag-id | aps-id | pw-id}:qtag1	
	qinq	{port-id | bundle-id | bpgrp-id | lag-id | pw-id}:qtag1.qtag2	
	atm	{port-id | aps-id}[:{vpi/vci | vpi | vpi1.vpi2 | cp.conn-prof-id}]	
		cp	keyword
		conn-prof-id	1 to 8000
	frame	port-id | aps-id:dlci	
	cisco-hdlc	slot/mda/port.channel	
	cem	slot/mda/port.channel	
	ima-grp	bundle-id[:{vpi/vci | vpi | vpi1.vpi2 | cp.conn-prof-id}]	
		cp	keyword

	conn-prof-id	1 to 8000
port-id	slot/mda/port[.channel]	
bundle-id	bundle-type-slot/ mda.*bundle-num*	
	bundle	keyword
	type	ima, fr, ppp
	bundle-num	1 to 336
bpgrp-id	bpgrp-type-bpgrp-num	
	bpgrp	keyword
	type	ima, ppp
	bpgrp-num	1 to 2000
aps-id	aps-group-id[.channel]	
	aps	keyword
	group-id	1 to 64
ccag-id	*ccag-id.path-id*[*cc-type*] *cc-id*	
	ccag	keyword
	id	1 to 8
	path-id	a, b
	cc-type	.sap-net, net-sap
	cc-id	0 to 4094
eth-tunnel	eth-tunnel-*id*[:eth-tun-sap-id]	
	id	1 to 1024
	eth-tun-sap-id	0 to 4094
lag-id	lag-id	
	lag	keyword
	id	1 to 800
pw-id	pw-*id*	
	pw	keyword
	id	1 to 10239
qtag1	*, 0 to 4094	
qtag2	*, 0 to 4094	
sap-id	pw-id:qtag1[.qtag2]	
	pw-	keyword
	id	identifier for the pw-port [1 to 10239]
	qtag1	value of the first 802.1 qtag
	qtag2	value of the second 802.1 qtag
vpi	0 to 4095 (NNI)	

	0 to 255 (UNI)	
vci	1, 2, 5 to 65535	
dlci	16 to 1022	
tunnel-id	tunnel-*id*.private|public:*tag*	
	tunnel	keyword
	id	1 to 16
	tag	0 to 4094

The following values apply to the 7450 ESS:

Values

sap-id	null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
	dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
	qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
	atm	[*port-id* | *aps-id*][:*vpi/vci*|*vpi*| *vpi1.vpi2*]	
	frame	[*port-id* | *aps-id*]:*dlci*	
	cisco-hdlc	*slot/mda/port.channel*	
	ima-grp	[*bundle-id*[:vpi/vci|vpi|*vpi1.vpi2*]	
	port-id	*slot/mda/port*[.*channel*]	
	bundle-id	bundle-*type-slot/mda.bundle-num*	
		bundle	keyword
		type	ima, fr, ppp
		bundle-num	1 to 336
	bpgrp-id	bpgrp-*type-bpgrp-num*	
		bpgrp	keyword
		type	ima, ppp
		bpgrp-num	1 to 2000
	aps-id	aps-*group-id*[.*channel*]	
		aps	keyword
		group-id	1 to 64
	ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
		ccag	keyword
		id	1 to 8
		path-id	a, b
		cc-type	.sap-net, .net-sap
		cc-id	0 to 4094
	eth-tunnel	eth-tunnel-*id*[:*eth-tun-sap-id*]	
		id	1 to 1024

	eth-tun-sap-id	0 to 4094
lag-id	lag-id	
	lag	keyword
	id	1 to 800
qtag1	0 to 4094	
qtag2	*, 0 to 4094	
sap-id	pw-<id>:<qtag1>[.<qtag2>]	
	pw	keyword
	id	identifier for the pw-port [1 to 10239]
	qtag1	value of the first 802.1 qtag
	qtag2	value of the second 802.1 qtag
vpi	NNI: 0 to 4095	
	UNI: 0 to 255	
vci	1, 2, 5 to 65535	
dlci	16 to 1022	

bundle-id — Specifies the multilink bundle to be associated with this IP interface. This parameter applies to the 7450 ESS and 7750 SR. The **bundle** keyword must be entered at the beginning of the parameter.

The command syntax must be configured as follows:

bundle-id: *bundle-type-slot-id/mda-slot.bundle-num*

bundle-id value range: 1 to 336

For example:

```
*A:ALA-12>config# port bundle-ppp-5/1.1
*A:ALA-12>config>port# multilink-bundle
```

bgprp-id — Specifies the bundle protection group ID to be associated with this IP interface. This parameter applies to the 7450 ESS and 7750 SR. The **bpgrp** keyword must be entered at the beginning of the parameter.

The command syntax must be configured as follows:

bpgrp-id: bpgrp-type-bpgrp-num

type: ima

bpgrp-num value range: 1 to 2000

For example:

```
*A:ALA-12>config# port bpgrp-ima-1
*A:ALA-12>config>service>vpls$ sap bpgrp-ima-1 create
```

qtag1, qtag2 — Specifies the encapsulation value used to identify the SAP on the port or sub-port. This parameter must be specifically defined.

Values qtag1: *, 0 to 4094
qtag2: *, null, 0 to 4094

The values depend on the encapsulation type configured for the interface. Table 30 describes the allowed values for the port and encapsulation types.

Table 30 **Permitted Values for Port Type and Encap Type**

Port Type	Encap-Type	Allowed Values	Comments
Ethernet	Null	0	The SAP is identified by the port.
Ethernet	Dot1q	0 to 4094	The SAP is identified by the 802.1Q tag on the port. Note that a 0 qtag1 value also accepts untagged packets on the dot1q port.
Ethernet	QinQ	qtag1: *, 0 to 4094 qtag2: *, null, 0 to 4094	The SAP is identified by two 802.1Q tags on the port. Note that The following combinations of qtag1.qtag2 accept untagged packets: “0.*”, “*.null”, “*.*”.
SONET/SDH	IPCP	-	The SAP is identified by the channel. No BCP is deployed and all traffic is IP.
SONET/SDH TDM	BCP-Null	0	The SAP is identified with a single service on the channel. Tags are assumed to be part of the customer packet and not a service delimiter.
SONET/SDH TDM	BCP-Dot1q	0 to 4094	The SAP is identified by the 802.1Q tag on the channel.
SONET/SDH TDM	Frame Relay	16 to 991	The SAP is identified by the data link connection identifier (DLCI). This port type applies to the 7750 SR only.
SONET/SDH ATM	ATM	vpi (NNI) 0 to 4095 vpi (UNI) 0 to 255 vci 1, 2, 5 to 65535	The SAP is identified by port or by PVPC or PVCC identifier (vpi, vpi/vci, or vpi range). This port type applies to the 7750 SR only.

sap ipsec-*id*.**private|public**:*tag* — This parameter associates an IPsec group SAP with this interface. This parameter applies to the 7750 SR only. 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.

Values 1 to 4095

pw-id — Specifies the SAP identifier for PW SAPs. This parameter applies to the 7450 ESS and 7750 SR.

4 Standards and Protocol Support

Note: The information presented is subject to change without notice.

Nokia assumes no responsibility for inaccuracies contained herein.

Access Node Control Protocol (ANCP)

draft-ietf-ancp-protocol-02, *Protocol for Access Node Control Mechanism in Broadband Networks*

RFC 5851, *Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks*

Application Assurance (AA)

3GPP Release 12 (ADC rules over Gx interfaces)

RFC 3507, *Internet Content Adaptation Protocol (ICAP)*

Asynchronous Transfer Mode (ATM)

AF-ILMI-0065.000, *Integrated Local Management Interface (ILMI) Version 4.0*

AF-PHY-0086.001, *Inverse Multiplexing for ATM (IMA) Specification Version 1.1*

AF-TM-0121.000, *Traffic Management Specification Version 4.1*

AF-TM-0150.00, *Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR*

GR-1113-CORE, *Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1*

GR-1248-CORE, *Generic Requirements for Operations of ATM Network Elements (NEs), Issue 3*

ITU-T I.432.1, *B-ISDN user-network interface - Physical layer specification: General characteristics (02/99)*

ITU-T I.610, *B-ISDN operation and maintenance principles and functions (11/95)*

RFC 1626, *Default IP MTU for use over ATM AAL5*

RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*

Bidirectional Forwarding Detection (BFD)

RFC 5880, *Bidirectional Forwarding Detection (BFD)*

RFC 5881, *Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop)*

RFC 5883, *Bidirectional Forwarding Detection (BFD) for Multihop Paths*

RFC 7130, *Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces*

Border Gateway Protocol (BGP)

draft-hares-idr-update-attrib-low-bits-fix-01, *Update Attribute Flag Low Bits Clarification*

draft-ietf-idr-add-paths-guidelines-08, *Best Practices for Advertisement of Multiple Paths in IBGP*

draft-ietf-idr-best-external-03, *Advertisement of the best external route in BGP*

draft-ietf-idr-bgp-flowspec-oid-03, *Revised Validation Procedure for BGP Flow Specifications*

draft-ietf-idr-bgp-gr-notification-01, *Notification Message support for BGP Graceful Restart*

draft-ietf-idr-bgp-optimal-route-reflection-10, *BGP Optimal Route Reflection (BGP-ORR)*

draft-ietf-idr-error-handling-03, *Revised Error Handling for BGP UPDATE Messages*

draft-ietf-idr-flowspec-interfaceset-03, *Applying BGP flowspec rules on a specific interface set*

draft-ietf-idr-flowspec-path-redirect-05, *Flowspec Indirection-id Redirect* (localised ID)

draft-ietf-idr-flowspec-redirect-ip-02, *BGP Flow-Spec Redirect to IP Action*

draft-ietf-idr-link-bandwidth-03, *BGP Link Bandwidth Extended Community*

draft-ietf-sidr-origin-validation-signaling-04, *BGP Prefix Origin Validation State Extended Community*

draft-uttaro-idr-bgp-persistence-03, *Support for Long-lived BGP Graceful Restart*

RFC 1772, *Application of the Border Gateway Protocol in the Internet*

RFC 1997, *BGP Communities Attribute*

RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*

RFC 2439, *BGP Route Flap Damping*

RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*

RFC 2858, *Multiprotocol Extensions for BGP-4*

RFC 2918, *Route Refresh Capability for BGP-4*

RFC 3107, *Carrying Label Information in BGP-4*

RFC 3392, *Capabilities Advertisement with BGP-4*

RFC 4271, *A Border Gateway Protocol 4 (BGP-4)*

RFC 4360, *BGP Extended Communities Attribute*

RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 4456, *BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)*

RFC 4486, *Subcodes for BGP Cease Notification Message*

RFC 4659, *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*

RFC 4684, *Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)*

RFC 4724, *Graceful Restart Mechanism for BGP* (helper mode)

RFC 4760, *Multiprotocol Extensions for BGP-4*

RFC 4798, *Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)*

RFC 4893, *BGP Support for Four-octet AS Number Space*

RFC 5004, *Avoid BGP Best Path Transitions from One External to Another*

RFC 5065, *Autonomous System Confederations for BGP*

RFC 5291, *Outbound Route Filtering Capability for BGP-4*

RFC 5396, *Textual Representation of Autonomous System (AS) Numbers* (asplain)

RFC 5549, *Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop*

RFC 5575, *Dissemination of Flow Specification Rules*

RFC 5668, *4-Octet AS Specific BGP Extended Community*

RFC 6286, *Autonomous-System-Wide Unique BGP Identifier for BGP-4*

RFC 6810, *The Resource Public Key Infrastructure (RPKI) to Router Protocol*

RFC 6811, *Prefix Origin Validation*

RFC 6996, *Autonomous System (AS) Reservation for Private Use*

RFC 7311, *The Accumulated IGP Metric Attribute for BGP*

RFC 7607, *Codification of AS 0 Processing*

RFC 7674, *Clarification of the Flowspec Redirect Extended Community*

RFC 7752, *North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP*

RFC 7854, *BGP Monitoring Protocol (BMP)*

RFC 7911, *Advertisement of Multiple Paths in BGP*

RFC 7999, *BLACKHOLE Community*

RFC 8092, *BGP Large Communities Attribute*

Circuit Emulation

RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)*

RFC 5086, *Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)*

RFC 5287, *Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks*

Ethernet

IEEE 802.1AB, *Station and Media Access Control Connectivity Discovery*

IEEE 802.1ad, *Provider Bridges*

IEEE 802.1ag, *Connectivity Fault Management*

IEEE 802.1ah, *Provider Backbone Bridges*

IEEE 802.1ak, *Multiple Registration Protocol*

IEEE 802.1aq, *Shortest Path Bridging*

IEEE 802.1ax, *Link Aggregation*

IEEE 802.1D, *MAC Bridges*

IEEE 802.1p, *Traffic Class Expediting*

IEEE 802.1Q, *Virtual LANs*

IEEE 802.1s, *Multiple Spanning Trees*

IEEE 802.1w, *Rapid Reconfiguration of Spanning Tree*

IEEE 802.1X, *Port Based Network Access Control*

IEEE 802.3ab, *1000BASE-T*

IEEE 802.3ac, *VLAN Tag*

IEEE 802.3ad, *Link Aggregation*

IEEE 802.3ae, *10 Gb/s Ethernet*

IEEE 802.3ah, *Ethernet in the First Mile*

IEEE 802.3ba, *40 Gb/s and 100 Gb/s Ethernet*

IEEE 802.3i, *Ethernet*

IEEE 802.3u, *Fast Ethernet*

IEEE 802.3x, *Ethernet Flow Control*

IEEE 802.3z, *Gigabit Ethernet*

ITU-T G.8031/Y.1342, *Ethernet Linear Protection Switching*

ITU-T G.8032/Y.1344, *Ethernet Ring Protection Switching*

ITU-T Y.1731, *OAM functions and mechanisms for Ethernet based networks*

Ethernet VPN (EVPN)

draft-ietf-bess-evpn-ac-df-01, *AC-Influenced Designated Forwarder Election for EVPN*

draft-ietf-bess-evpn-pref-df-01, *Preference-based EVPN DF Election*

draft-ietf-bess-evpn-prefix-advertisement-11, *IP Prefix Advertisement in EVPN* draft-ietf-bess-evpn-proxy-arp-nd-04, *Operational Aspects of Proxy-ARP/ND in EVPN Networks*

draft-ietf-bess-evpn-vpls-seamless-integ-03, *(PBB-)EVPN Seamless Integration with (PBB-)VPLS*

draft-snr-bess-pbb-evpn-isid-cmacflush-01, *PBB-EVPN ISID-based CMAC-Flush*

RFC 7432, *BGP MPLS-Based Ethernet VPN*

RFC 7623, *Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)*

RFC 8214, *Virtual Private Wire Service Support in Ethernet VPN*

RFC 8317, *Ethernet-Tree (E-Tree) Support in Ethernet VPN (EVPN) an Provider Backbone Bridging EVPN (PBB-EVPN)*

RFC 8365, *A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)*

Frame Relay

ANSI T1.617 Annex D, *DSS1 - Signalling Specification For Frame Relay Bearer Service*

FRF.1.2, *PVC User-to-Network Interface (UNI) Implementation Agreement*

FRF.12, *Frame Relay Fragmentation Implementation Agreement*

FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*

FRF.5, *Frame Relay/ATM PVC Network Interworking Implementation*

FRF2.2, *PVC Network-to-Network Interface (NNI) Implementation Agreement*

ITU-T Q.933 Annex A, *Additional procedures for Permanent Virtual Connection (PVC) status management*

Generalized Multiprotocol Label Switching (GMPLS)

draft-ietf-ccamp-rsvp-te-srlg-collect-04, *RSVP-TE Extensions for Collecting SRLG Information*

RFC 3471, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description*

RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions*

RFC 4204, *Link Management Protocol (LMP)*

RFC 4208, *Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model*

RFC 4872, *RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery*

RFC 5063, *Extensions to GMPLS Resource Reservation Protocol (RSVP) Graceful Restart* (helper mode)

gRPC Remote Procedure Calls (gRPC)

gnmi.proto, *gRPC Network Management Interface (gNMI), version 0.4.0*

gRPC Network Management Interface (gNMI), *Capabilities, Get, Set, Subscribe (ONCE, SAMPLE, ON_CHANGE)*

Intermediate System to Intermediate System (IS-IS)

draft-ietf-isis-mi-02, *IS-IS Multi-Instance*

draft-kaplan-isis-ext-eth-02, *Extended Ethernet Frame Size Support*

ISO/IEC 10589:2002, Second Edition, Nov. 2002, *Intermediate system to Intermediate system intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)*

RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*

RFC 2973, *IS-IS Mesh Groups*

RFC 3359, *Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System*

RFC 3719, *Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)*

RFC 3787, *Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS)*

RFC 4971, *Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information*

RFC 5120, *M-ISIS: Multi Topology (MT) Routing in IS-IS*

RFC 5130, *A Policy Control Mechanism in IS-IS Using Administrative Tags*

RFC 5301, *Dynamic Hostname Exchange Mechanism for IS-IS*

RFC 5302, *Domain-wide Prefix Distribution with Two-Level IS-IS*

RFC 5303, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies*

RFC 5304, *IS-IS Cryptographic Authentication*

RFC 5305, *IS-IS Extensions for Traffic Engineering TE*

RFC 5306, *Restart Signaling for IS-IS* (helper mode)

RFC 5307, *IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*

RFC 5308, *Routing IPv6 with IS-IS*

RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*

RFC 5310, *IS-IS Generic Cryptographic Authentication*

RFC 6213, *IS-IS BFD-Enabled TLV*

RFC 6232, *Purge Originator Identification TLV for IS-IS*

RFC 6233, *IS-IS Registry Extension for Purges*

RFC 6329, *IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging*

RFC 7775, *IS-IS Route Preference for Extended IP and IPv6 Reachability*

RFC 7794, *IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability*

RFC 8202, *IS-IS Multi-Instance* (single topology)

Internet Protocol (IP) — Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, *Operational management of Loop Free Alternates*

RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*

RFC 7431, *Multicast-Only Fast Reroute*

RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*

Internet Protocol (IP) — General

draft-grant-tacacs-02, *The TACACS+ Protocol*

RFC 768, *User Datagram Protocol*

RFC 793, *Transmission Control Protocol*

RFC 854, *Telnet Protocol Specifications*

RFC 1350, *The TFTP Protocol (revision 2)*

RFC 2347, *TFTP Option Extension*

RFC 2348, *TFTP Blocksize Option*

RFC 2349, *TFTP Timeout Interval and Transfer Size Options*

RFC 2428, *FTP Extensions for IPv6 and NATs*

RFC 2784, *Generic Routing Encapsulation (GRE)*

RFC 2890, *Key and Sequence Number Extensions to GRE*

RFC 4250, *The Secure Shell (SSH) Protocol Assigned Numbers*

RFC 4251, *The Secure Shell (SSH) Protocol Architecture*

RFC 4252, *The Secure Shell (SSH) Authentication Protocol* (publickey, password)

RFC 4253, *The Secure Shell (SSH) Transport Layer Protocol*

RFC 4254, *The Secure Shell (SSH) Connection Protocol*

RFC 4632, *Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*

RFC 5082, *The Generalized TTL Security Mechanism (GTSM)*

RFC 5656, *Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer* (ECDSA)

RFC 5925, *The TCP Authentication Option*

RFC 5926, *Cryptographic Algorithms for the TCP Authentication Option (TCP-AO)*

RFC 6398, *IP Router Alert Considerations and Usage* (MLD)

RFC 6528, *Defending against Sequence Number Attacks*

Internet Protocol (IP) — Multicast

cisco-ipmulticast/pim-autorp-spec01, *Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast* (version 1)

draft-dolganow-bess-mvpn-expl-track-01, *Explicit Tracking with Wild Card Routes in Multicast VPN*

draft-ietf-bier-mvpn-11, *Multicast VPN Using BIER*

draft-ietf-idmr-traceroute-ipm-07, *A "traceroute" facility for IP Multicast*

draft-ietf-l2vpn-vpls-pim-snooping-07, *Protocol Independent Multicast (PIM) over Virtual Private LAN Service (VPLS)*

draft-ietf-mboned-mtrace-v2-17, *Mtrace Version 2: Traceroute Facility for IP Multicast*

RFC 1112, *Host Extensions for IP Multicasting*

RFC 2236, *Internet Group Management Protocol, Version 2*

RFC 2365, *Administratively Scoped IP Multicast*

RFC 2375, *IPv6 Multicast Address Assignments*

RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*

RFC 3306, *Unicast-Prefix-based IPv6 Multicast Addresses*

RFC 3376, *Internet Group Management Protocol, Version 3*

RFC 3446, *Anycast Rendevous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*

RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*

RFC 3618, *Multicast Source Discovery Protocol (MSDP)*

RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*

RFC 3956, *Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address*

RFC 3973, *Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)* (auto-RP groups)

RFC 4541, *Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches*

RFC 4604, *Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast*

RFC 4607, *Source-Specific Multicast for IP*

RFC 4608, *Source-Specific Protocol Independent Multicast in 232/8*

RFC 4610, *Anycast-RP Using Protocol Independent Multicast (PIM)*

RFC 4611, *Multicast Source Discovery Protocol (MSDP) Deployment Scenarios*

RFC 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*

RFC 5186, *Internet Group Management Protocol Version 3 (IGMPv3) / Multicast Listener Discovery Version 2 (MLDv2) and Multicast Routing Protocol Interaction*

RFC 5384, *The Protocol Independent Multicast (PIM) Join Attribute Format*

RFC 5496, *The Reverse Path Forwarding (RPF) Vector TLV*

RFC 6037, *Cisco Systems' Solution for Multicast in MPLS/BGP IP VPNs*

RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*

RFC 6513, *Multicast in MPLS/BGP IP VPNs*

RFC 6514, *BGP Encodings and Procedures for Multicast in MPLS/IP VPNs*

RFC 6515, *IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs*

RFC 6516, *IPv6 Multicast VPN (MVPN) Support Using PIM Control Plane and Selective Provider Multicast Service Interface (S-PMSI) Join Messages*

RFC 6625, *Wildcards in Multicast VPN Auto-Discover Routes*

RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path*

RFC 7246, *Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context*

RFC 7385, *IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points*

RFC 7716, *Global Table Multicast with BGP Multicast VPN (BGP-MVPN) Procedures*

RFC 7761, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*

RFC 8279, *Multicast Using Bit Index Explicit Replication (BIER)*

RFC 8296, *Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks* (MPLS encapsulation)

RFC 8401, *Bit Index Explicit Replication (BIER) Support via IS-IS*

Internet Protocol (IP) — Version 4

RFC 791, *Internet Protocol*

RFC 792, *Internet Control Message Protocol*

RFC 826, *An Ethernet Address Resolution Protocol*

RFC 951, *Bootstrap Protocol (BOOTP)*

RFC 1034, *Domain Names - Concepts and Facilities*

RFC 1035, *Domain Names - Implementation and Specification*

RFC 1191, *Path MTU Discovery* (router specification)

RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*

RFC 1534, *Interoperation between DHCP and BOOTP*

RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*

RFC 1812, *Requirements for IPv4 Routers*

RFC 1918, *Address Allocation for Private Internets*

RFC 2003, *IP Encapsulation within IP*

RFC 2131, *Dynamic Host Configuration Protocol*

RFC 2132, *DHCP Options and BOOTP Vendor Extensions*

RFC 2401, *Security Architecture for Internet Protocol*

RFC 3021, *Using 31-Bit Prefixes on IPv4 Point-to-Point Links*

RFC 3046, *DHCP Relay Agent Information Option (Option 82)*

RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*

RFC 4884, *Extended ICMP to Support Multi-Part Messages* (ICMPv4 and ICMPv6 Time Exceeded)

Internet Protocol (IP) — Version 6

RFC 2464, *Transmission of IPv6 Packets over Ethernet Networks*

RFC 2529, *Transmission of IPv6 over IPv4 Domains without Explicit Tunnels*

RFC 3122, *Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification*

RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*

RFC 3587, *IPv6 Global Unicast Address Format*

RFC 3596, *DNS Extensions to Support IP version 6*

RFC 3633, *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*

RFC 3646, *DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*

RFC 3736, *Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6*

RFC 3971, *SEcure Neighbor Discovery (SEND)*

RFC 3972, *Cryptographically Generated Addresses (CGA)*

RFC 4007, *IPv6 Scoped Address Architecture*

RFC 4193, *Unique Local IPv6 Unicast Addresses*

RFC 4291, *Internet Protocol Version 6 (IPv6) Addressing Architecture*

RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*

RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*

RFC 4862, *IPv6 Stateless Address Autoconfiguration* (router functions)

RFC 4890, *Recommendations for Filtering ICMPv6 Messages in Firewalls*

RFC 4941, *Privacy Extensions for Stateless Address Autoconfiguration in IPv6*

RFC 5007, *DHCPv6 Leasequery*

RFC 5095, *Deprecation of Type 0 Routing Headers in IPv6*

RFC 5722, *Handling of Overlapping IPv6 Fragments*

RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6* (IPv6)

RFC 5952, *A Recommendation for IPv6 Address Text Representation*

RFC 6092, *Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service* (Internet Control and Management, Upper-Layer Transport Protocols, UDP Filters, IPsec and Internet Key Exchange (IKE), TCP Filters)

RFC 6106, *IPv6 Router Advertisement Options for DNS Configuration*

RFC 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*

RFC 8021, *Generation of IPv6 Atomic Fragments Considered Harmful*

RFC 8200, *Internet Protocol, Version 6 (IPv6) Specification*

RFC 8201, *Path MTU Discovery for IP version 6*

Internet Protocol Security (IPsec)

draft-ietf-ipsec-isakmp-mode-cfg-05, *The ISAKMP Configuration Method*

draft-ietf-ipsec-isakmp-xauth-06, *Extended Authentication within ISAKMP/Oakley (XAUTH)*

RFC 2401, *Security Architecture for the Internet Protocol*

RFC 2403, *The Use of HMAC-MD5-96 within ESP and AH*

RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH*

RFC 2405, *The ESP DES-CBC Cipher Algorithm With Explicit IV*

RFC 2406, *IP Encapsulating Security Payload (ESP)*

RFC 2407, *IPsec Domain of Interpretation for ISAKMP (IPsec DoI)*

RFC 2408, *Internet Security Association and Key Management Protocol (ISAKMP)*

RFC 2409, *The Internet Key Exchange (IKE)*

RFC 2410, *The NULL Encryption Algorithm and Its Use With IPsec*

RFC 3526, *More Modular Exponential (MODP) Diffie-Hellman group for Internet Key Exchange (IKE)*

RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*

RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*

RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*

RFC 3947, *Negotiation of NAT-Traversal in the IKE*

RFC 3948, *UDP Encapsulation of IPsec ESP Packets*

RFC 4106, *The Use of Galois/Counter Mode (GCM) in IPsec ESP*

RFC 4210, *Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)*

RFC 4211, *Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)*

RFC 4301, *Security Architecture for the Internet Protocol*

RFC 4303, *IP Encapsulating Security Payload*

RFC 4307, *Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)*

RFC 4308, *Cryptographic Suites for IPsec*

RFC 4434, *The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE)*

RFC 4543, *The Use of Galois Message Authentication Code (GMAC) in IPsec ESP and AH*

RFC 4868, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPSec*

RFC 4945, *The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2 and PKIX*

RFC 5019, *The Lightweight Online Certificate Status Protocol (OCSP) Profile for High-Volume Environments*

RFC 5280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*

RFC 5282, *Using Authenticated Encryption Algorithms with the Encrypted Payload of the IKEv2 Protocol*

RFC 5903, *ECP Groups for IKE and IKEv2*

RFC 5998, *An Extension for EAP-Only Authentication in IKEv2*

RFC 6379, *Suite B Cryptographic Suites for IPsec*

RFC 6380, *Suite B Profile for Internet Protocol Security (IPsec)*

RFC 6712, *Internet X.509 Public Key Infrastructure -- HTTP Transfer for the Certificate Management Protocol (CMP)*

RFC 6960, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*

RFC 7296, *Internet Key Exchange Protocol Version 2 (IKEv2)*

RFC 7321, *Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)*

RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*

RFC 7427, *Signature Authentication in the Internet Key Exchange Version 2 (IKEv2)*

RFC 7468, *Textual Encodings of PKIX, PKCS, and CMS Structures*

Label Distribution Protocol (LDP)

draft-ietf-mpls-ldp-ip-pw-capability-09, *Controlling State Advertisements Of Non-negotiated LDP Applications*

draft-pdutta-mpls-ldp-adj-capability-00, *LDP Adjacency Capabilities*

draft-pdutta-mpls-ldp-v2-00, *LDP Version 2*

draft-pdutta-mpls-mldp-up-redundancy-00, *Upstream LSR Redundancy for Multi-point LDP Tunnels*

draft-pdutta-mpls-multi-ldp-instance-00, *Multiple LDP Instances*

draft-pdutta-mpls-tldp-hello-reduce-04, *Targeted LDP Hello Reduction*

RFC 3037, *LDP Applicability*

RFC 3478, *Graceful Restart Mechanism for Label Distribution Protocol* (helper mode)

RFC 5036, *LDP Specification*

RFC 5283, *LDP Extension for Inter-Area Label Switched Paths (LSPs)*

RFC 5443, *LDP IGP Synchronization*

RFC 5561, *LDP Capabilities*

RFC 5919, *Signaling LDP Label Advertisement Completion*

RFC 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*

RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*

RFC 6826, *Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*

RFC 7032, *LDP Downstream-on-Demand in Seamless MPLS*

RFC 7552, *Updates to LDP for IPv6*

Layer Two Tunneling Protocol (L2TP) Network Server (LNS)

draft-mammoliti-l2tp-accessline-avp-04, *Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions*

RFC 2661, *Layer Two Tunneling Protocol "L2TP"*

RFC 2809, *Implementation of L2TP Compulsory Tunneling via RADIUS*

RFC 3438, *Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers: Internet Assigned Numbers Authority (IANA) Considerations Update*

RFC 3931, *Layer Two Tunneling Protocol - Version 3 (L2TPv3)*

RFC 4719, *Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)*

RFC 4951, *Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"*

Management

draft-ieft-snmpv3-update-mib-05, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*

draft-ietf-isis-wg-mib-06, *Management Information Base for Intermediate System to Intermediate System (IS-IS)*

draft-ietf-mboned-msdp-mib-01, *Multicast Source Discovery protocol MIB*

draft-ietf-mpls-ldp-mib-07, *Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP)*

draft-ietf-mpls-lsr-mib-06, *Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2*

draft-ietf-mpls-te-mib-04, *Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base*

draft-ietf-ospf-mib-update-08, *OSPF Version 2 Management Information Base*

draft-ietf-vrrp-unified-mib-06, *Definitions of Managed Objects for the VRRP over IPv4 and IPv6* (IPv6)

ianaaddressfamilynumbers-mib, *IANA-ADDRESS-FAMILY-NUMBERS-MIB*

ianagmplstc-mib, *IANA-GMPLS-TC-MIB*

ianaiftype-mib, *IANAifType-MIB*

ianaiprouteprotocol-mib, *IANA-RTPROTO-MIB*

IEEE8021-CFM-MIB, *IEEE P802.1ag(TM) CFM MIB*

IEEE8021-PAE-MIB, *IEEE 802.1X MIB*

IEEE8023-LAG-MIB, *IEEE 802.3ad MIB*

LLDP-MIB, *IEEE P802.1AB(TM) LLDP MIB*

openconfig-bgp.yang version 3.0.1, *BGP Module*

openconfig-bgp-common.yang version 3.0.1, *BGP Common Module*

openconfig-bgp-common-multiprotocol.yang version 3.0.1, *BGP Common Multiprotocol Module*

openconfig-bgp-common-structure.yang version 3.0.1, *BGP Common Structure Module*

openconfig-bgp-global.yang version 3.0.1, *BGP Global Module*

openconfig-bgp-neighbor.yang version 3.0.1, *BGP Neighbor Module*

openconfig-bgp-peer-group.yang version 3.0.1, *BGP Peer Group Module*

openconfig-bgp-policy.yang version 4.0.1, *BGP Policy Module*

openconfig-if-aggregate.yang version 2.0.0, *Interfaces Aggregated Model*

openconfig-if-ethernet.yang version 2.0.0, *Interfaces Ethernet Model*

openconfig-if-ip.yang version 2.0.0, *Interfaces IP Module*

openconfig-if-ip-ext.yang version 2.0.0, *Interfaces IP Extensions Module*

openconfig-interfaces.yang version 2.0.0, *Interfaces Module* openconfig-isis.yang version 0.3.0, *IS-IS Module*

openconfig-isis-lsp.yang version 0.3.0, *IS-IS LSP Module*

openconfig-isis-routing.yang version 0.3.0, *IS-IS Routing Module*

openconfig-lacp.yang version 1.1.0, *LACP Module*

openconfig-lldp.yang version 0.1.0, *LLDP Module*

openconfig-local-routing.yang version 1.0.1, *Local Routing Module*

openconfig-network-instance.yang version 0.8.0, *Network Instance Module*

openconfig-routing-policy.yang version 3.0.0, *Routing Policy Module*

openconfig-vlan.yang version 2.0.0, *VLAN Module*

RFC 1157, *A Simple Network Management Protocol (SNMP)*

RFC 1212, *Concise MIB Definitions*

RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II*

RFC 1215, *A Convention for Defining Traps for use with the SNMP*

RFC 1724, *RIP Version 2 MIB Extension*

RFC 1901, *Introduction to Community-based SNMPv2*

RFC 2021, *Remote Network Monitoring Management Information Base Version 2 using SMIv2*

RFC 2115, *Management Information Base for Frame Relay DTEs Using SMIv2*

RFC 2206, *RSVP Management Information Base using SMIv2*

RFC 2213, *Integrated Services Management Information Base using SMIv2*

RFC 2494, *Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type*

RFC 2514, *Definitions of Textual Conventions and OBJECT-IDENTITIES for ATM Management*

RFC 2515, *Definitions of Managed Objects for ATM Management*

RFC 2570, *SNMP Version 3 Framework*

RFC 2571, *An Architecture for Describing SNMP Management Frameworks*

RFC 2572, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*

RFC 2573, *SNMP Applications*

RFC 2574, *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*

RFC 2575, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*

RFC 2578, *Structure of Management Information Version 2 (SMIv2)*

RFC 2579, *Textual Conventions for SMIv2*

RFC 2580, *Conformance Statements for SMIv2*

RFC 2787, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol*

RFC 2819, *Remote Network Monitoring Management Information Base*

RFC 2856, *Textual Conventions for Additional High Capacity Data Types*

RFC 2863, *The Interfaces Group MIB*

RFC 2864, *The Inverted Stack Table Extension to the Interfaces Group MIB*

RFC 2933, *Internet Group Management Protocol MIB*

RFC 3014, *Notification Log MIB*

RFC 3164, *The BSD syslog Protocol*

RFC 3165, *Definitions of Managed Objects for the Delegation of Management Scripts*

RFC 3231, *Definitions of Managed Objects for Scheduling Management Operations*

RFC 3273, *Remote Network Monitoring Management Information Base for High Capacity Networks*

RFC 3416. *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*

RFC 3417, *Transport Mappings for the Simple Network Management Protocol (SNMP)* (SNMP over UDP over IPv4)

RFC 3419, *Textual Conventions for Transport Addresses*

RFC 3498, *Definitions of Managed Objects for Synchronous Optical Network (SONET) Linear Automatic Protection Switching (APS) Architectures*

RFC 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework*

RFC 3592, *Definitions of Managed Objects for the Synchronous Optical Network/ Synchronous Digital Hierarchy (SONET/SDH) Interface Type*

RFC 3593, *Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals*

RFC 3635, *Definitions of Managed Objects for the Ethernet-like Interface Types*

RFC 3637, *Definitions of Managed Objects for the Ethernet WAN Interface Sublayer*

RFC 3826, *The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model*

RFC 3877, *Alarm Management Information Base (MIB)*

RFC 3895, *Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types*

RFC 3896, *Definitions of Managed Objects for the DS3/E3 Interface Type*

RFC 4001, *Textual Conventions for Internet Network Addresses*

RFC 4022, *Management Information Base for the Transmission Control Protocol (TCP)*

RFC 4113, *Management Information Base for the User Datagram Protocol (UDP)*

RFC 4220, *Traffic Engineering Link Management Information Base*

RFC 4273, *Definitions of Managed Objects for BGP-4*

RFC 4292, *IP Forwarding Table MIB*

RFC 4293, *Management Information Base for the Internet Protocol (IP)*

RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*

RFC 4511, *Lightweight Directory Access Protocol (LDAP): The Protocol*

RFC 4513, *Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms* (TLS)

RFC 4631, *Link Management Protocol (LMP) Management Information Base (MIB)*

RFC 4878, *Definitions and Managed Objects for Operations, Administration, and Maintenance (OAM) Functions on Ethernet-Like Interfaces*

RFC 5101, *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information*

RFC 5102, *Information Model for IP Flow Information Export*

RFC 5246, *The Transport Layer Security (TLS) Protocol Version 1.2* (TLS client, RSA public key)

RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*

RFC 6425, *Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping*

RFC 6991, *Common YANG Data Types*

RFC 7420, *Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module*

RFC 7950, *The YANG 1.1 Data Modeling Language*

SFLOW-MIB, *sFlow MIB Version 1.3 (Draft 5)*

Multiprotocol Label Switching — Transport Profile (MPLS-TP)

RFC 5586, *MPLS Generic Associated Channel*

RFC 5921, *A Framework for MPLS in Transport Networks*

RFC 5960, *MPLS Transport Profile Data Plane Architecture*

RFC 6370, *MPLS Transport Profile (MPLS-TP) Identifiers*

RFC 6378, *MPLS Transport Profile (MPLS-TP) Linear Protection*

RFC 6426, *MPLS On-Demand Connectivity and Route Tracing*

RFC 6427, *MPLS Fault Management Operations, Administration, and Maintenance (OAM)*

RFC 6428, *Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile*

RFC 6478, *Pseudowire Status for Static Pseudowires*

RFC 7213, *MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing*

Multiprotocol Label Switching (MPLS)

draft-ietf-teas-sr-rsvp-coexistence-rec-02, *Recommendations for RSVP-TE and Segment Routing LSP co-existence*

RFC 3031, *Multiprotocol Label Switching Architecture*

RFC 3032, *MPLS Label Stack Encoding*

RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*

RFC 4023, *Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)*

RFC 4182, *Removing a Restriction on the use of MPLS Explicit NULL*

RFC 5332, *MPLS Multicast Encapsulations*

RFC 5884, *Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)*

RFC 6790, *The Use of Entropy Labels in MPLS Forwarding*

RFC 7510, *Encapsulating MPLS in UDP*

Network Address Translation (NAT)

draft-ietf-behave-address-format-10, *IPv6 Addressing of IPv4/IPv6 Translators*

draft-ietf-behave-v6v4-xlate-23, *IP/ICMP Translation Algorithm*

draft-miles-behave-l2nat-00, *Layer2-Aware NAT*

draft-nishitani-cgn-02, *Common Functions of Large Scale NAT (LSN)*

RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*

RFC 5382, *NAT Behavioral Requirements for TCP*

RFC 5508, *NAT Behavioral Requirements for ICMP*

RFC 6146, *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*

RFC 6333, *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*

RFC 6334, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite*

RFC 6887, *Port Control Protocol (PCP)*

RFC 6888, *Common Requirements For Carrier-Grade NATs (CGNs)*

RFC 7915, *IP/ICMP Translation Algorithm*

Network Configuration Protocol (NETCONF)

RFC 5277, *NETCONF Event Notifications*

RFC 6022, *YANG Module for NETCONF Monitoring*

RFC 6241, *Network Configuration Protocol (NETCONF)*

RFC 6242, *Using the NETCONF Protocol over Secure Shell (SSH)*

RFC 6243, *With-defaults Capability for NETCONF*

RFC 7895, *YANG Module Library*

Open Shortest Path First (OSPF)

draft-ietf-ospf-ospfv3-lsa-extend-13, *OSPFv3 LSA Extendibility*

RFC 1586, *Guidelines for Running OSPF Over Frame Relay Networks*

RFC 1765, *OSPF Database Overflow*

RFC 2328, *OSPF Version 2*

RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*

RFC 3509, *Alternative Implementations of OSPF Area Border Routers*

RFC 3623, *Graceful OSPF Restart Graceful OSPF Restart* (helper mode)

RFC 3630, *Traffic Engineering (TE) Extensions to OSPF Version 2*

RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*

RFC 4222, *Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance*

RFC 4552, *Authentication/Confidentiality for OSPFv3*

RFC 4576, *Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 5185, *OSPF Multi-Area Adjacency*

RFC 5187, *OSPFv3 Graceful Restart* (helper mode)

RFC 5243, *OSPF Database Exchange Summary List Optimization*

RFC 5250, *The OSPF Opaque LSA Option*

RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*

RFC 5340, *OSPF for IPv6*

RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*

RFC 5838, *Support of Address Families in OSPFv3*

RFC 6987, *OSPF Stub Router Advertisement*

RFC 7684, *OSPFv2 Prefix/Link Attribute Advertisement*

RFC 7770, *Extensions to OSPF for Advertising Optional Router Capabilities*

OpenFlow

TS-007, *OpenFlow Switch Specification Version 1.3.1* (OpenFlow-hybrid switches)

Path Computation Element Protocol (PCEP)

draft-alvarez-pce-path-profiles-04, *PCE Path Profiles*

draft-ietf-pce-segment-routing-08, *PCEP Extensions for Segment Routing*

draft-ietf-pce-stateful-pce-14, *PCEP Extensions for Stateful PCE*

RFC 5440, *Path Computation Element (PCE) Communication Protocol (PCEP)*

Point-to-Point Protocol (PPP)

RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*

RFC 1377, *The PPP OSI Network Layer Control Protocol (OSINLCP)*

RFC 1661, *The Point-to-Point Protocol (PPP)*

RFC 1662, *PPP in HDLC-like Framing*

RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*

RFC 1989, *PPP Link Quality Monitoring*

RFC 1990, *The PPP Multilink Protocol (MP)*

RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*

RFC 2153, *PPP Vendor Extensions*

RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*

RFC 2615, *PPP over SONET/SDH*

RFC 2686, *The Multi-Class Extension to Multi-Link PPP*

RFC 2878, *PPP Bridging Control Protocol (BCP)*

RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*

RFC 5072, *IP Version 6 over PPP*

Policy Management and Credit Control

3GPP TS 29.212 Release 11, *Policy and Charging Control (PCC); Reference points* (Gx support as it applies to wireline environment (BNG))

RFC 3588, *Diameter Base Protocol*

RFC 4006, *Diameter Credit-Control Application*

Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, *OAM Procedures for VPWS Interworking*

MFA Forum 12.0.0, *Multiservice Interworking - Ethernet over MPLS*

MFA Forum 13.0.0, *Fault Management for Multiservice Interworking v1.0*

MFA Forum 16.0.0, *Multiservice Interworking - IP over MPLS*

MFA Forum 9.0.0, *The Use of Virtual trunks for ATM/MPLS Control Plane Interworking*

RFC 3916, *Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)*

RFC 3985, *Pseudo Wire Emulation Edge-to-Edge (PWE3)*

RFC 4385, *Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN*

RFC 4446, *IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)*

RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*

RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*

RFC 4619, *Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks*

RFC 4717, *Encapsulation Methods for Transport Asynchronous Transfer Mode (ATM) over MPLS Networks*

RFC 4816, *Pseudowire Emulation Edge-to-Edge (PWE3) Asynchronous Transfer Mode (ATM) Transparent Cell Transport Service*

RFC 5085, *Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires*

RFC 5659, *An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge*

RFC 5885, *Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)*

RFC 6073, *Segmented Pseudowire*

RFC 6310, *Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping*

RFC 6391, *Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network*

RFC 6575, *Address Resolution Protocol (ARP) Mediation for IP Interworking of Layer 2 VPNs*

RFC 6718, *Pseudowire Redundancy*

RFC 6829, *Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6*

RFC 6870, *Pseudowire Preferential Forwarding Status bit*

RFC 7023, *MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking*

RFC 7267, *Dynamic Placement of Multi-Segment Pseudowires*

Quality of Service (QoS)

RFC 2430, *A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)*

RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*

RFC 2598, *An Expedited Forwarding PHB*

RFC 3140, *Per Hop Behavior Identification Codes*

RFC 3260, *New Terminology and Clarifications for Diffserv*

Remote Authentication Dial In User Service (RADIUS)

RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*

RFC 2866, *RADIUS Accounting*

RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support*

RFC 2868, *RADIUS Attributes for Tunnel Protocol Support*

RFC 2869, *RADIUS Extensions*

RFC 3162, *RADIUS and IPv6*

RFC 4818, *RADIUS Delegated-IPv6-Prefix Attribute*

RFC 5176, *Dynamic Authorization Extensions to RADIUS*

RFC 6911, *RADIUS attributes for IPv6 Access Networks*

RFC 6929, *Remote Authentication Dial-In User Service (RADIUS) Protocol Extensions*

Resource Reservation Protocol — Traffic Engineering (RSVP-TE)

draft-newton-mpls-te-dynamic-overbooking-00, *A Diffserv-TE Implementation Model to dynamically change booking factors during failure events*

RFC 2702, *Requirements for Traffic Engineering over MPLS*

RFC 2747, *RSVP Cryptographic Authentication*

RFC 2961, *RSVP Refresh Overhead Reduction Extensions*

RFC 3097, *RSVP Cryptographic Authentication -- Updated Message Type Value*

RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*

RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions* (IF_ID RSVP_HOP object with unnumbered interfaces and RSVP-TE graceful restart helper procedures)

RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*

RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*

RFC 3906, *Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels*

RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

RFC 4124, *Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering*

RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*

RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)*

RFC 4950, *ICMP Extensions for Multiprotocol Label Switching*

RFC 5151, *Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions*

RFC 5712, *MPLS Traffic Engineering Soft Preemption*

RFC 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks*

Routing Information Protocol (RIP)

RFC 1058, *Routing Information Protocol*

RFC 2080, *RIPng for IPv6*

RFC 2082, *RIP-2 MD5 Authentication*

RFC 2453, *RIP Version 2*

Segment Routing (SR)

draft-filsfils-spring-segment-routing-policy-05, *Segment Routing Policy for Traffic Engineering*

draft-francois-rtgwg-segment-routing-ti-lfa-04, *Topology Independent Fast Reroute using Segment Routing*

draft-gredler-idr-bgp-ls-segment-routing-ext-03, *BGP Link-State extensions for Segment Routing*

draft-ietf-idr-segment-routing-te-policy-02, *Advertising Segment Routing Policies in BGP*

draft-ietf-isis-segment-routing-extensions-04, *IS-IS Extensions for Segment Routing* draft-ietf-mpls-spring-lsp-ping-02, *Label Switched Path (LSP) Ping/Trace for Segment Routing Networks Using MPLS Dataplane*

draft-ietf-ospf-segment-routing-extensions-04, *OSPF Extensions for Segment Routing*

draft-ietf-spring-conflict-resolution-05, *Segment Routing MPLS Conflict Resolution*

draft-ietf-spring-segment-routing-ldp-interop-09, *Segment Routing interworking with LDP*

Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH)

ANSI T1.105.03, *Jitter Network Interfaces*

ANSI T1.105.06, *Physical Layer Specifications*

ANSI T1.105.09, *Network Timing and Synchronization*

ITU-T G.703, *Physical/electrical characteristics of hierarchical digital interfaces*

ITU-T G.707, *Network node interface for the synchronous digital hierarchy (SDH)*

ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC)*

ITU-T G.823, *The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy*

ITU-T G.824, *The control of jitter and wander within digital networks which are based on the 1544 kbit/s hierarchy*

ITU-T G.825, *The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)*

ITU-T G.841, *Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum 1, issued in July 2002*

ITU-T G.957, *Optical interfaces for equipments and systems relating to the synchronous digital hierarchy*

Time Division Multiplexing (TDM)

ANSI T1.403, *DS1 Metallic Interface Specification*

ANSI T1.404, *DS3 Metallic Interface Specification*

Timing

GR-1244-CORE, *Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005*

GR-253-CORE, *SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000*

IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*

ITU-T G.781, *Synchronization layer functions, issued 09/2008*

ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003*

ITU-T G.8261, *Timing and synchronization aspects in packet networks, issued 04/2008*

ITU-T G.8262, *Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007*

ITU-T G.8264, *Distribution of timing information through packet networks, issued 10/2008*

ITU-T G.8265.1, *Precision time protocol telecom profile for frequency synchronization, issued 10/2010*

ITU-T G.8275.1, *Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014*

RFC 5905, *Network Time Protocol Version 4: Protocol and Algorithms Specification*

Two-Way Active Measurement Protocol (TWAMP)

RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP)* (server, unauthenticated mode)

RFC 5938, *Individual Session Control Feature for the Two-Way Active Measurement Protocol (TWAMP)*

RFC 6038, *Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features*

Virtual Private LAN Service (VPLS)

RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*

RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*

RFC 5501, *Requirements for Multicast Support in Virtual Private LAN Services*

RFC 6074, *Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)*

RFC 7041, *Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging*

RFC 7117, *Multicast in Virtual Private LAN Service (VPLS)*

Voice and Video

DVB BlueBook A86, *Transport of MPEG-2 TS Based DVB Services over IP Based Networks*

ETSI TS 101 329-5 Annex E, *QoS Measurement for VoIP - Method for determining an Equipment Impairment Factor using Passive Monitoring*

ITU-T G.1020 Appendix I, *Performance Parameter Definitions for Quality of Speech and other Voiceband Applications Utilizing IP Networks - Mean Absolute Packet Delay Variation & Markov Models*

ITU-T G.107, *The E Model - A computational model for use in planning*

ITU-T P.564, *Conformance testing for voice over IP transmission quality assessment models*

RFC 3550 Appendix A.8, *RTP: A Transport Protocol for Real-Time Applications* (estimating the interarrival jitter)

RFC 4585, *Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)*

RFC 4588, *RTP Retransmission Payload Format*

Wireless Local Area Network (WLAN) Gateway

3GPP TS 23.402, *Architecture enhancements for non-3GPP accesses* (S2a roaming based on GPRS)

Customer Document and Product Support

Customer Documentation

Customer Documentation Welcome Page

Technical Support

Product Support Portal

Documentation Feedback

Customer Documentation Feedback

© 2018 Nokia.
3HE 14138 AAAB TQZZA 01