NOKIA

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

SERVICES OVERVIEW GUIDE
RELEASE 22.10.R1

3HE 18402 AAAD TQZZA 01
Issue 01

October 2022

Nokia is committed to diversity and inclusion. We are continuously reviewing our customer documentation and consulting with standards bodies to ensure that terminology is inclusive and aligned with the industry. Our future customer documentation will be updated accordingly.

Table of contents

1 Getting started 8
1.1 About this guide 8
1.2 Services configuration process 8
1.3 Conventions 9
1.3.1 Precautionary and information messages 9
1.3.2 Options or substeps in procedures and sequential workflows 9

2 Services overview 11
2.1 Introduction 11
2.1.1 Service types 11
2.1.2 Service policies 12
2.1.2.1 Multipoint shared queuing 12
2.2 Nokia service model 17
2.3 Service entities 18
2.3.1 Customers 18
2.3.2 SAPs 19
2.3.2.1 SAP encapsulation types and identifiers 19
2.3.2.2 Ethernet encapsulations 19
2.3.2.3 Default SAP on a dot1q port 21
2.3.2.4 QinQ SAPs 22
2.3.2.5 Services and SAP encapsulations 26
2.3.2.6 SAP configuration considerations 27
2.3.2.7 G.8032 protected Ethernet rings 28
2.3.2.8 SAP bandwidth CAC 28
2.3.3 Connection profile VLAN SAPs 30
2.3.3.1 Using connection-profile-vlan in dot1q ports 33
2.3.3.2 Using connection-profile-vlan in QinQ ports 33
2.3.4 Service distribution points 34
2.3.4.1 SDP binding 35
2.3.4.2 Spoke and mesh SDPs 35
2.3.4.3 SDP using BGP route tunnel 35
2.3.4.4 SDP keepalives 36
2.3.4.5 SDP administrative groups 36

2.3.4.6 SDP selection rules ... 37
2.3.4.7 Class-based forwarding ... 38
2.3.4.8 Source IPv4 address configuration in GRE SDP and GRE tunnel ... 39
2.3.4.9 GRE SDP tunnel fragmentation and reassembly ... 44
2.3.4.10 GRE SDP termination on router interface IP address ... 46
2.3.5 SAP and MPLS binding loopback with MAC swap ... 47
2.3.6 Promiscuous ETH-LBM mode of operation ... 51
2.4 Multi-service sites ... 52
2.5 G.8031 Protected Ethernet Tunnels ... 53
2.5.1 OAM considerations ... 56
2.5.2 QoS considerations ... 56
2.5.3 Mirroring and Lawful Intercept considerations ... 56
2.5.4 Support service and solution combinations ... 56
2.5.5 LAG emulation using Ethernet tunnels ... 57
2.6 G.8032 Ethernet ring protection switching ... 57
2.6.1 Overview of G.8032 operation ... 58
2.6.2 Ethernet ring sub-rings ... 62
2.6.2.1 Virtual and non-virtual channel ... 63
2.6.2.2 LAG support ... 69
2.6.3 OAM considerations ... 69
2.6.4 Support service and solution combinations ... 70
2.7 Internal objects created for L2TP and NAT ... 70
2.8 Ethernet unnumbered interfaces ... 70
2.9 ECMP and weighted ECMP for services using RSVP and SR-TE LSPs ... 71
2.10 NGE ... 71
2.10.1 NGE overview ... 71
2.10.1.1 NGE key groups and encryption partitions ... 73
2.10.1.2 Network services platform management ... 74
2.10.2 Key groups ... 75
2.10.2.1 Key group algorithms ... 76
2.10.2.2 Security associations ... 77
2.10.3 Services encryption ... 77
2.10.3.1 Services encryption overview ... 78
2.10.3.2 Assigning key groups to services ... 80
2.10.3.3 VPRN Layer 3 spoke SDP encryption and MP-BGP-based VPRN encryption interaction ... 81

2.10.3.4 L2 Service encryption using PW templates....82
2.10.3.5 Pseudowire switching for NGE traffic....82
2.10.3.6 Pseudowire control word for NGE traffic....82
2.10.3.7 WLAN-GW encryption....83
2.10.3.8 NGE and RFC 8277....83
2.10.3.9 NGE for NG-MVPN....83
2.10.4 NGE packet overhead and MTU considerations....84
2.10.5 Statistics....86
2.10.6 Remote network monitoring support....86
2.10.7 Configuration notes....87
2.10.7.1 Enabling NGE for an SDP or VPRN service....87
2.10.7.2 Enabling NGE for a router interface....87
2.10.7.3 Changing NGE from one key group to another key group for an SDP or VPRN service....87
2.10.7.4 Changing NGE from one key group to another key group for a router interface..88
2.10.7.5 Disabling NGE for an SDP or VPRN service....88
2.10.7.6 Disabling NGE for a router interface....88
2.11 Service creation process overview....88
2.12 Deploying and provisioning services....89
2.12.1 Building the core network....89
2.12.2 Performing service administration....90
2.12.3 Provisioning services....90
2.13 General configuration notes....90
2.14 Configuring global service entities with CLI....91
2.14.1 Service model entities....91
2.14.2 Basic configurations....92
2.14.3 Common configuration tasks....93
2.14.3.1 Configuring customers....93
2.14.3.2 Configuring an SDP....96
2.15 Ethernet connectivity fault management (ETH-CFM)....99
2.15.1 Facility MEPs....101
2.15.1.1 Common actionable failures....102
2.15.1.2 General detection, processing and reaction....103
2.15.1.3 Port-based MEP....105
2.15.1.4 LAG based MEP....112
2.15.1.5 Tunnel-based MEP....118

2.15.1.6 Router interface MEP ... 127
2.15.1.7 Hardware support ... 130
2.15.2 ETH-CFM and MC-LAG ... 130
2.15.2.1 ETH-CFM and MC-LAG default behavior ... 131
2.15.2.2 Linking ETH-CFM to MC-LAG state ... 132
2.15.3 ETH-CFM feature: CCM hold timers ... 138
2.15.4 Configuring ETH-CFM parameters ... 139
2.16 Configuring NGE with CLI ... 142
2.16.1 Basic NGE configuration overview ... 142
2.16.2 Configuring NGE components ... 142
2.16.2.1 Configuring the global encryption label ... 143
2.16.2.2 Configuring a key group ... 143
2.16.2.3 Assigning a key group to an SDP, VPRN service, or PW template ... 144
2.17 Global service entity management tasks ... 146
2.17.1 Modifying customer accounts ... 146
2.17.2 Deleting customers ... 146
2.17.3 Modifying SDPs ... 147
2.17.4 Deleting SDPs ... 147
2.18 NGE management tasks ... 147
2.18.1 Modifying a key group ... 148
2.18.2 Removing a key group ... 149
2.18.2.1 Removing a key group from an SDP, VPRN service, or PW template ... 149
2.18.3 Changing key groups ... 150
2.18.4 Changing the key group for an SDP, VPRN service, or PW template ... 150
2.18.5 Deleting a key group from an NGE node ... 151

3 Standards and protocol support ... 152
3.1 Access Node Control Protocol (ANCP) ... 152
3.2 Application Assurance (AA) ... 152
3.3 Bidirectional Forwarding Detection (BFD) ... 152
3.4 Border Gateway Protocol (BGP) ... 152
3.5 Broadband Network Gateway (BNG) Control and User Plane Separation (CUPS) ... 154
3.6 Certificate management ... 154
3.7 Circuit emulation ... 155
3.8 Ethernet ... 155
3.9 Ethernet VPN (EVPN) ... 156

3.10 gRPC Remote Procedure Calls (gRPC)....156
3.11 Intermediate System to Intermediate System (IS-IS)....156
3.12 Internet Protocol (IP) Fast Reroute (FRR)....158
3.13 Internet Protocol (IP) general....158
3.14 Internet Protocol (IP) multicast....159
3.15 Internet Protocol (IP) version 4....161
3.16 Internet Protocol (IP) version 6....161
3.17 Internet Protocol Security (IPsec)....162
3.18 Label Distribution Protocol (LDP)....164
3.19 Layer Two Tunneling Protocol (L2TP) Network Server (LNS)....164
3.20 Multiprotocol Label Switching (MPLS)....164
3.21 Multiprotocol Label Switching - Transport Profile (MPLS-TP)....165
3.22 Network Address Translation (NAT)....165
3.23 Network Configuration Protocol (NETCONF)....166
3.24 Open Shortest Path First (OSPF)....166
3.25 OpenFlow....167
3.26 Path Computation Element Protocol (PCEP)....167
3.27 Point-to-Point Protocol (PPP)....168
3.28 Policy management and credit control....168
3.29 Pseudowire (PW)....168
3.30 Quality of Service (QoS)....169
3.31 Remote Authentication Dial In User Service (RADIUS)....169
3.32 Resource Reservation Protocol - Traffic Engineering (RSVP-TE)....170
3.33 Routing Information Protocol (RIP)....171
3.34 Segment Routing (SR)....171
3.35 Simple Network Management Protocol (SNMP)....172
3.36 Timing....174
3.37 Two-Way Active Measurement Protocol (TWAMP)....175
3.38 Virtual Private LAN Service (VPLS)....175
3.39 Voice and video....175
3.40 Wireless Local Area Network (WLAN) gateway....176
3.41 Yet Another Next Generation (YANG)....176
3.42 Yet Another Next Generation (YANG) OpenConfig Modules....176

1 Getting started

1.1 About this guide

This guide describes subscriber services, and mirroring support provided by SR-series 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.

Note: Unless otherwise indicated, this guide uses classic CLI command syntax and configuration examples.

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

- 7450 ESS
- 7750 SR
- 7950 XRS
- Virtualized Service Router

For a list of unsupported features by platform and chassis, see the *SR OS R22.x.Rx Software Release Notes*, part number 3HE 18412 000 *x* TQZZA.

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

Note:
The SR OS CLI trees and command descriptions can be found in the following guides:

- *7450 ESS, 7750 SR, 7950 XRS, and VSR Classic CLI Command Reference Guide*
- *7450 ESS, 7750 SR, 7950 XRS, and VSR Clear, Monitor, Show, and Tools Command Reference Guide* (for both MD-CLI and Classic CLI)
- *7450 ESS, 7750 SR, 7950 XRS, and VSR MD-CLI Command Reference Guide*

Note:
This guide generically covers Release 22.*x*.R*x* content and may contain some content that will be released in later maintenance loads. See the *SR OS R22.x.Rx Software Release Notes*, part number 3HE 18412 000 *x* TQZZA for information about features supported in each load of the Release 22.*x*.R*x* software.

1.2 Services configuration process

Table 1: Configuration process 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 1: 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

1.3 Conventions

This section describes the general conventions used in this guide.

1.3.1 Precautionary and information messages

The following are information symbols used in the documentation.

DANGER: Danger warns that the described activity or situation may result in serious personal injury or death. An electric shock hazard could exist. Before you begin work on this equipment, be aware of hazards involving electrical circuitry, be familiar with networking environments, and implement accident prevention procedures.

WARNING: Warning indicates that the described activity or situation may, or will, cause equipment damage, serious performance problems, or loss of data.

Caution: Caution indicates that the described activity or situation may reduce your component or system performance.

Note: Note provides additional operational information.

Tip: Tip provides suggestions for use or best practices.

1.3.2 Options or substeps in procedures and sequential workflows

Options in a procedure or a sequential workflow are indicated by a bulleted list. In the following example, at step 1, the user must perform the described action. At step 2, the user must perform one of the listed options to complete the step.

Example: Options in a procedure

1. User must perform this step.
2. This step offers three options. User must perform one option to complete this step.

- This is one option.
- This is another option.
- This is yet another option.

Substeps in a procedure or a sequential workflow are indicated by letters. In the following example, at step 1, the user must perform the described action. At step 2, the user must perform two substeps (a. and b.) to complete the step.

Example: Substeps in a procedure

1. User must perform this step.
2. User must perform all substeps to complete this action.
   a. This is one substep.
   b. This is another substep.

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)**

    This is a Layer 2 point-to-point VLL service for Ethernet frames.
  - **IP pipe (Ipipe)**

    This provides 7750 SR and 7450 ESS IP connectivity between a host attached to a point-to-point access circuit (PPP) with routed IPv4 encapsulation and a host attached to an Ethernet interface.

  See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide* for more information about VLL services.
- **Virtual Private LAN Service (VPLS)**

  This is 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.

  See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide* for more information about VPLS.
- **Internet Enhanced Service (IES)**

This is a direct Internet access service where the customer is assigned an IP interface for Internet connectivity.

See 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)**

  This is a Layer 3 IP multipoint-to-multipoint VPN service as defined in RFC 2547bis.

  See 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**

  SAP QoS policies 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**

  Filter policies allow for 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**

  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**

  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: Unicast service queue mapping to multiple destination based hardware queues 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

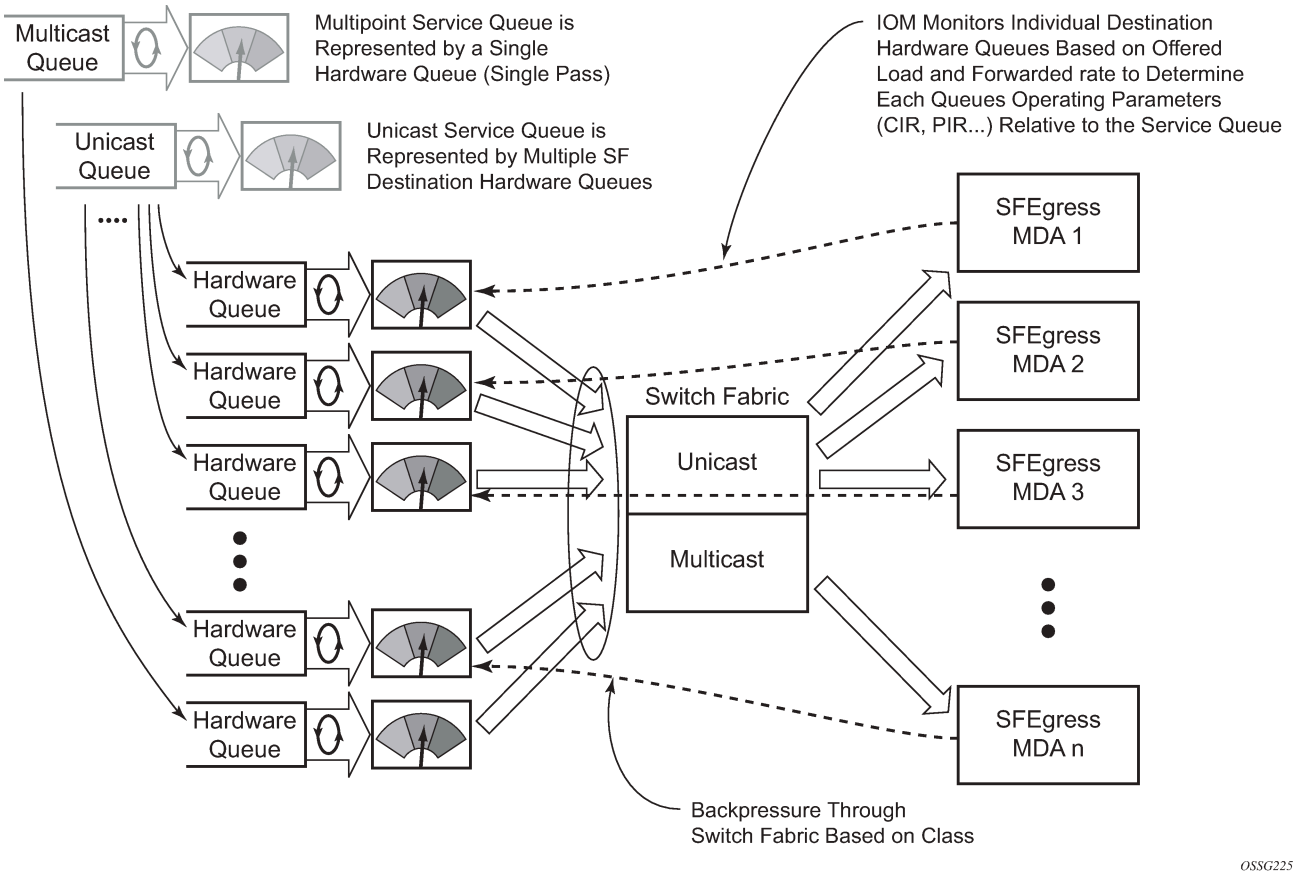

2.1.2.1.3 Ingress shared queuing

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: Unicast service queuing with shared queuing enabled 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: Multipoint queue behavior with shared queuing enabled.

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

Figure 2: Unicast service queuing with shared queuing enabled

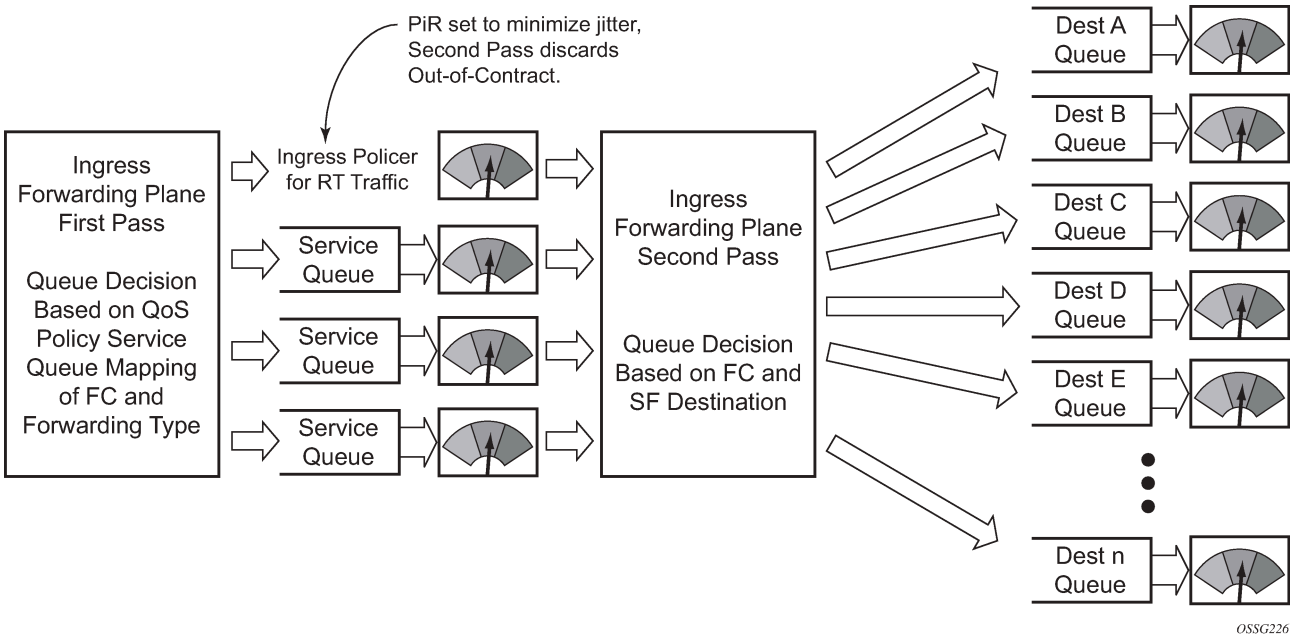

Figure 3: Multipoint queue behavior with shared queuing enabled

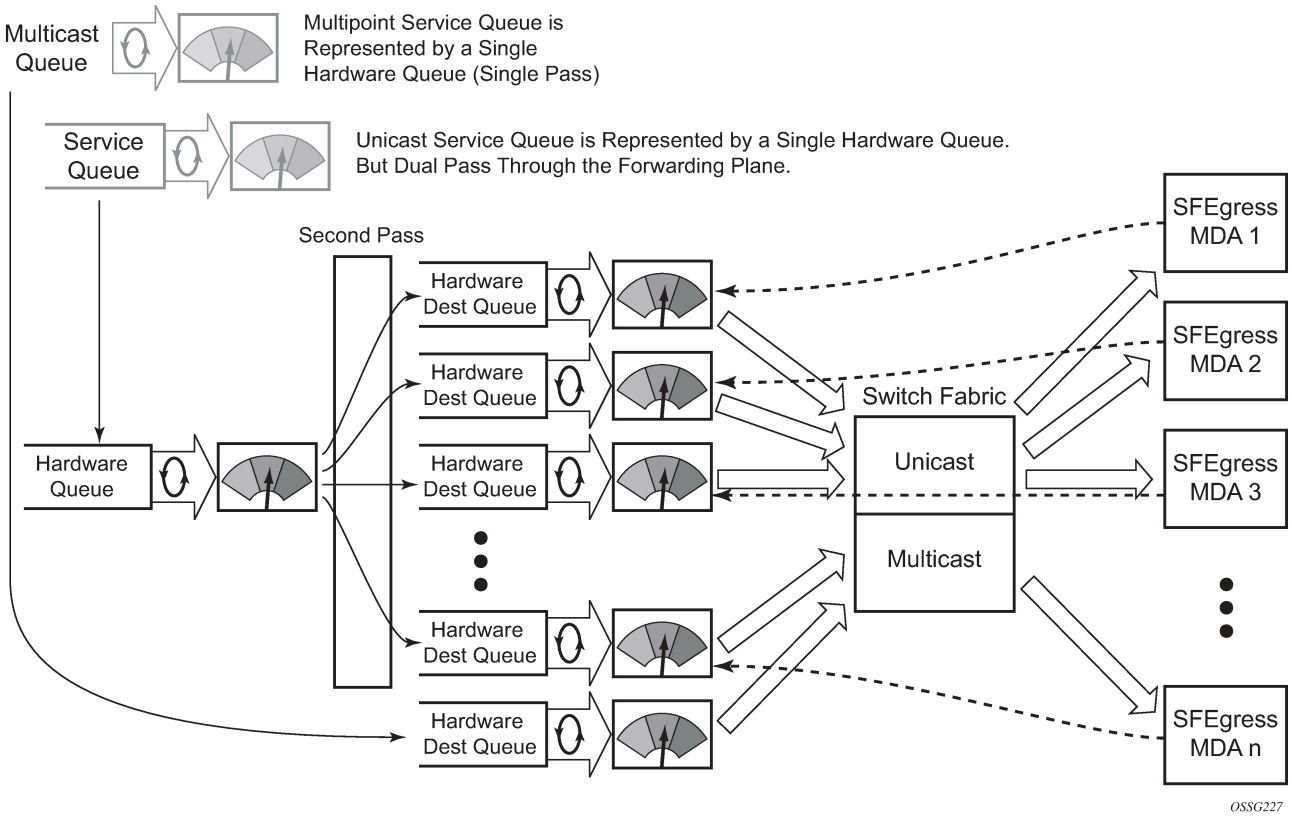

2.1.2.1.4 Ingress multipoint shared queuing

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: Multipoint shared queuing using first pass unicast queues 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 because of dual pass ingress processing.

Figure 4: Multipoint shared queuing using first pass unicast queues

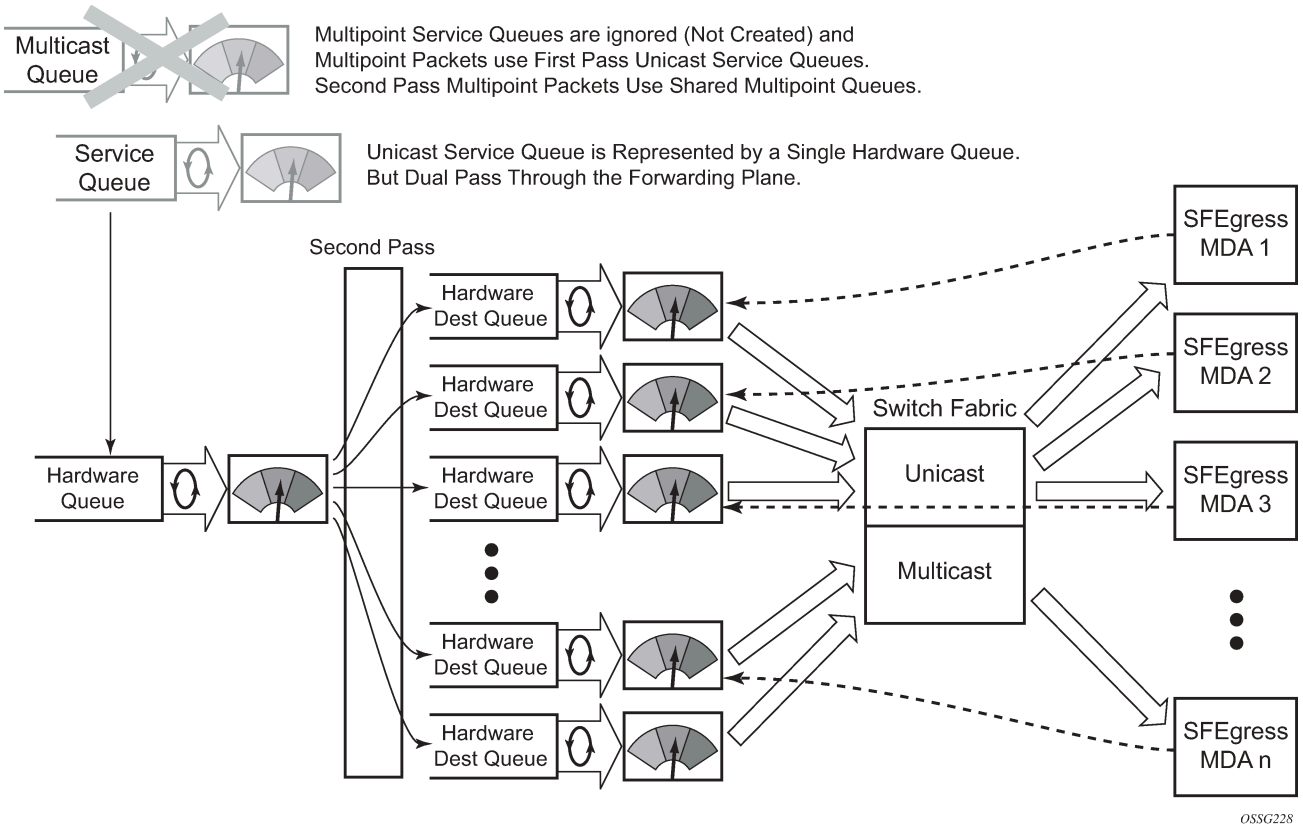

2.2 Nokia service model

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 instead of multiple ports on multiple devices. It is easier to change one tunnel instead of several services.
- The operational integrity of a logical entity (such as a service tunnel and service end points) can be verified instead of 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
- SAPs
- Service distribution points (for distributed services only)

Figure 5: Service entities

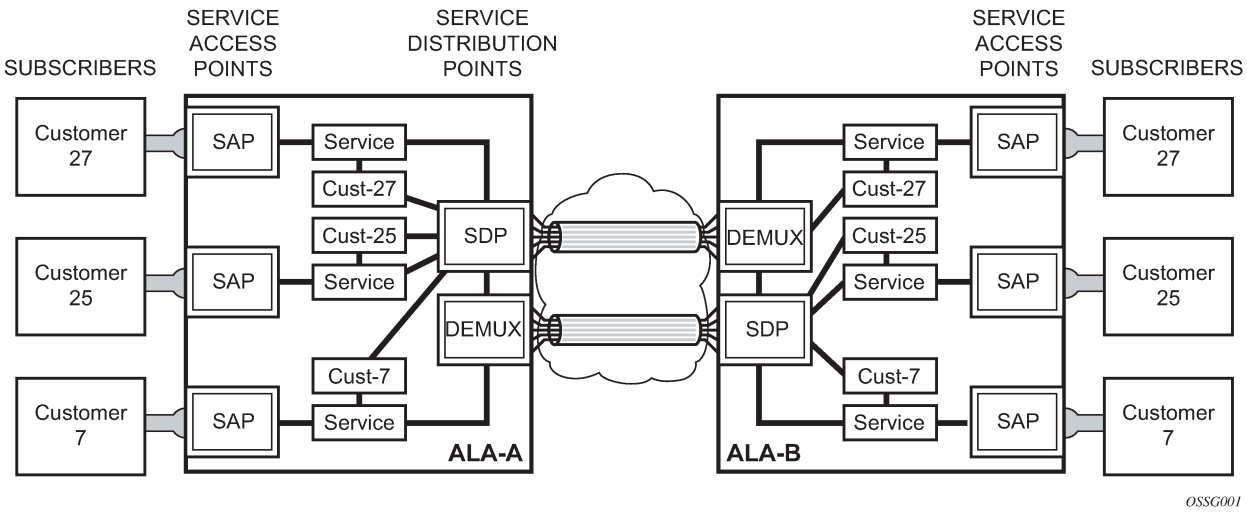

2.3.1 Customers

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 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: 7750 SR/7950 XRS service access point (SAP)). 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 before provisioning a service (See 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)

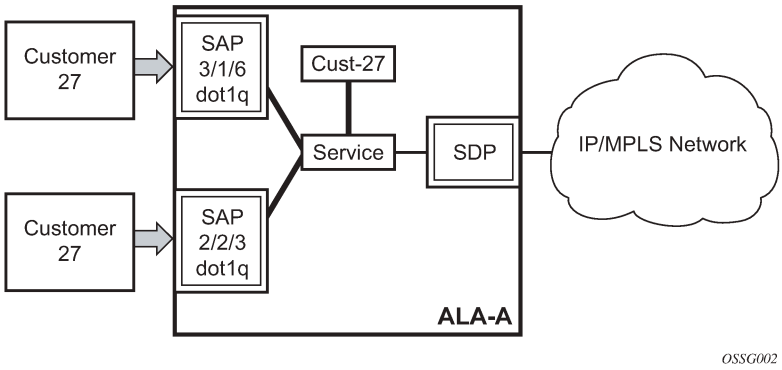

A SAP can also be associated with a pseudowire port instead of an access port. Such SAPs are called pseudowire SAPs. This is only applicable to IES, VPRN, and Epipe services. Pseudowire ports represent pseudowires in enhanced subscriber management (ESM). For a description of pseudowire ports, see the *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 specific 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**

  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**

  Dot1q supports multiple services for one customer or services for multiple customers (Figure 6: 7750 SR/7950 XRS service access point (SAP) and Figure 7: 7750 SR/7950 XRS and 7450 ESS multiple SAPs on a single port/channel). 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)**

  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)**

  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)**

  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.

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

- **Internet Protocol Control Protocol (IPCP)**

  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)**

  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)**

  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.

Figure 7: 7750 SR/7950 XRS and 7450 ESS multiple SAPs on a single port/channel

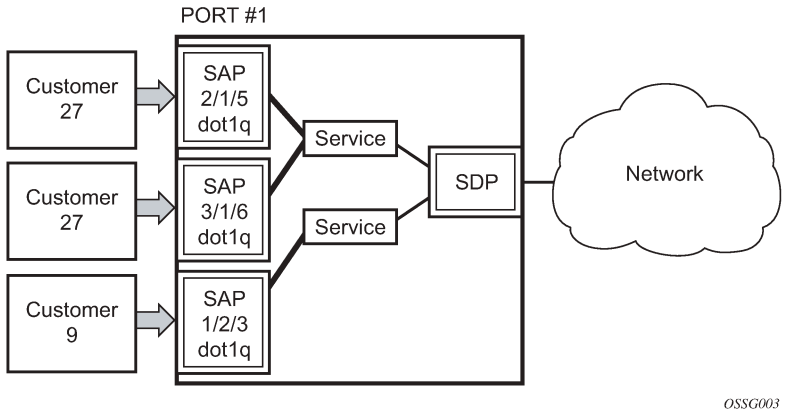

2.3.2.3 Default SAP on a dot1q port

On dot1q-encapsulated ports where a default SAP is configured, all packets with q-tags not matching any explicitly defined SAPs are 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 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 are transparently forwarded.
- This type of SAP and a SAP defined by explicit null encapsulation (for example, 1/1/1:0) are mutually exclusive. This avoids conflict as to which SAP untagged frames should be associated.

In a Dot1q port SAP with a non-zero or non-default tag, the tag (referred to as service-delimiting tag) is stripped off on ingress and pushed on egress. For example, the tag is popped from frames received on SAP 1/1/1:10 with a tag that contains VID 10. A tag with VID 10 is pushed onto frames that are sent out of SAP 1/1/1:10.

In case of Dot1q port SAPs with a zero or default tag, no tag is stripped off on ingress, and no tag is pushed on egress. For instance, tags are not stripped off from frames entering 1/1/1:* or 1/1/1:0, and tags are not pushed either on frames egressing those SAPs.

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 - [*, null, 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 singly-tagged frames in a QinQ port. This SAP accepts single tags in the range 0 to 4095 as well as untagged traffic. This SAP never pushes any tags on egress.
- '*.*' is defined as a default sap for doubly-tagged frames in a QinQ port. This SAP accepts untagged, singly tagged, and doubly tagged frames with tags in the range 0..4095. This SAP never pushes any tags on egress.
- 'null.null' is defined as a default SAP for untagged frames only in a QinQ port. This SAP accepts only untagged frames and never pushes any tags on egress. This SAP has higher priority than '*.null' and, or '*.*' when configured on the same QinQ port, therefore it captures untagged frames even if '*.null' or '*.*' are configured.
- '0.*' can also be used as a default SAP and captures untagged frames, doubly-tagged frames with qtag1 0 (and any value on qtag2) and singly-tagged frames with qtag 0. SAP '0.*' and 'null.null' cannot be configured on the same QinQ port.
- 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. The system never pushes any qtag2 on egress for 1/1/110.0 and 1/1/1:10.*, only qtag1 is pushed. The x.0 accepts only 0 as second tag or not tag (and nothing else), while x.* accepts anything as second tag or no tag.

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

Table 2: 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.null**	**:*.null**	**:*.***
x.y	1st		2nd				3rd

Incoming frame *qtag1.qtag2*	System/port settings [new-qinq-untagged-sap=YES]						
	SAP lookup precedence order						
	:X.Y	:X.0	:X.*	:0.*	:null.null	:*.null	:*.*
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	4th

The following considerations apply to the information described in Table 2: SAP lookup precedence order for incoming frames:

- All SAP types (:X.Y, :X.0, :X.*, :0.* or :null.null, :*.null and :*.*) are supported in the same QinQ port (with the exception of :0.* and :null.null being incompatible) and, in the table, they are ordered from the most specific (left 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'
  - 0 means VID 0 or untagged
- On egress, the system pushes a tag with a VID value for X and Y, whereas no tag is pushed on egress for values 0, * or null. For example:
  - On SAP 1/1/1:10.20, the system pushes tags with VIDs 10 and 20 (outer and inner respectively) on egress.
  - On SAP 1/1/1:10.0 or 1/1/1:10.* the system pushes only one tag with VID 10 on egress.
  - On SAPs 1/1/1:0.*, 1/1/1:null.null, 1/1/1:*.null or 1/1/1:*.* the system never pushes any tags on egress.
- 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 2: SAP lookup precedence order for incoming frames shows the lookup behavior for ingress frames and priority across SAPs in case more than one can match a specific ingress frame. The SAP lookup result for a specific 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 specific frame with VIDs (0, 300) yields SAP 1/1/1:0.* regardless of its operational status.
  - The frame only matches SAP 1/1/1:*.* when the 0.* SAP is removed from the configuration.
- The following applies to VLAN tag handling:
  - The system does not strip-off any tags for frames entering the default SAPs (:0.*, :null.null, :*.null or :*.*).

– No extra tags are added when the system transmits frames on the default SAPs (:0.*, :null.null, :*.null or :*.*).

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

As shown in Figure 8: Example 1 SAP classification QinQ ports, 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

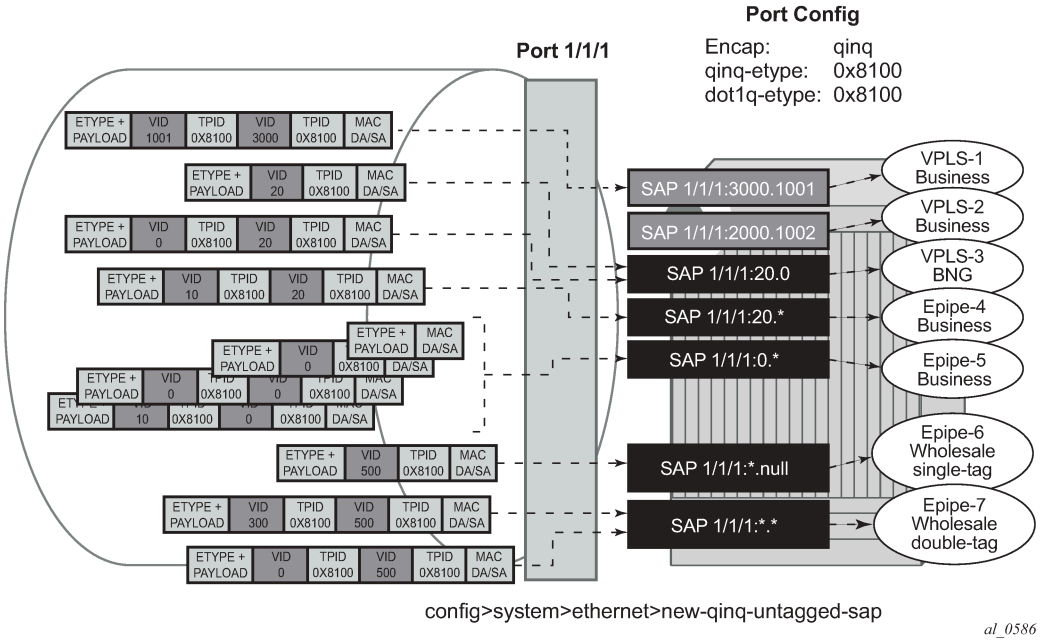

Based on the SAPs configuration described above, the incoming traffic is classified in the following way - notation (outer-VID, inner-VID):

- (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)

Figure 9: Example 2 SAP classification QinQ ports 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

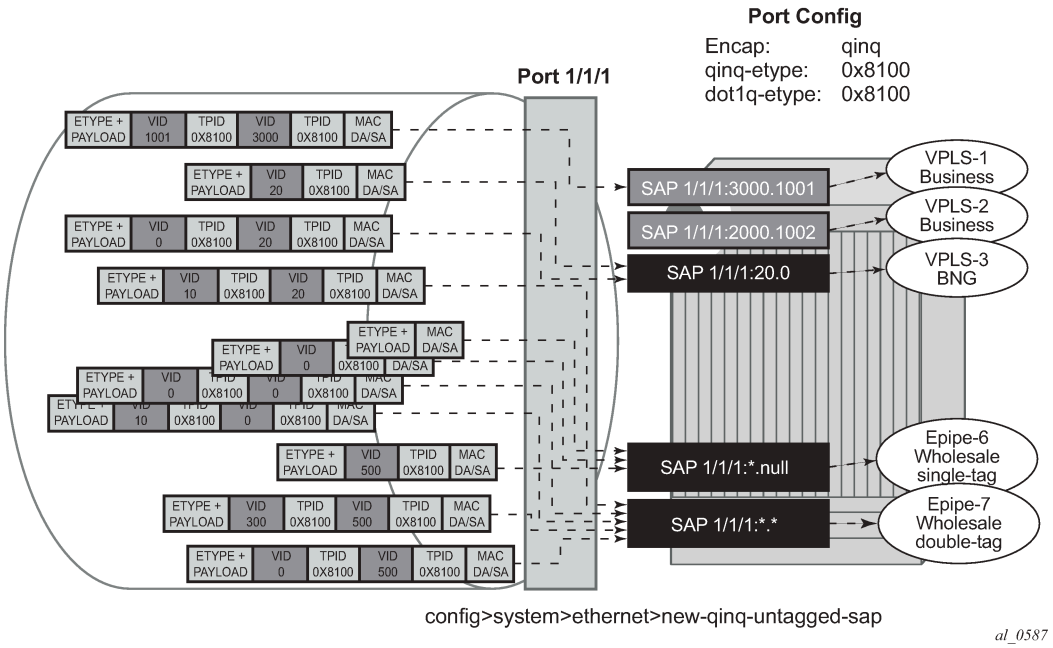

As described in Figure 9: Example 2 SAP classification QinQ ports, assuming the **new-qinq-untagged-sap** command is enabled, the following SAPs are defined on the same port:

- 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 does 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.null, :*.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 four SAPs (null.null, *.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 are transmitted out all four 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 2: SAP lookup precedence order for incoming frames 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 3: Service and SAP encapsulations.

Table 3: 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	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 specific 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 are also deleted. For Internet Enhanced Service (IES), the IP interface must be shut down 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 are 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 about Ethernet rings, see G.8032 Ethernet ring protection switching.

2.3.2.8 SAP bandwidth CAC

This feature provides a bandwidth Connection Admission Control (CAC) function per port or LAG based on an administrator bandwidth configured on a SAP and on the associated port or LAG. A booking factor is provided to allow overbooking or underbooking of the sum of the SAP bandwidth compared to the port or LAG bandwidth.

The administrator bandwidth is a statically configured abstract value that could represent either the ingress or egress bandwidth, or both.

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

This feature 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 administrator bandwidth configured, in which case it is excluded from the CAC function. Configuring an administrator bandwidth on a SAP enforces the CAC function.

An administrator bandwidth can only be configured on a SAP that is connected to a port or LAG on which an administrator bandwidth is already configured. When a LAG is configured, the administrator 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 administrator bandwidth on the port or LAG, with the booking factor applied, minus the sum of administrator bandwidth configured on its SAPs. An attempt to increase a SAP's administrator bandwidth fails if there is insufficient available bandwidth on its port or LAG.

Use the following CLI commands to configure the administrator bandwidth and booking factor for the port or LAG.

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

```
                    bandwidth <bw-value>
```

Dynamic changes in administrator bandwidth and booking factor are possible without having to disable the SAP, port, or LAG.

A SAP is allocated bandwidth on a port or LAG regardless of whether the SAP and port or LAG are administratively or operationally up or down. The administrator bandwidth must be removed from the SAP configuration to free up its bandwidth on the port or LAG. Actions such as clearing the card or MDA, power cycling the card, or removing and reinserting a card or MDA do not change the CAC state of the SAP and port or LAG.

2.3.2.8.1 CAC enforcement

The CAC is enforced when an administrator bandwidth is configured on a SAP, which may be when the administrator bandwidth is initially configured or when an existing administrator bandwidth value is modified.

The CAC enforcement is achieved by comparing the newly requested SAP administrator bandwidth (the incremental administrator bandwidth being configured above any currently assigned administrator bandwidth) with the available administrator 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.

Example

In the following example, a port is configured with an administrator bandwidth of 500 Mb/s, and a SAP on that port is configured with a bandwidth of 10 Mb/s. The **show** output displays these configured values together with the available and booked administrator bandwidth for the port. An

increase of the SAP administrator bandwidth to 600 Mb/s is attempted; the operation fails because of insufficient available administrator bandwidth on the port.

The booking factor for the port is increased to 200% and the increase of the SAP administrator bandwidth to 600 Mb/s is successful as the available administrator bandwidth for the port becomes 1 Gb/s. The booked administrator bandwidth for the port is 600 Mb/s and so the available administrator bandwidth for the port 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 specific 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 example shows the command usage 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
```

The following is an example configuration output:

```
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>
```

For VLAN manipulation, 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:

- A default SAP consumes less resources; a default SAP consumes one SAP instance, whereas a CP SAP consumes SAP instances equal to the number of VLANs in the range. To check the number of SAP instances used by the system, run the following CLI commands:

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

```
*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>
===============================================================================
```

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

Figure 10: VLAN tag handling shows customer VID processing by SAPs with service-delimiting VIDs, and by CP SAPs. In this example, 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

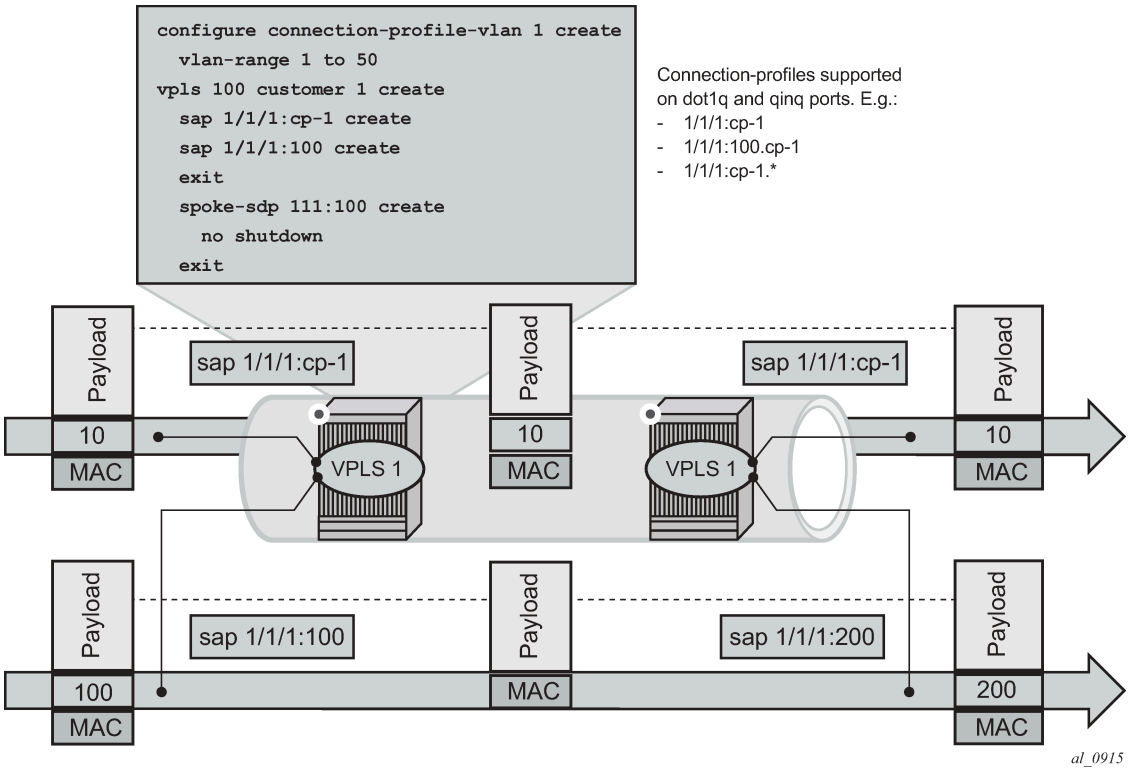

A **connection-profile-vlan** allows the configuration of VLAN ranges with the following characteristics:

- 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** configuration. Overlapping VLAN ranges can exist across different connection profiles if 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 VLAN within the connection-profile-vlan.
- The legacy OAM commands (**mac-ping**, **mac-trace**, **mac-purge**, and **mac-populate**) are not supported for CP SAPs.

2.3.3.1 Using connection-profile-vlan in dot1q ports

Table 4: SAP lookup matching order for dot1q ports describes the SAP lookup matching order that is applied when **connection-profile-vlan** is used in dot1q ports.

Table 4: 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 5: SAP lookup matching order for QinQ ports describes the SAP lookup matching order that is applied when **connection-profile-vlan** is used in QinQ ports.

Table 5: 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

The following considerations apply when connection profile VLAN (CP VLAN) is used 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 with an explicitly defined inner or outer tag in a specific port. For example, assuming that X and Y are tags defined in "CP", a specific port can be defined with ":X.CP" or ":Y.CP", but not with ":X.Y" and ":X.CP" or ":CP.*" and "X.*" in the same port.
- 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 is picked, irrespective of its oper-status, as long as the SAP is still created. Therefore, if the SAP is down, the frames do 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 do 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 unidirectional (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 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. When 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: GRE service distribution point (SDP) pointing from ALA-A to ALA-B) must be specified in the service creation process to "bind" the service to the tunnel (the SDP). Otherwise, service traffic is not directed to a far-end point and the far-end devices 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

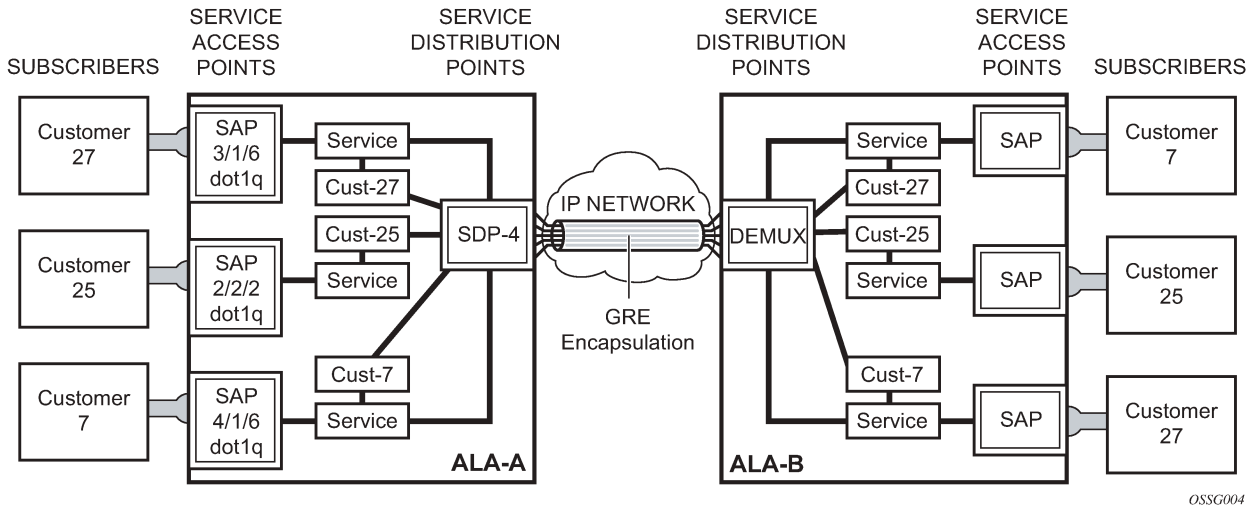

2.3.4.2 Spoke and mesh SDPs

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 because of 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 specific 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 is immediately brought operationally down.

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

When a response is received that indicates the error has cleared and the hold down time interval has expired, the SDP is eligible to be put into the operationally up state. If no other condition prevents the operational change, the SDP enters 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 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 is enforced.

Any changes made to the admin group **sdp-include** and **sdp-exclude** constraints are only 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 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 sends 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 is 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 are 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 is not 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

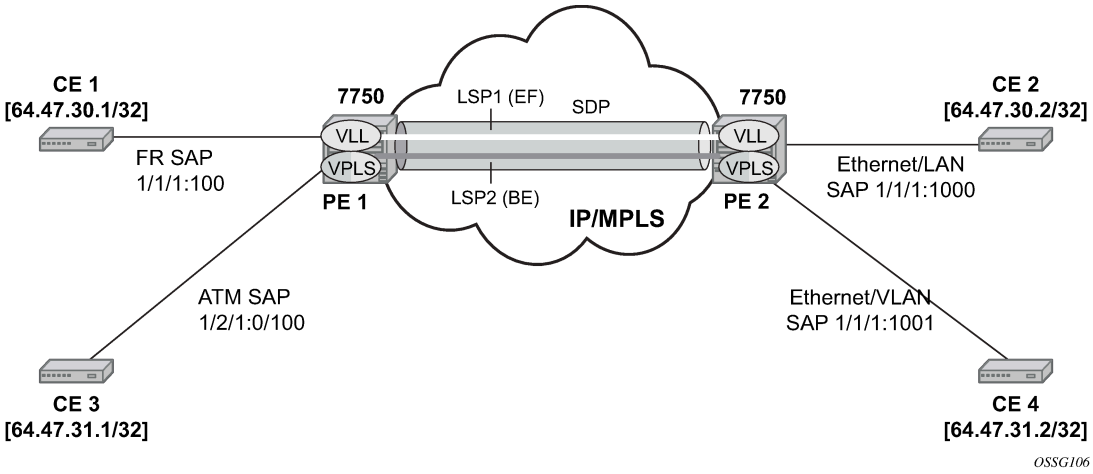

Figure 12: Class-based forwarding over SDP LSPs illustrates the use of class-based forwarding to direct packets of a service to specific RSVP or static LSPs that are part of the same SDP based on the packets' forwarding class. The forwarding class of the packet is the one assigned to the packet as a 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 for 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 supports. A forwarding class can only be assigned to one LSP in a specific SDP, meaning that only one LSP within an SDP supports a specific class of service. However, multiple classes of services can be assigned to an LSP. Both RSVP and static LSPs are allowed. All subclasses are 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 leaves the SR on an SDP that is configured for class-based forwarding, the outgoing LSP is 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 are forwarded to the LSP which is the result of hashing the VLL service ID. Because 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 is 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, 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:

config>**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 **config**>**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 because 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

  (**config**>**service**>**vprn**>**interface**>**spoke-sdp**)
- VPLS with provisioned a 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 a 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:

config>**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. After the new value is configured, the system address is not 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 toward the maximum of 15 distinct address values per system that are used by all GRE SDPs under the **config**>**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 **config**>**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 described 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 uses one of these sessions, or auto-creates 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: Mismatched T-LDP control plane parameters results in both GRE SDP1 and SDP2 to remain down because the targeted Hello adjacency and LDP session does not come up between the two LDP LSRs.

Figure 13: Mismatched T-LDP control plane parameters

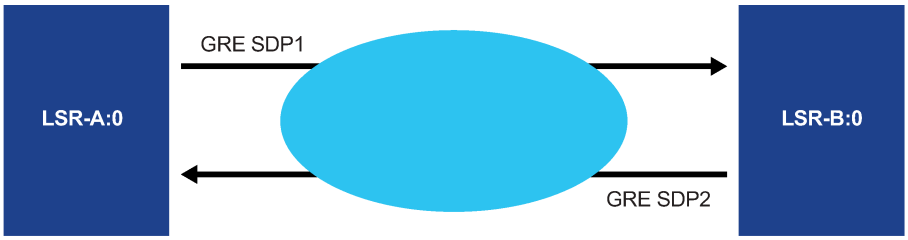

The user must match the local transport address of the T-LDP session to the local-end address of the GRE SDP in both the local and remote PE routers. This can be achieved by manually configuring a T-LDP session to the peer, or by auto-creating a T-LDP session with the targeted peer template feature, 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: Proper setting of T-LDP control plane parameters shows an example of a correct configuration.

Figure 14: Proper setting of T-LDP control plane parameters

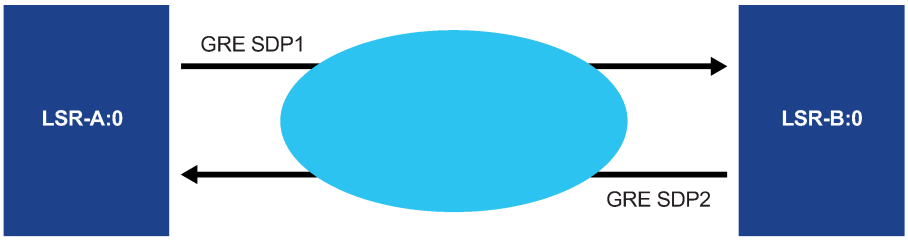

GRE SDP1:
• Local-end = A1
• Far-end = B1
T-LDP session (manually provisioned or auto-created via peer template):
• Local transport address = A1
• Peer transport address = B1

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 still come up. For example, the setup in Figure 15: Source address mismatch between control and data planes 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

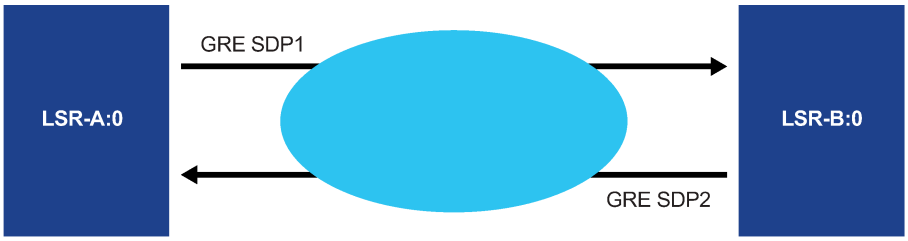

GRE SDP1:
• Local-end = A1
• Far-end = system2
T-LDP session (SDP auto-created):
• Local transport address = system1
• Peer transport address = system2

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.

2.3.4.9 GRE SDP tunnel fragmentation and reassembly

GRE SDP tunnel fragmentation and reassembly allows those services for which fragmentation is typically not available to make use of IP layer fragmentation and reassembly of GRE SDP tunnels that carry those services. It enables services that require larger MTUs to be carried over network segments where the MTU is less than the MTU of their respective GRE SDP tunnel packets. As a result, GRE SDP tunnel packets that require fragmentation are either fragmented at the source node or within the MTU restricted network, and reassembled on the GRE SDP terminating node.

GRE SDP tunnel fragmentation is supported on the following platforms:

- VSR-I
- VSR-a

2.3.4.9.1 GRE SDP tunnel fragmentation

The **allow-fragmentation** command enables GRE SDP tunnel fragmentation and can be configured under the following:

- SDP with type GRE (MPLS is not supported)
- PW template with **auto-gre-sdp**

The services that are supported with GRE SDP fragmentation and reassembly include the following:

- Epipe VLL services using GRE SDP tunnels
- BGP-VPLS
- BGP-VPWS

When enabled, GRE SDP tunnel fragmentation is applied to any generated GRE SDP tunnel packet where its size is greater than the MTU of the network interface needed to egress the node. The network interface chosen to egress the node is based on the SDP bindings for the service.

For GRE SDPs and SDP bindings using PW templates with **auto-gre-sdp**, if a received SAP packet size is greater than the service MTU, the packet is dropped per normal SAP ingress processing.

If GRE-SDP tunnel fragmentation is enabled when configuring **allow-fragmentation**, the DF-bit is cleared for unfragmented and fragmented GRE SDP tunnel packets. This allows downstream routers to fragment the packets further if needed.

2.3.4.9.2 GRE SDP tunnel reassembly

To reassemble GRE SDP tunnel fragments on the node terminating the service carried in the GRE SDP tunnel, the BB-ISA application is used to receive fragments, reassemble them, and return the reassembled GRE SDP tunnel packet for regular ingress processing.

To enable reassembly of GRE SDP fragmented packets, the following items must be configured:

1. Enable the BB-ISA application on the node (see the *7450 ESS, 7750 SR, and VSR Multiservice ISA and ESA Guide* for more information).
2. Configure an IP filter, using **config**>**filter**>**ip-filter** with:
   - **default-action forward**
   - **match protocol gre** and **fragment** set to true
   - **action** set to **reassemble**
3. Configure a NAT group under **config**>**isa**>**nat-group** with the **active-mda-limit** set to 1 and referencing the **mda** where the BB-ISA is configured.
4. Add the filter to the router interface or interfaces where GRE SDP packets and fragments are expected to be received.
5. Enable reassembly to base routing by configuring **config**>**router**>**reassembly group** *nat-group* **to-base-network**.

For example:

```
#--bb-isa-----------------------------------------------
config
    card 1
        mda 2
            mda-type isa-bb-v
            no shutdown
        exit
        no shutdown
    exit
#--ip-filter-------------------------------------------------
    filter
        ip-filter 10 create
            default-action forward
            entry 1 create
                match protocol gre
                    fragment true
                exit
                action
                    reassemble
                exit
            exit
        exit
    exit
#--nat-group-------------------------------------------------
    isa
        nat-group 1 create
            active-mda-limit 1
            mda 1/2
            no shutdown
        exit
    exit
#--router-base-----------------------------------------------
    router Base
        interface "system"
```

```
            address 91.91.91.91/32
            no shutdown
        exit
        interface "to-PE2"
            address 61.61.61.61/24 gre-termination
            port 1/1/11:123
            ingress
               filter ip 10
            exit
            no shutdown
        exit
    exit
exit
```

2.3.4.9.3 NGE considerations

GRE SDP tunnel fragmentation and reassembly is supported in conjunction with NGE. Encryption is applied to the GRE SDP (NGE) tunnel packet before fragmentation on egress and not to each fragment. Decryption occurs after the GRE SDP (NGE) tunnel packet is reassembled by the BB-ISA application and returned to base routing for processing.

GRE SDP tunnel reassembly is supported when the GRE SDP tunnel terminates on the router interface IP address. See GRE SDP termination on router interface IP address for information about GRE SDP termination on a router interface IP address.

2.3.4.10 GRE SDP termination on router interface IP address

An operator can choose to terminate GRE SDP tunnel packets directly on a router interface IP address. This is enabled by using the **config router interface address** <*IP address*> **gre-termination** command and is supported on the following platforms:

- VSR-I
- VSR-a

Services that support GRE SDP termination on the router interface IP address include the following:

- VPRN services (spoke SDP and auto-bind-tunnel)
- T-LDP signaled Ethernet VLL services
- T-LDP signaled VPLS services
- BGP-VPLS services
- BGP-VPWS services

Only GRE SDP tunnel packets that have a destination IP that is equal to the /31 value of the IP address configured on the router interface can terminate on the router interface. For example, if the address is set to 1.2.3.4/24, only the single /31 address 1.2.3.4 is used for the GRE SDP tunnel packet to terminate on. GRE SDP termination on router interface IP addresses is not supported on loopback interfaces.

When T-LDP is used to signal services, the **local-lsr-id** must be set to the /31 value of the IP address of the router interface used for the GRE SDP tunnel packets to terminate on.

When BGP is used to advertise routes for services, the **local-address** for BGP sessions must be set to the /31 value of the IP address of the router interface used for the GRE-SDP tunnel packets to terminate on.

2.3.5 SAP and MPLS binding loopback with MAC swap

SAPs and MPLS SDP bindings within Ethernet services, Epipe, and VPLS can be placed into a loopback mode, which allows packets to be looped back toward the source of the traffic. The feature 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, PBB Epipe, VPLS, and I-VPLS service constructs support both ingress and egress loopbacks on Ethernet SAPs or MPLS SDP bindings.

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

An ingress loopback configured on the entity has the following effects on traffic forwarding on the entity:

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

Essentially an ingress loopback function isolates the SAP or MPLS SDP binding from the rest of the service. Figure 16: Ingress loopback packet processing uses a simple Epipe service example to show the various touch points for a packet that is processed by an ingress loopback as it moves through the network element.

Figure 16: Ingress loopback packet processing

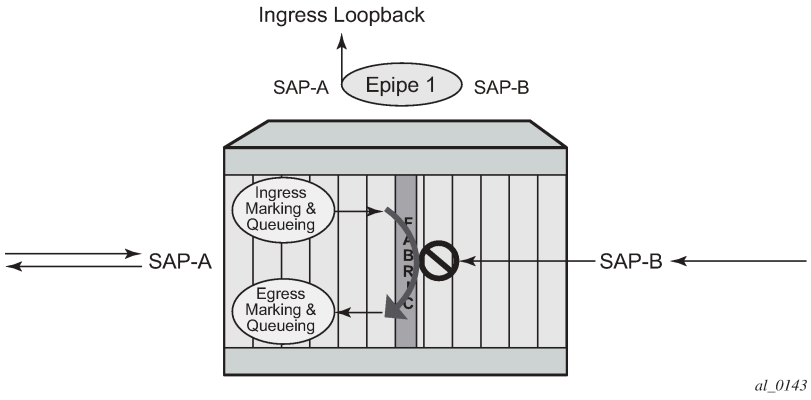

An egress loopback configured on the entity has the following effects on the traffic forwarding on the entity:

- Traffic arriving on any service SAP or SDP binding that is forwarded to an egress loopback is looped back into the service.
- Traffic attempting to gain access to the service from that entity (ingress the network element from the entity) is 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.

Note: Extreme care must be used when considering the application of an egress loopback in a VPLS or I-VPLS service. Because a VPLS service relies on MAC based forwarding, any

packet that arrives at an egress loopback is reflected back into the service, which uses MAC based forwarding to apply the correct forwarding decision. In a live multipoint service with active endpoints this could have a major negative impact on the service and the clients connected to this service. Even if the forwarding database is primed, any arriving broadcast, unknown or multicast traffic arrives on the egress loopback and is reflected back into the service causing (at the very least) duplication of this type of traffic in the service.

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

Figure 17: Egress loopback packet processing

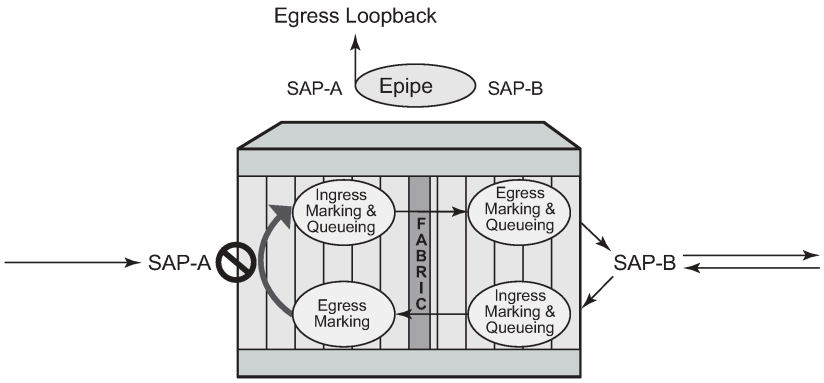

The operational state of the SAP or MPLS SDP binding does not change as a result of the loopback function. This means a SAP or MPLS SDP binding that is operationally up does not change state strictly because the loopback started or stopped. Of course control protocols that are attempting to gain access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

Exercise caution when considering the use of control protocols in a service with enabled loopbacks. The operator must understand that control protocol interruptions can significantly impact 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, interrupting the control protocol causes the SAP to fail. Other SAPs linking their state to a failed SAP react to that failure as well. This loopback function is per Ethernet SAP or MPLS SDP binding. That is, all traffic that is extracted and sent to the CPM before the loopback process is looped back 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 blocked. However, there may be instances where it is essential to continue running 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 support various MAC swap functions. Table 6: MAC-Swap configuration and options lists the actions and functions based on the configured **mac-swap** and associated options.

Table 6: 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 member ports that are administratively up. If any of these conditions are not met, the loopback function fails.

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 enters an operationally down state when one of the SAPs is non-operational. The service state remains or is returned to an operational state if the **ignore-oper-down** command is configured under the non-operational SAP. A VPLS service remains 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 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 never arrives on the SAP to be looped.

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

Use the *tools* hierarchy to configure this functionality. 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 survives 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 loopback function does 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 and functions that use mirroring are mutually exclusive.

Figure 18: Active loopback mode shows of 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

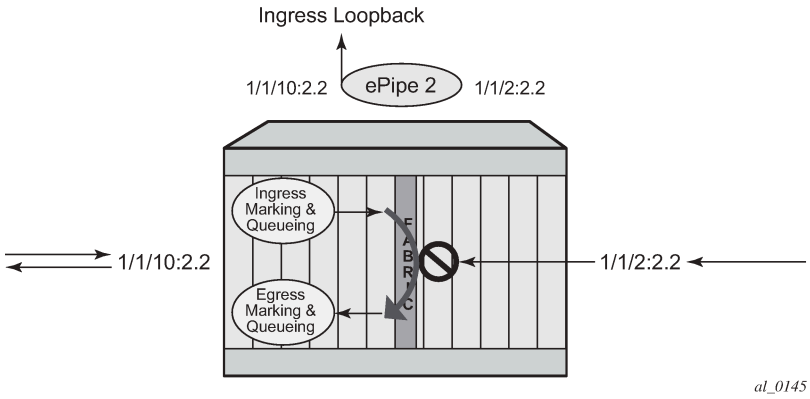

The following is an example output of the active loopback mode based on Figure 18: Active loopback mode.

```
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 operates 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 does 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, are not incremented. CFM debugging is not 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 access the additional resources required to accommodate high-speed service activation processing. Normal processing of ETH-CFM packets occurs for all other ETH-CFM PDUs that arrive on the MEP with the same domain level. The MEP also processes and terminates the lower levels as per normal processing. To ensure correct 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.

The ETH-CFM level of the high-speed ETH-LBM stream must match the level of a MEP configured with the **lbm-svc-act-responder** command. It must not target any lower ETH-CFM level that the MEP terminates. When the service activation test is complete, the MEP can be returned to standard processing by removing this command. If there is available bandwidth, the MEP responds to other ETH-CFM PDUs, such as ETH-DMM marker packets, using standard processing.

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.

There is interaction between the **lbm-svc-act-responder** command and the **tools perform service id loopback eth** command. Nokia recommends that either the **lbm-svc act-responder** or the **tools perform service id loopback eth** command be used at any time within a service. If both commands must be configured, and the target reflection point is the MAC Swap Loopback function, the inbound stream of data must not include ETH-CFM traffic that is equal to or lower than the domain level of any configured MEP which would otherwise extract and process the ETH-CFM message.

If the reflection target is a MEP configured with the **lbm-svc-act-responder** command, the mode (ingress or egress) of the SAP or SDP specified with the tools command and the MEP direction (up or down) must match when the functions are enabled on the same reflection point. The domain level of the inbound ETH-LBM must be the same as that of the MEP configured with **lbm-svc-act-responder**. At no time should the two functions be conflicting with each other along the path of the stream. Such a conflict can lead to unpredictable and possibly destabilizing situations.

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 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 because 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 because 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 toward the customer site, or within the Ethernet backbone. E-Line, 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

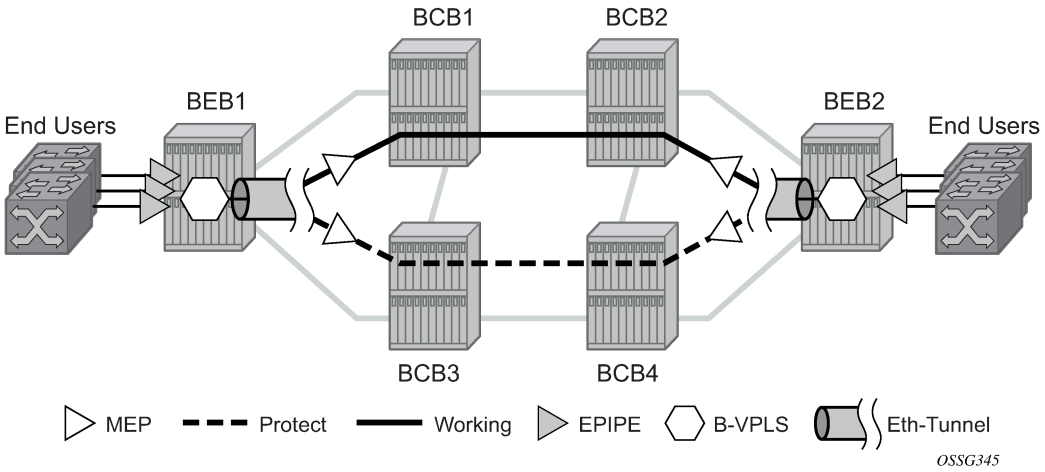

In the implementation, the Ethernet tunnel is a logical interface for a SAP defined Layer 2 service similar to a LAG. The implementation offers ITU G.8031 1:1 compliance as well as some added capabilities such as fate sharing and emulated LAG support. Other added capabilities include:

- synchronization between services such that both send and receive on the same Ethernet path in stable state
- revertive/non-revertive choices
- emulated-LAG coexistence

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 fail to receive the CC messages exchanged on that path or 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 declares the associated path down, then sends an indication to the remote protection module to switch the transmit direction to the backup path.

To address unidirectional failures, the RDI bit is 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 continues 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* 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

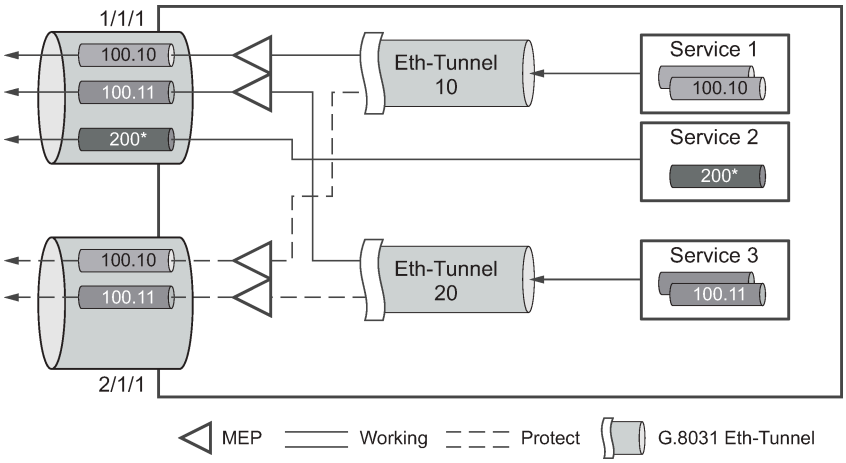

Figure 21: PBB fate sharing Ethernet tunnels

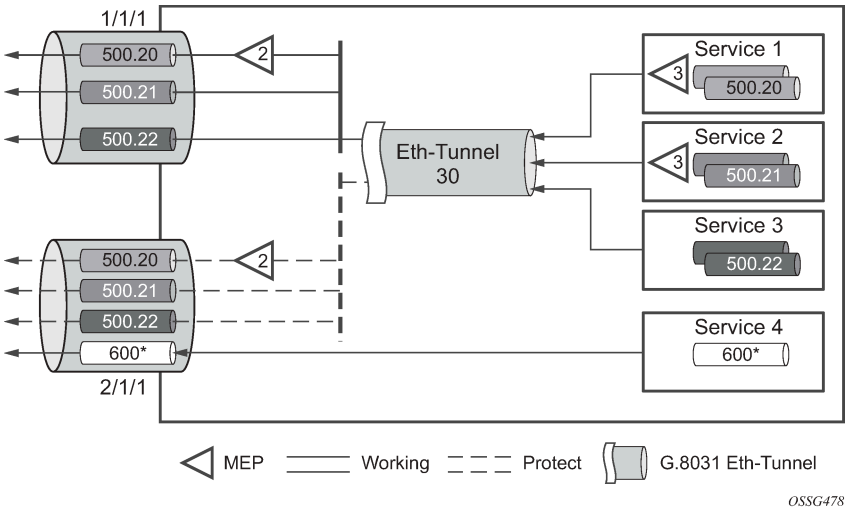

One of the advantages of access redundancy using Ethernet APS is that because it operates at the VLAN level protection mechanisms can be varied between services supported on the physical port. For example, it is possible to provide a protected service for “Premium” customers and also provide non-protected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

Ethernet tunnel CC messages transmitted over the SAP queues using the default egress QoS policy use NC (network class) as a forwarding class. If user traffic is assigned to the NC forwarding class, it competes 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 links, 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 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. Because of 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 Loss Of Signal (LOS) 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. When 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 is 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 blocked 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: 0-1 G.8032 ring in the initial state depicts this operational state:

Figure 22: 0-1 G.8032 ring in the initial state

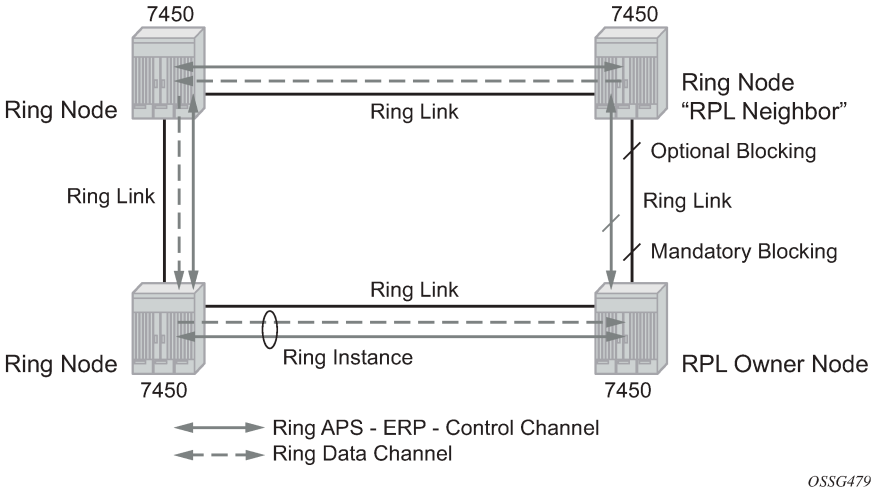

When a ring failure occurs, a node or nodes detecting the failure (enabled by Y.1731 OAM CC monitoring) send R-APS message in both directions. This allows the nodes at both ends of the failed link to block forwarding to the failed link preventing it from becoming active. In revertive mode, the RPL Owner 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: 0-1 G.8032 ring in the protecting state depicts the failed link scenario.

Figure 23: 0-1 G.8032 ring in the protecting state

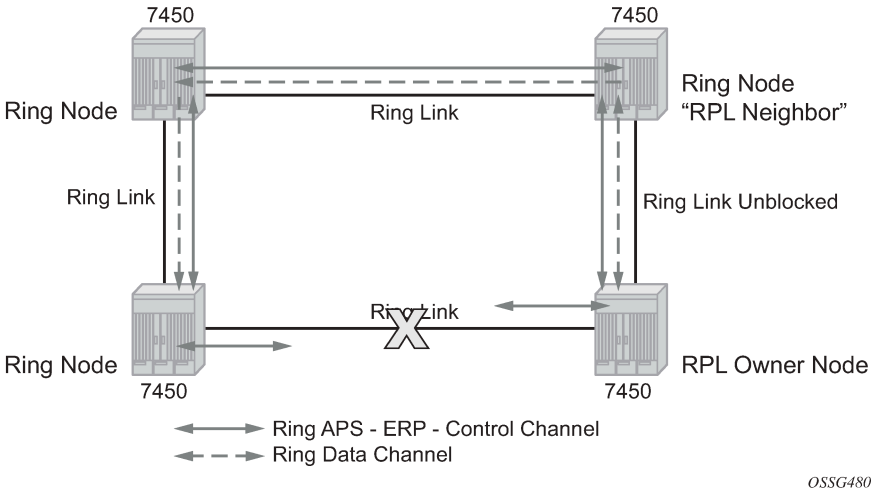

When the failed link recovers, the nodes that blocked the link again send the R-APS messages indicating no failure this time. This in turn triggers RPL owner to block the RPL link and indicate the blocked RPL link the ring in R-APS message, which when received by the nodes at the recovered link cause 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, reconvergence 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, after 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 toward 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 needed. 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 fail to receive the CCMs exchanged on that segment or 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 needed 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 declares the associated ring link down and the G.8032 state machine sends 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: 0-3 ring example 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

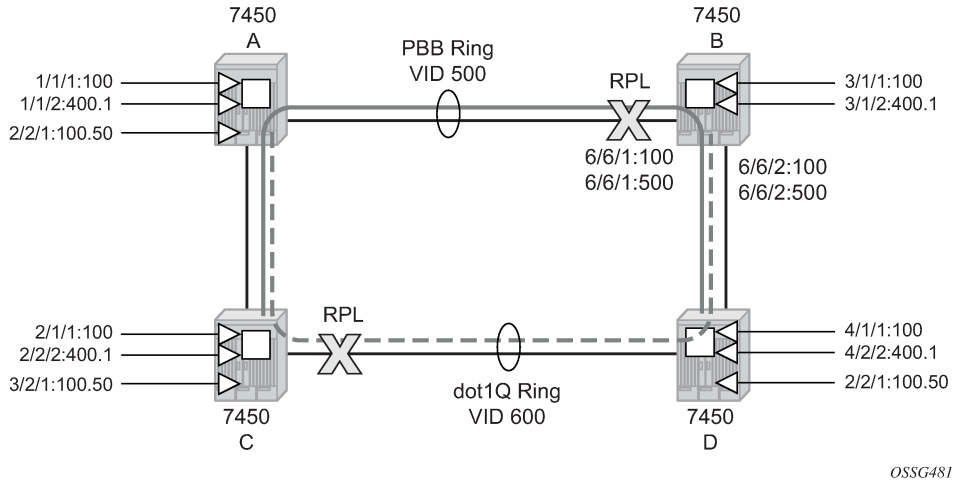

Example configuration:

```
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: 0-4 G.8032 sub-ring 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

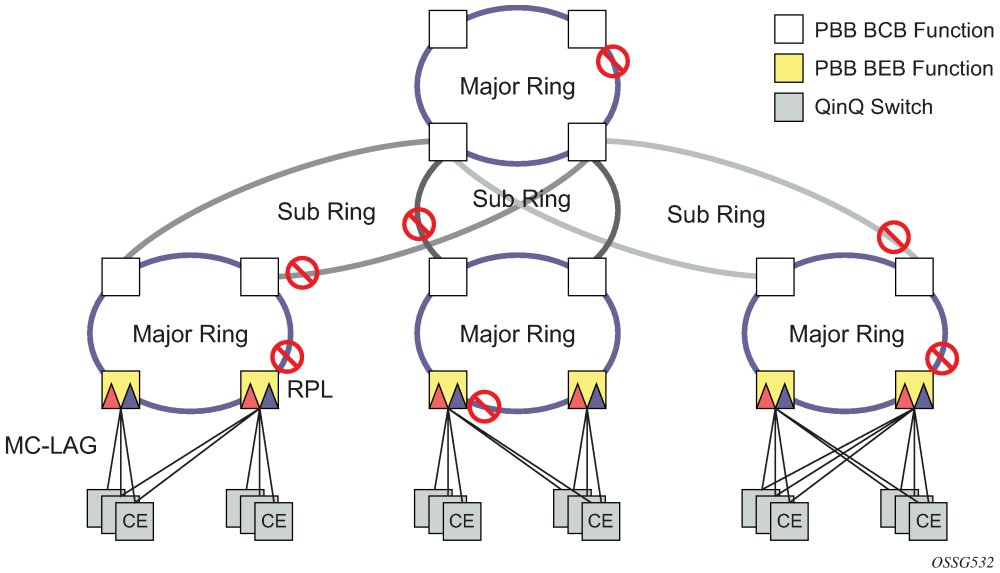

Sub-Rings and Major Rings run similar state machines for the ring logic, however there are some differences. When Sub-Rings protect a link, the flush messages are propagated to the major ring. (A special configuration allows control of this option on the router.) When major rings change topology, the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

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 because 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

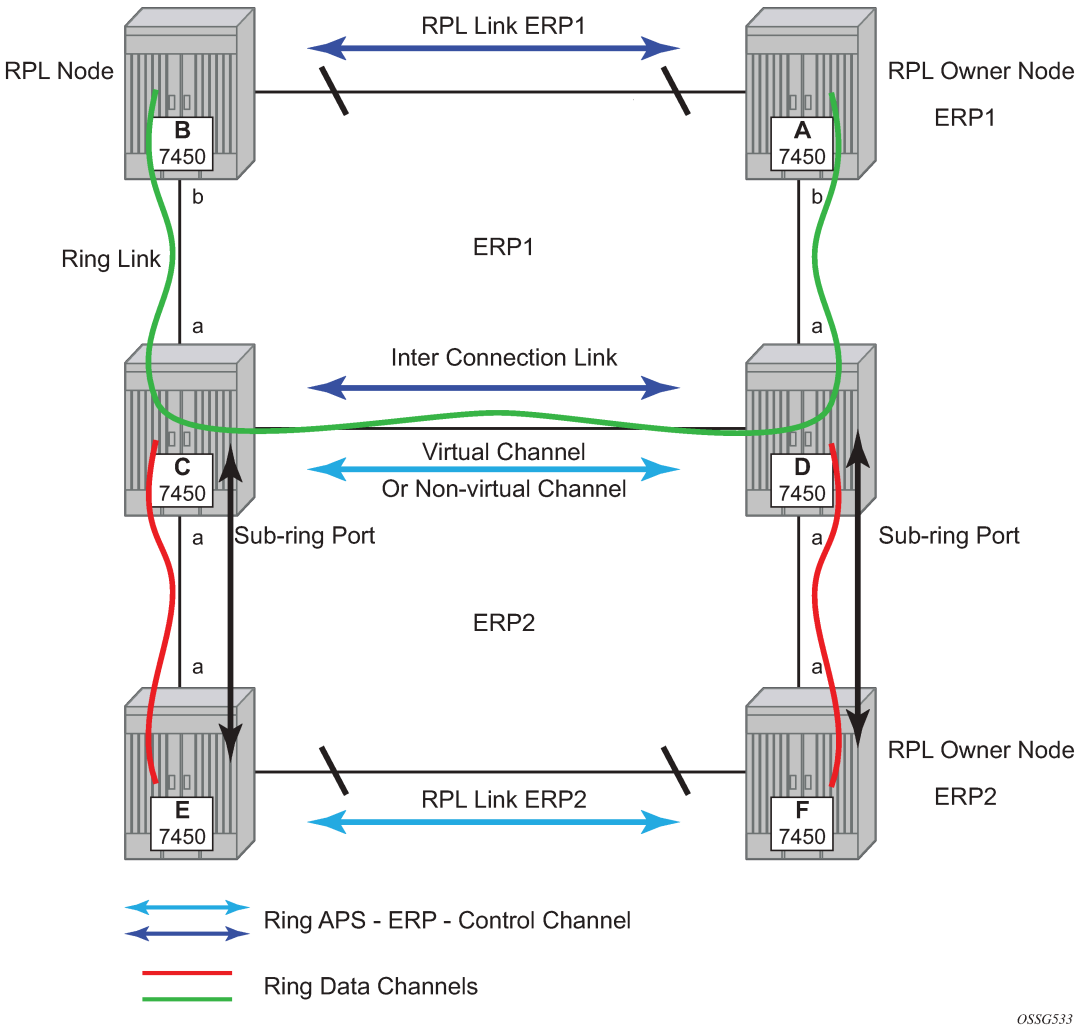

Sub-ring configuration is similar to major ring configuration and consists of three parts: Ethernet-ring instance configuration, control VPLS configuration, and data VPLS configuration (data instance or data channel). The Ethernet-ring configuration of a sub-ring is tied to a major ring and only one 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 needed 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 an example 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: 0-6 sub-ring homed to VPLS. 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

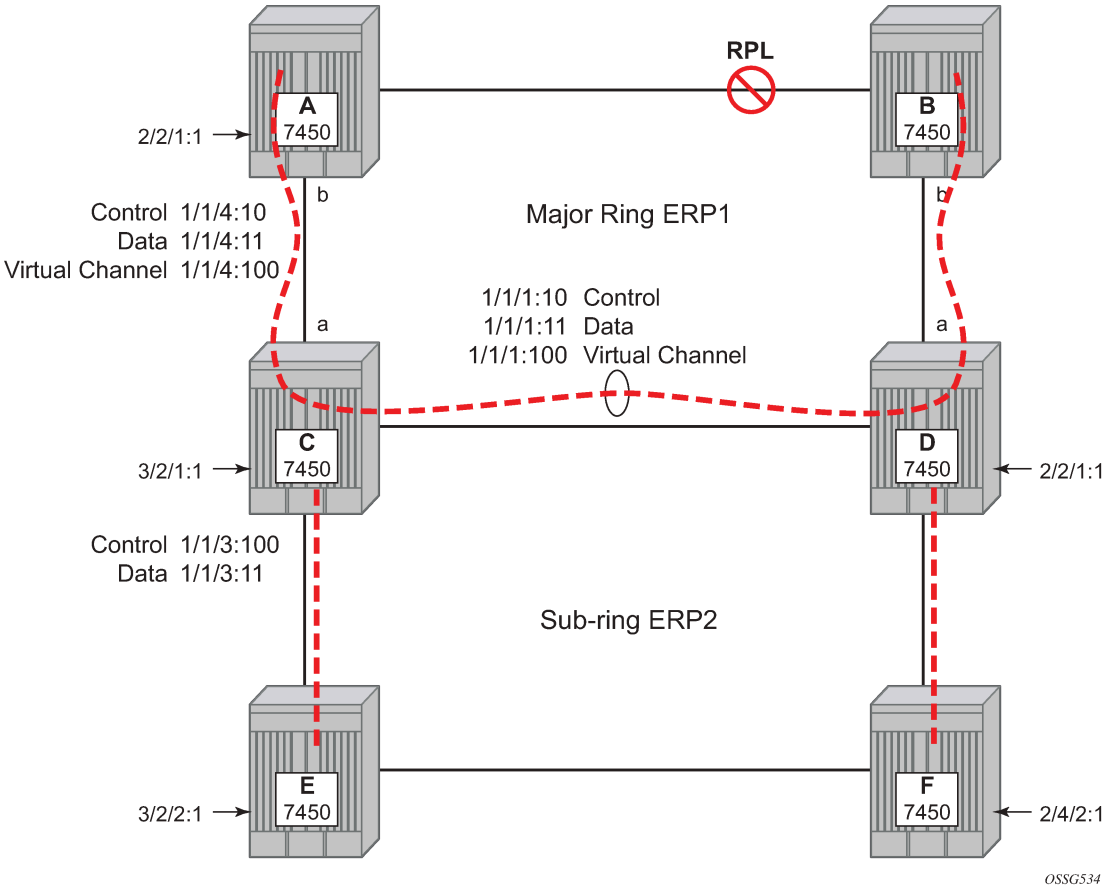

The following is an example sub-ring configuration for VPLS (at PE1):

```
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: 0-7 multi ring hierarchy 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

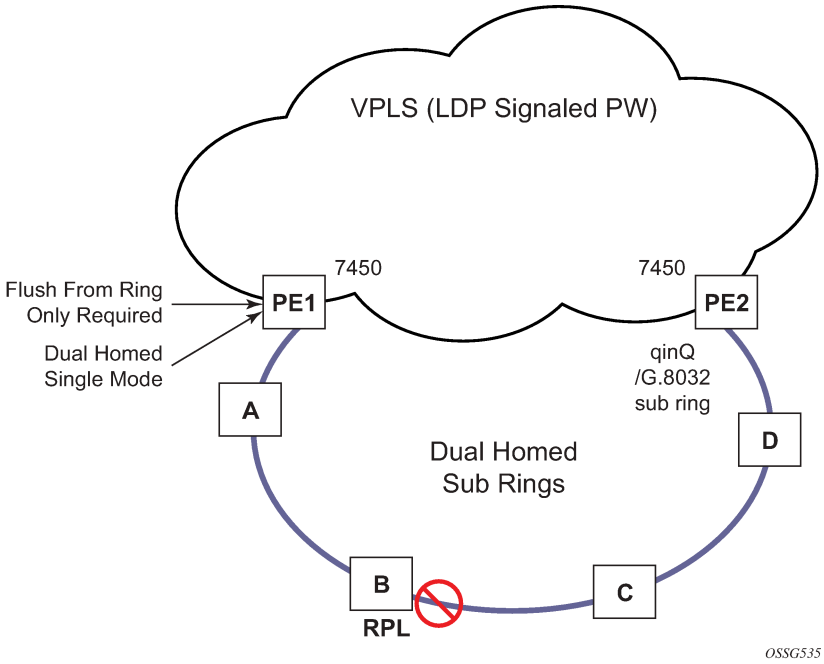

The major rings are connected by sub-rings to the top level major ring. These sub-rings require virtual channel and do not work with non-virtual channel. Ring flushing is contained to major rings, or in the case of a sub-ring link or node failure, to the sub-ring and the directly attached major rings.

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 specific 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 path 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 are created with the actual parameters on each card.

Ethernet ring CC messages transmitted over the SAP queues using the default egress QoS policy use NC (network class) as a forwarding class. If user traffic is assigned to the NC forwarding class, it competes

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*.

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:

- appear in CLI show commands and MIB walks output.
- appear in the output of **info detail** commands but are never 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 is never 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 is automatically 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. See 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 (including LDP over RSVP) and SR-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 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: NGE network with NSP NFM-P management shows an NGE network with NSP NFM-P services, control plane configuration, and key management.

Figure 29: NGE network with NSP NFM-P management

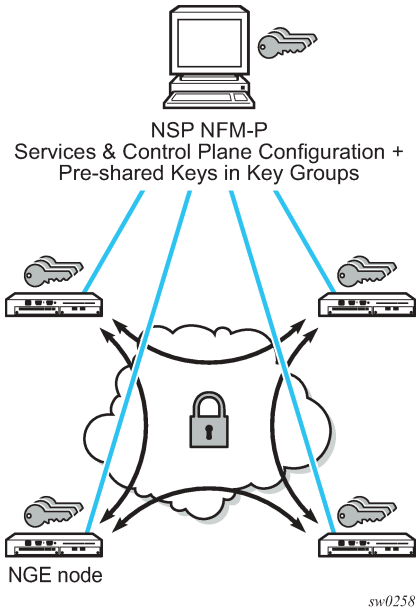

NGE provides five main types of encryption to secure an IP/MPLS network:

- **SDP encryption**

  This is 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**
  - **unicast VPRN**

    This is MP-BGP-based VPRN-level encryption using auto-bind of LDP, GRE, RSVP-TE, MPLS (LDP or RSVP-TE), or segment routing (SR-ISIS, SR-OSPF, and SR-TE) tunnels.

  - **multicast VPRN**

    NG-mVPN using mLDP with auto-discovery

- **router interface**

  This is Layer 3 user plane and control plane encryption.
- **WLAN-GW group interface**

  This is L2oMPLSoGRE level encryption from WLAN access points (APs) that support NGE.
- **PW template encryption**

  This is BGP-VPLS- and BGP-VPWS-based MPLS services encryption, which uses a PW template with **auto-gre-sdp** configured.

Note: See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide* for information about configuring NGE on router interfaces. See the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide* for information about configuring group encryption on the WLAN-GW group interface.

NGE is supported on the following platforms:

- VSR-I
- VSR-a

WLAN-GW group interfaces enabled with NGE is further supported on the following platforms:

- 7750 SR-7
- 7750 SR-12
- 7750 SR-12e
- 7750 SR-1e
- 7750 SR-2e
- 7750 SR-3e

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: Key group partitioning 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

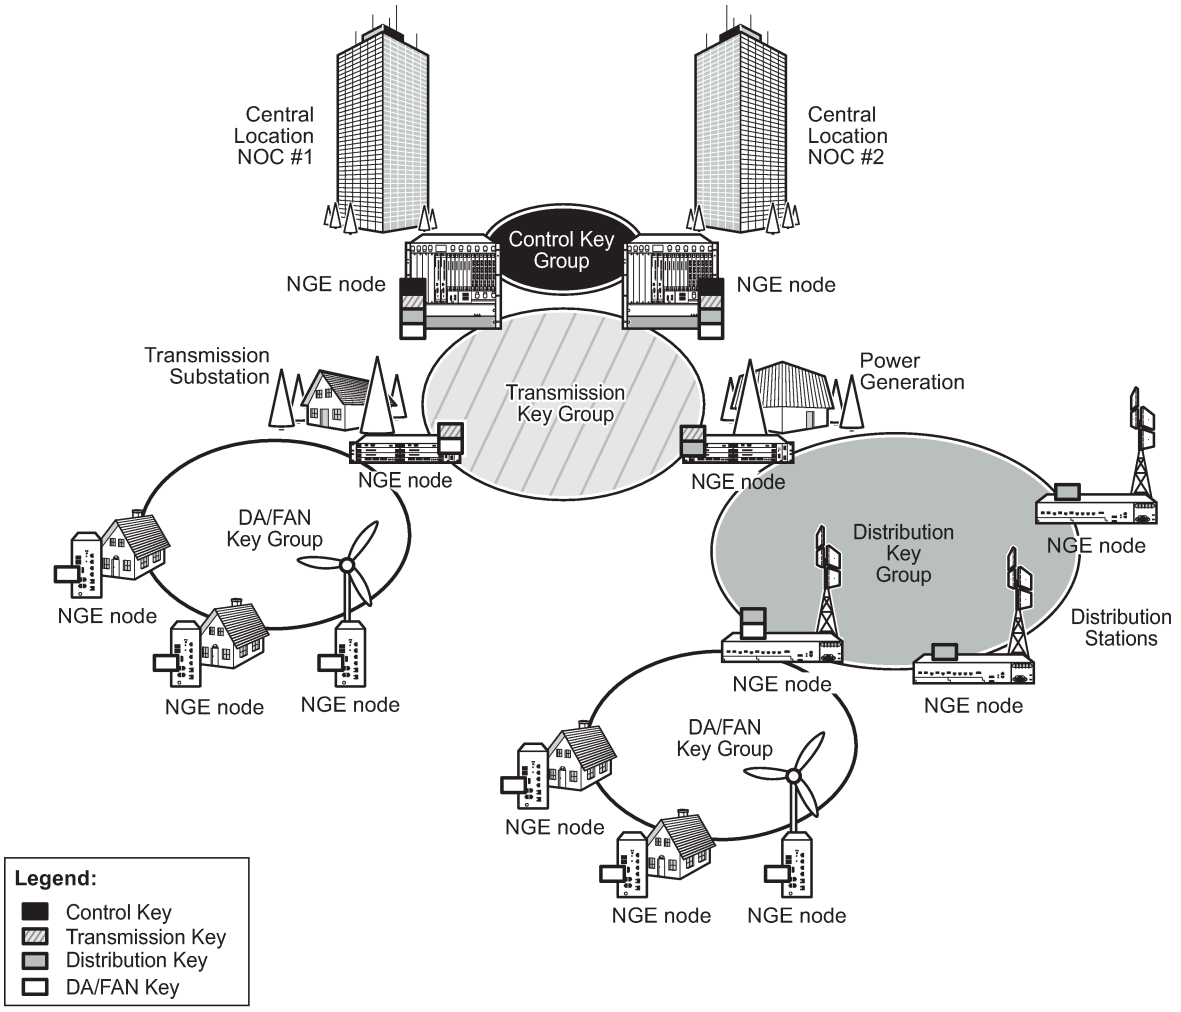

Another application of key group partitioning allows different parts of an organization to have their own method of end-to-end communication without the need to share encryption keys between each organization. If two partitions need to communicate between themselves, gateway nodes configured with both 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 needed, or all key groups can be rekeyed network-wide at the same time.

For more information about NSP NFM-P management, see the "Service Management" section in the *NSP NFM-P User Guide*.

2.10.2 Key groups

Users can partition the network based on security requirements by organizing encryption keys into distinct key groups. 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: Key groups and a typical NGE packet illustrates the use of key groups (KGs), SAs, and security parameter indexes (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. KG1 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

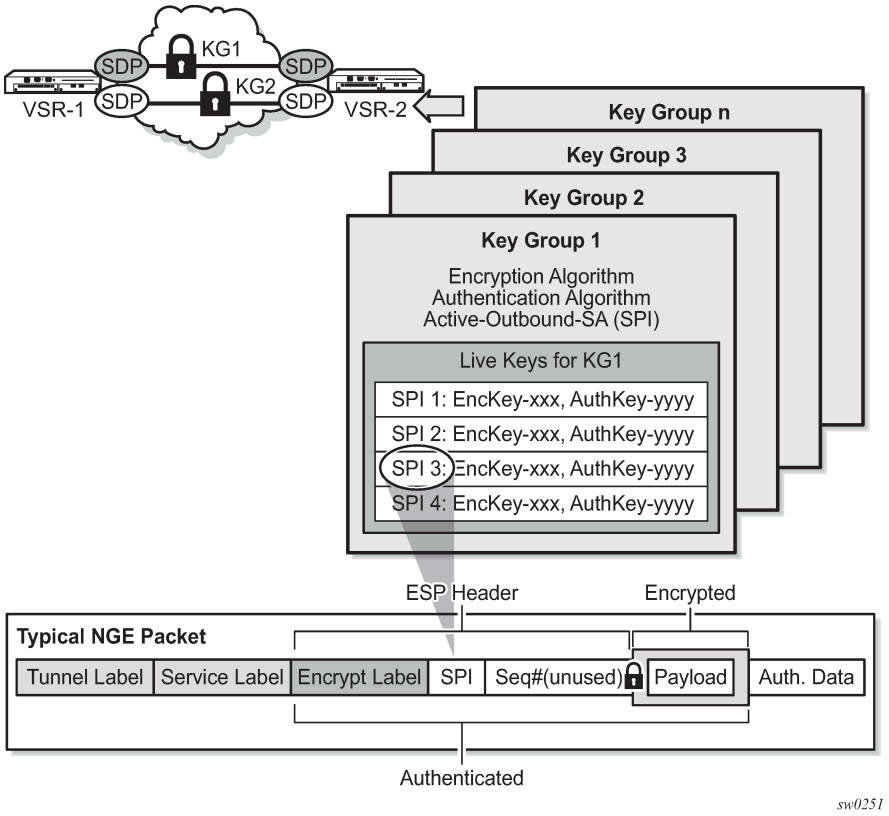

2.10.2.1 Key group algorithms

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: Key groups and a typical NGE packet illustrates these relationships.

As shown in Figure 31: Key groups and a typical NGE packet , 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
- L2oMPLSoGRE VLL service terminating on a WLAN-GW group interface
- BGP-VPLS and BGP-VPWS using a PW template with **auto-gre-sdp** configured
- NG-MVPN

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, resolving routes using spoke SDPs (**spoke-sdp**) or auto-bind SDPs (**auto-bind-tunnel**) is supported using LDP, GRE, RSVP-TE, or segment routing (SR-ISIS, SR-OSPF, or SR-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: NGE MPLS/GRE/MPLSoUDP label stack illustrates the NGE MPLS, GRE, or MPLSoUDP label stack.

Figure 32: NGE MPLS/GRE/MPLSoUDP label stack

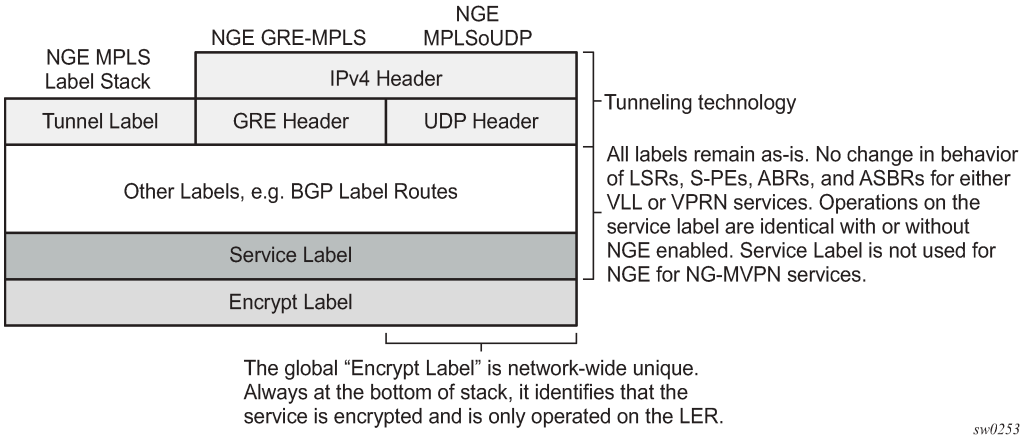

Figure 33: NGE and packet formats illustrates VPRN and PW (with control word) packet formats using NGE.

Figure 33: NGE and packet formats

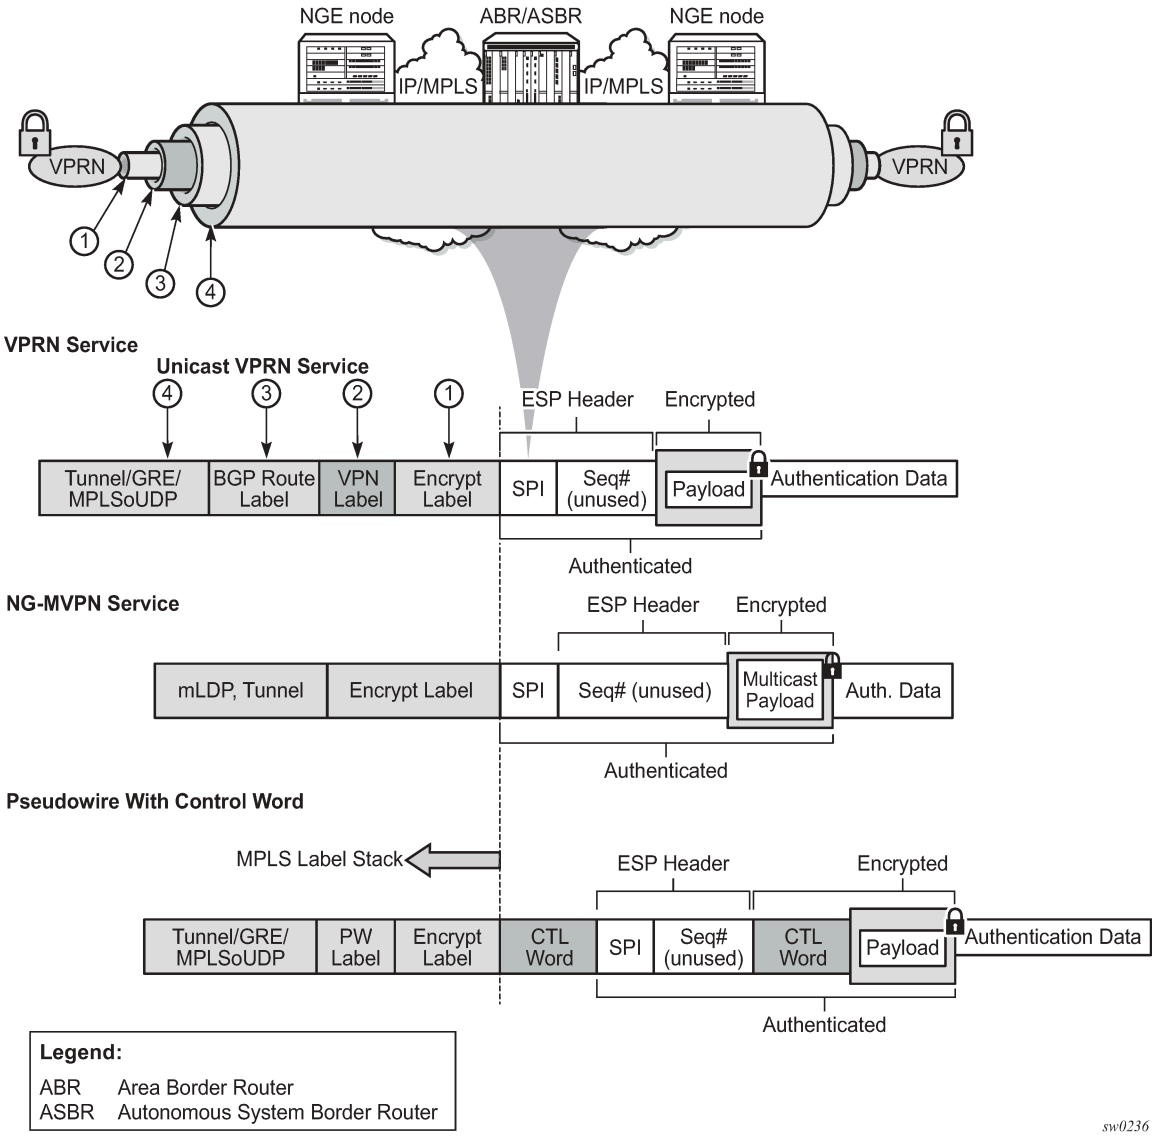

2.10.3.2 Assigning key groups to services

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

Figure 34: Inbound and outbound key group assignments

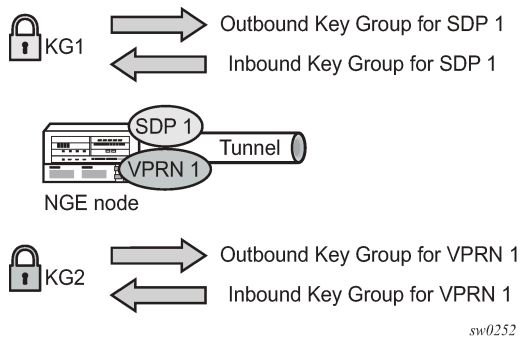

The outbound key group identifies which key group to use for traffic that egresses the node for the service. The inbound key group ensures that ingress traffic is using the correct key group for the service.

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. See the *NSP NFM-P User Guide* for more information.

2.10.3.3 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 (either with spoke SDPs using **spoke-sdp** or auto-bind SDPs using **auto-bind-tunnel**) 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**.

2.10.3.4 L2 Service encryption using PW templates

NGE **encryption-keygroup** configuration is supported on PW templates to enable the encryption of MPLS services that are based on BGP-VPLS and BGP-VPWS and that have **auto-gre-sdp** enabled on the PW template. All services configured using the PW template that have both NGE and **auto-gre-sdp** enabled are encrypted.

When changing the **encryption-keygroup** on a PW template, the change does not take effect immediately. The operator must execute the following command after each change to the **encryption-keygroup** for it to take effect:

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

2.10.3.5 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).

See "Pseudowire Switching" in the *7450 ESS, 7750 SR, 7950 XRS, and VSR OAM and Diagnostics Guide* for more information.

2.10.3.6 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.

2.10.3.7 WLAN-GW encryption

NGE is supported on the WLAN-GW to provide encryption of traffic to and from WLAN APs that support NGE, such as the 7705 SAR-Hm. As shown in Figure 35: Terminating NGE-protected WLAN AP traffic destined for the WLAN-GW, the application uses Epipe pseudowire services, as described in Pseudowire switching for NGE traffic and Pseudowire control word for NGE traffic, with L2oMPLSoGRE transport and NGE applied to the GRE-SDP.

Figure 35: Terminating NGE-protected WLAN AP traffic destined for the WLAN-GW

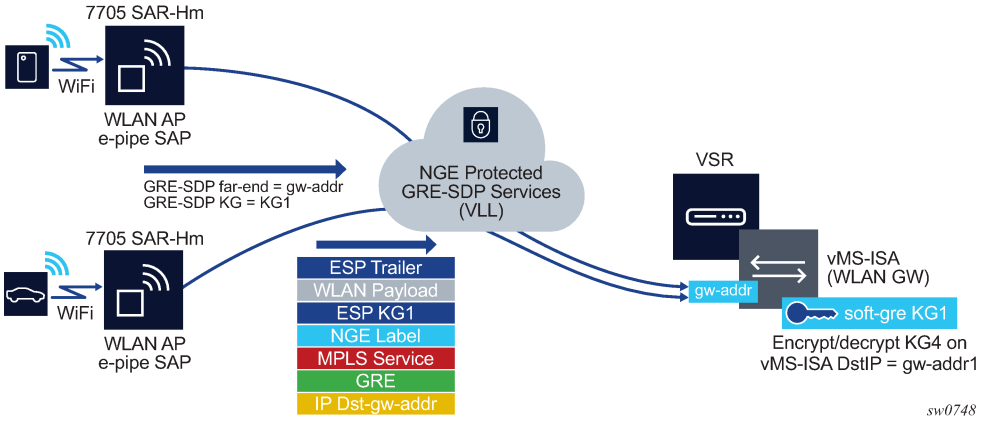

In Figure 35: Terminating NGE-protected WLAN AP traffic destined for the WLAN-GW, the same key group, KG1, is configured on:

- the WLAN-GW gw-address under the IES or VPRN soft-GRE group interface for the vMS-ISA
- the GRE SDPs that are bound to any WLAN AP Epipes that are terminating soft-GRE services on the WLAN-GW group interface

Traffic from an authenticated user on the SAR-Hm WLAN AP is encrypted and an NGE label is added to the packet after the Epipe service label. The packet format is shown in Figure 35: Terminating NGE-protected WLAN AP traffic destined for the WLAN-GW.

The WLAN-GW group interface is configured with the same inbound and outbound key group as the GRE-SDP used for the Epipe from the WLAN AP. Any L2oMPLSoGRE packet received by the WLAN-GW on the NGE-enabled group interface must be encrypted with NGE per the above format. All other supported WLAN-GW packet types (that is, those not using L2oMPLSoGRE) are not impacted by the NGE configuration and can pass through the WLAN-GW without NGE (such as L2oGRE packets).

2.10.3.8 NGE and RFC 8277

When RFC 8277 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).

2.10.3.9 NGE for NG-MVPN

NGE is supported for NG-MVPN services with multicast configurations that include:

- I-PMSI
- S-PMSI
- C-multicast signaling
- mLDP and RSVP-TE multicast tunnel LSPs

See the 7450 ESS, 7750 SR, 7950 XRS, and VSR Multicast Routing Protocols Guide for information about Multicast VPN (MVPN).

When R-VPLS is configured for the VPRN, the source of an NG-MVPN multicast stream can originate within a VPLS service and can be NGE encrypted before entering the I-PMSI or S-PMSI. The receiver of an NG-MVPN multicast stream can be within a VPLS service and can be NGE decrypted before being sent over the VPLS service.

When NGE is enabled on a VPRN with NG-MVPN-based services, transit nodes (LSRs) have no knowledge that NGE is being employed, nor that the NGE encryption label is being used with an ESP header after the NGE label. Features that inspect packet contents to make further decisions are not supported and must be disabled for mLDP multicast paths that need to carry NG-MVPN traffic that is NGE encrypted.

These features include:

- ingress multicast path management
- IP-based LSR hashing

The above restriction includes any 3rd party routing function that inspects the contents after the mLDP or RVSP-TE transport label and expects a non-encrypted payload on which to make hashing decisions.

2.10.4 NGE packet overhead and MTU considerations

NGE adds overhead packets to services. Table 7: NGE overhead for MPLS shows the additional overhead for the worst-case scenario of MPLS services encryption. Table 8: NGE overhead for router interface shows the additional overhead for the worst-case scenario of router interface. Additional overhead depends on which encryption and authentication algorithms are chosen.

Table 7: NGE overhead for MPLS

Item	Number of bytes
Encryption label	4
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 8: 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 7: NGE overhead for MPLS 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 described in Table 9: Accounting for NGE overhead SDP and service MTU — calculation examples for MPLS-based and GRE-based services. For details on configuring MTU, see the "MTU Configuration Guidelines" section in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide*.

The calculations in Table 9: Accounting for NGE overhead SDP and service MTU — calculation examples 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 9: 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):

Service type	MTU values with and without NGE enabled
	= 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 is 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 restrictions. For more information about configuring NGE using the NSP NFM-P, see the *NSP NFM-P User Guide*.

Consider the following restrictions 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.

2.10.7.1 Enabling NGE for an SDP or VPRN service

Procedure

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.

2.10.7.2 Enabling NGE for a router interface

Procedure

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

Step 2. Configure the outbound key group.

Step 3. Configure the inbound key group.

2.10.7.3 Changing NGE from one key group to another key group for an SDP or VPRN service

Procedure

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.

2.10.7.4 Changing NGE from one key group to another key group for a router interface

Procedure

Step 1. Remove the inbound key group.

Step 2. Configure the new outbound key group.

Step 3. Configure the new inbound key group.

2.10.7.5 Disabling NGE for an SDP or VPRN service

Procedure

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.

2.10.7.6 Disabling NGE for a router interface

Procedure

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 36: Service creation and implementation flow displays the overall process to provision core and subscriber services.

Figure 36: Service creation and implementation flow

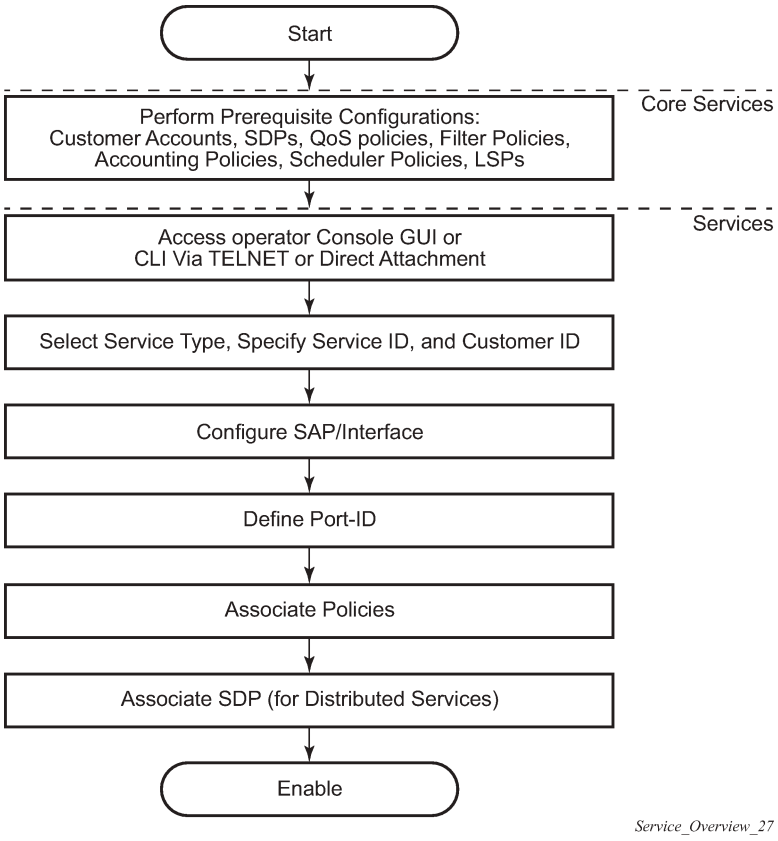

2.12 Deploying and provisioning services

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 Building the core network

About this task

The first phase of deploying and provisioning services is the core network construction.

Procedure

Step 1. Build the IP or IP/MPLS core network.

Step 2. Configure the routing protocols.

Step 3. Configure MPLS LSPs (if MPLS is used).

Step 4. Construct the core SDP service tunnel mesh for the services.

What to do next

Performing service administration

2.12.2 Performing service administration

Prerequisites

Perform the first phase for deploying and provisioning services, that is, building the core network.

About this task

The second phase of deploying and provisioning services is the service administration. The service administration includes preliminary policy and SDP configurations to control traffic flow, operator access, and to manage fault conditions and alarm messages.

Procedure

Step 1. Configure group and user access privileges.

Step 2. Build templates for QoS, filter, or accounting policies needed to support the core services.

What to do next

Provisioning services

2.12.3 Provisioning services

Prerequisites

Perform the second phase for deploying and provisioning services, that is, performing service administration.

About this task

The third phase of deploying and provisioning services is the service provisioning.

Procedure

Step 1. Provision customer account information.

Step 2. If necessary, build any customer-specific QoS, filter, or accounting policies.

Step 3. Provision the services on the service edge routers by defining SAPs, binding policies to the SAPs, and binding the service to appropriate SDPs as necessary.

See Configuring customers and Configuring an SDP.

2.13 General configuration notes

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

- core tasks
- subscriber tasks

The core tasks must be performed before the subscriber tasks.

Core tasks include the following:

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

Subscriber tasks include the following:

1. Create Cpipe, Epipe, IES, Ipipe, VPLS or VPRN services on the 7750 SR.
2. Create Epipe, IES, Ipipe, VPLS or VPRN services on the 7450 ESS or 7950 XRS.
3. Bind SDPs.
4. Configure interfaces (where required) and SAPs.
5. 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
  - **Circuit Emulation services (Cpipe)**

    See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide* for more information.
  - **Ethernet Pipe (Epipe) services**

    See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide* for more information.
  - **IP Interworking VLL (Ipipe) services**

    See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide* for more information.
  - **Virtual Private LAN Service (VPLS)**

    See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide* for more information.
  - **Internet Enhanced Service (IES)**

    See 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**

    See 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**

    See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide* for more information.
  - **VPLS SAP**

    See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide* for more information.
  - **IES SAP**

    See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for more information.
  - **VPRN Interface SAP**

    See 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 specific service when 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:

```
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.

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

```
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

About this task

This procedure provides a brief overview of the tasks that must be performed to configure a basic SPD.

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. After an SDP is created, services can be associated with 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.
- To configure an MPLS SDP, LSPs must be configured first and then the LSP-to-SDP association must be explicitly created. Auto-LSPs such as one-hop-p2p and mesh-p2p, and PCE initiated LSPs cannot be used by a configured MPLS SDP.
- 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.

Procedure

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. 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. See 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:

```
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 8277 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 fails.

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.

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 time, the service manager programs only one type of LSP in the line card that activates it to forward service packets according to the following priority order:

In the case of the RSVP or LDP SDP, the service manager programs the NHLFEs 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 re-programs 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 reverts 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 takes effect only at the next use of the timer. Any timer which is outstanding at the time of the change is restarted with the new value.

If class based forwarding is enabled for this SDP, the forwarding of the packets over the RSVP LSPs is 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 prefers the LDP LSP type over the BGP LSP type. The service manager re-programs 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 specific /32 prefix, only a single route exists in the routing table: the IGP route or the BGP route. Therefore, either the LDP FEC or the BGP label route is active at any 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 because there is no situation where both LSP types are active for the same /32 prefix.

- **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
- **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.
- **BGP LSP type**
  One RFC 8277-labeled BGP prefix programmed by the service manager. The ingress line card can use more than one next-hop for the 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 specific cases, the different functions 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 provides configuration examples for each of the functions. It also provides the various OAM command line options and show commands to operate the network. The individual service guides provide the complete CLI configuration and description of the commands to build the necessary constructs and management points.

Table 10: ETH-CFM acronym expansions lists and expands the acronyms used in this section.

Table 10: ETH-CFM acronym expansions

Acronym	Expansion	Supported platform
1DM	One way Delay Measurement (Y.1731)	All

Acronym	Expansion	Supported platform
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
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

Acronym	Expansion	Supported platform
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. Because 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 11: Defect conditions and priority settings 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 11: Defect conditions and priority settings

Defect	Low priority defect	Description	Causes	Priority
DefNone	n/a	No faults in the association.	Normal operations.	n/a
DefRDICCM	allDef	Remote Defect Indication.	Feedback mechanism to inform unidirectional faults exist. It provides the feedback loop to the node with the unidirectional failure conditions.	1
DefMACStatus (default)	macRemErrXcon	MAC Layer.	Remote MEP is indicating a remote port or interface not operational.	2
DefRemoteCCM	remErrXon	No communication from remote peer.	MEP is not receiving CCM from a configured peer. The timeout of CCM occurs at 3.5x the local CC interval. As per the specification, this value is not configurable.	3

Defect	Low priority defect	Description	Causes	Priority
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 before 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 before 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** 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**

  This 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. 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**

    This 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**

    This 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**

    This enters faulty state and further impacts 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 affected only by the **tunnel-fault** configuration option.

  – **router IP interface**

    This 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 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 is unpredictable. When an explicit endpoint is configured, service-based MEPs are required when AIS generation is the needed 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 needed 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 is terminated and processed by the port-based MEP.

Do not use MEPs configured with level 0 to validate logical transport or services. Consider 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 are 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. Because 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 because of ETH-CFM, the MEP continues to transmit and receive 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 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 after 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 37: Fault handling non-member port 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 37: Fault handling non-member port

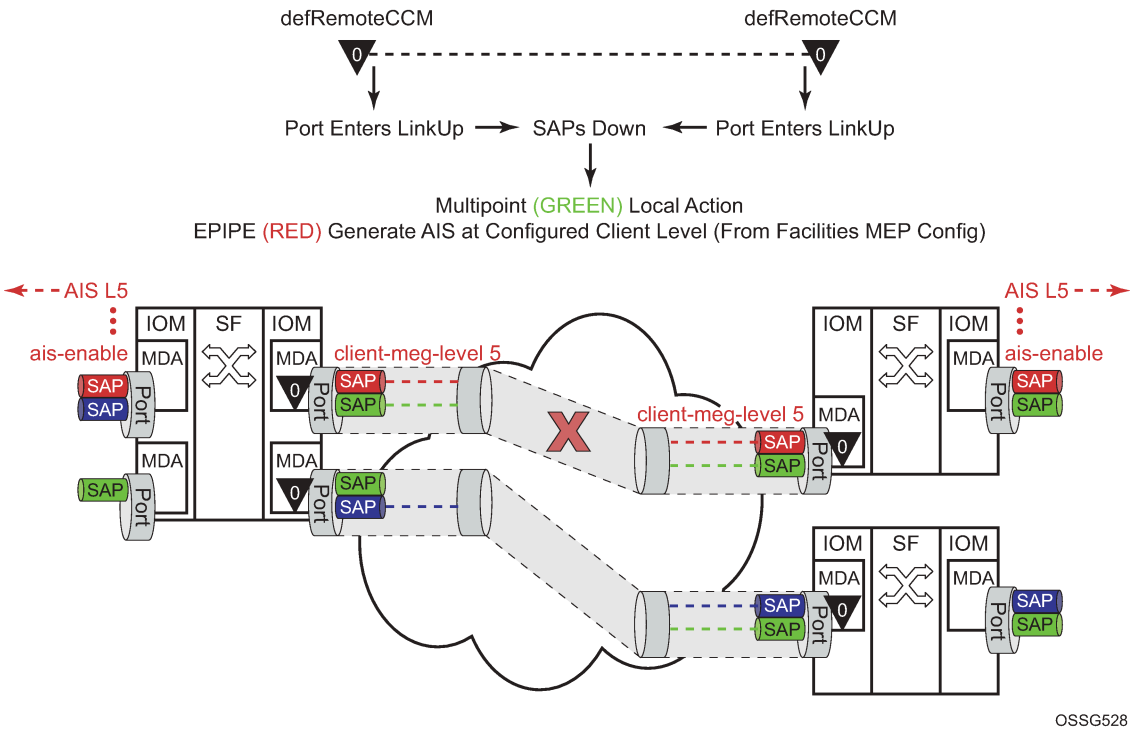

A debounce function has been implemented to prevent notifying every port state change if a port bounces multiple times within a window. Up to four notifications are accepted in a three second window. If the third port state is a down state change, the fourth is ignored. If the fourth port state change is a down state change, it is processed. After that, no further state changes are 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 is 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.

Figure 38: Port-Based MEP example 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 38: Port-Based MEP example

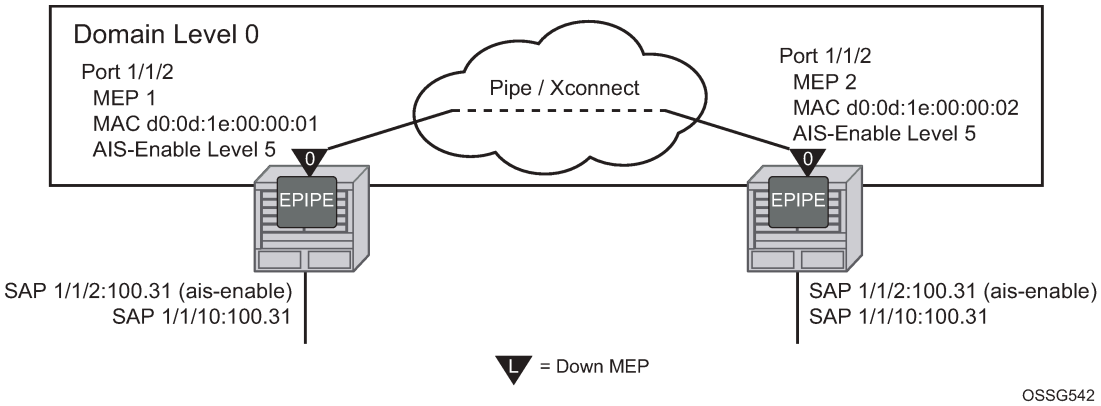

Configure port-based MEPs with the **facility-fault** option and **ais-enable client-meg-level** command. When the MEP enters any defect state, an AIS is generated to any Epipe service that has the ais-enable configured under the **sap>eth-cfm** hierarchy.

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 correct 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.

Because 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.

Because 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 39: Fault handling LAG MEP, 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 39: Fault handling LAG MEP

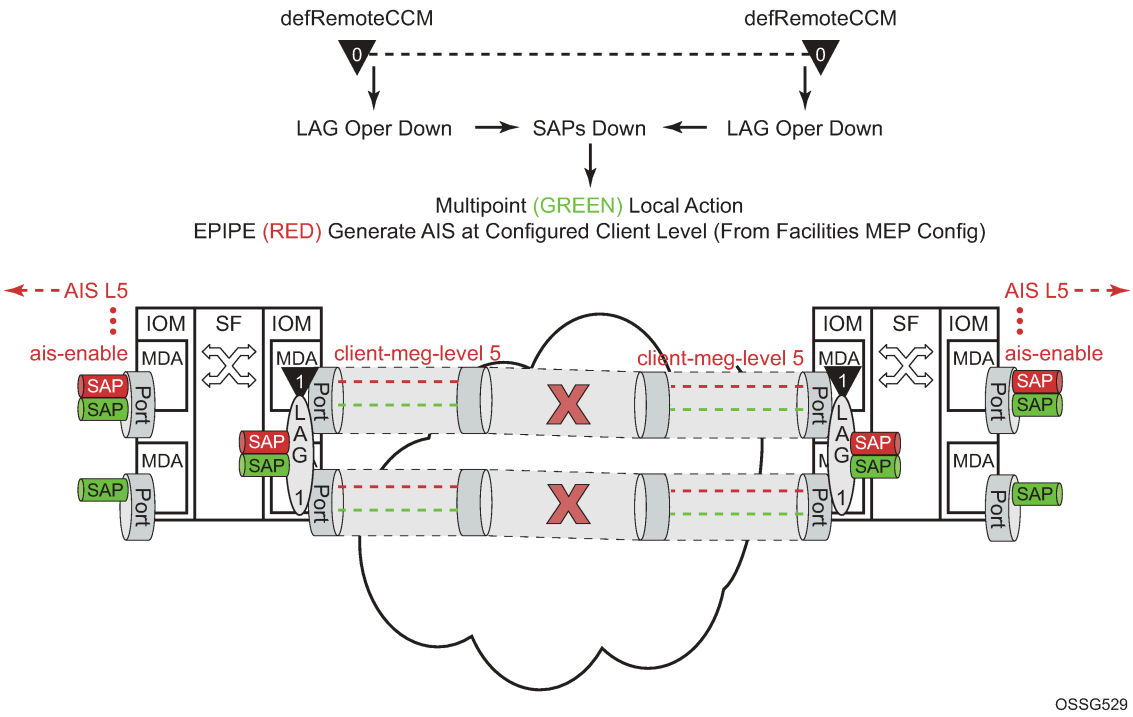

LAG-based MEP are supported for MultiChassis LAG (MC-LAG) configurations.

A LAG facility MEP must not be configured with **facility-fault** when it is applied to an MC-LAG. Traffic goes into a 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.

Figure 40: LAG MEP example uses a port-based MEP to validate port-to-port connectivity.

Figure 40: LAG MEP example

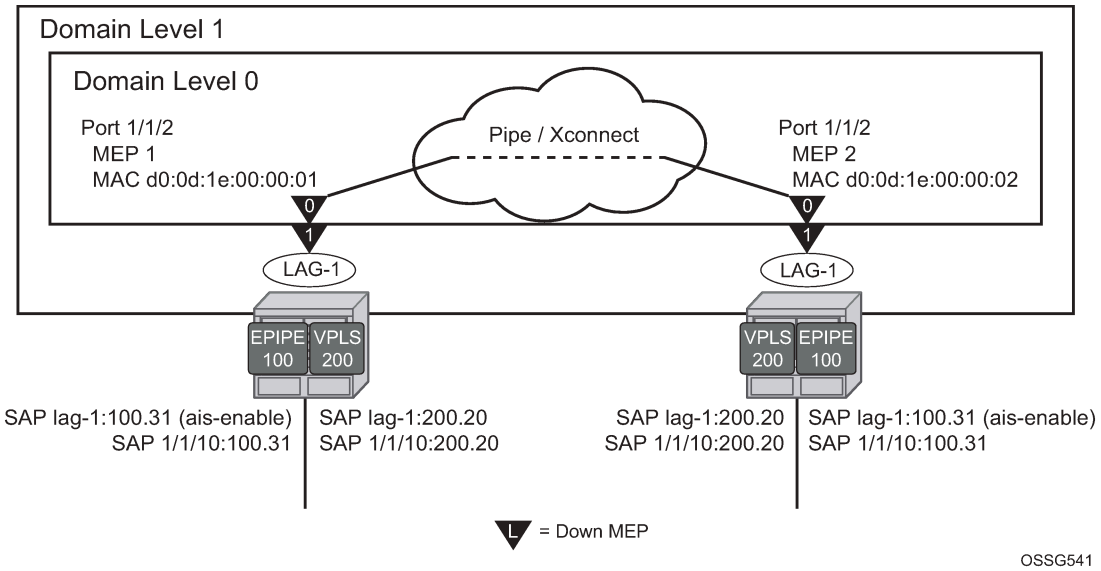

With the introduction of the LAG, the port no longer has direct control over the services SAPs. The ais-enable command has been disabled from the port for this reason. The low-priority-defect condition has been modified to react to all defect conditions "allDef", avoiding the unidirectional issue demonstrated 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 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 are still generated out the mate SAP.

Output from one of the nodes is included below. Because 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. 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 needed 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, because 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. 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 enables 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 does 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 uses 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 41: Tunnel concepts and encapsulation 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 41: Tunnel concepts and encapsulation

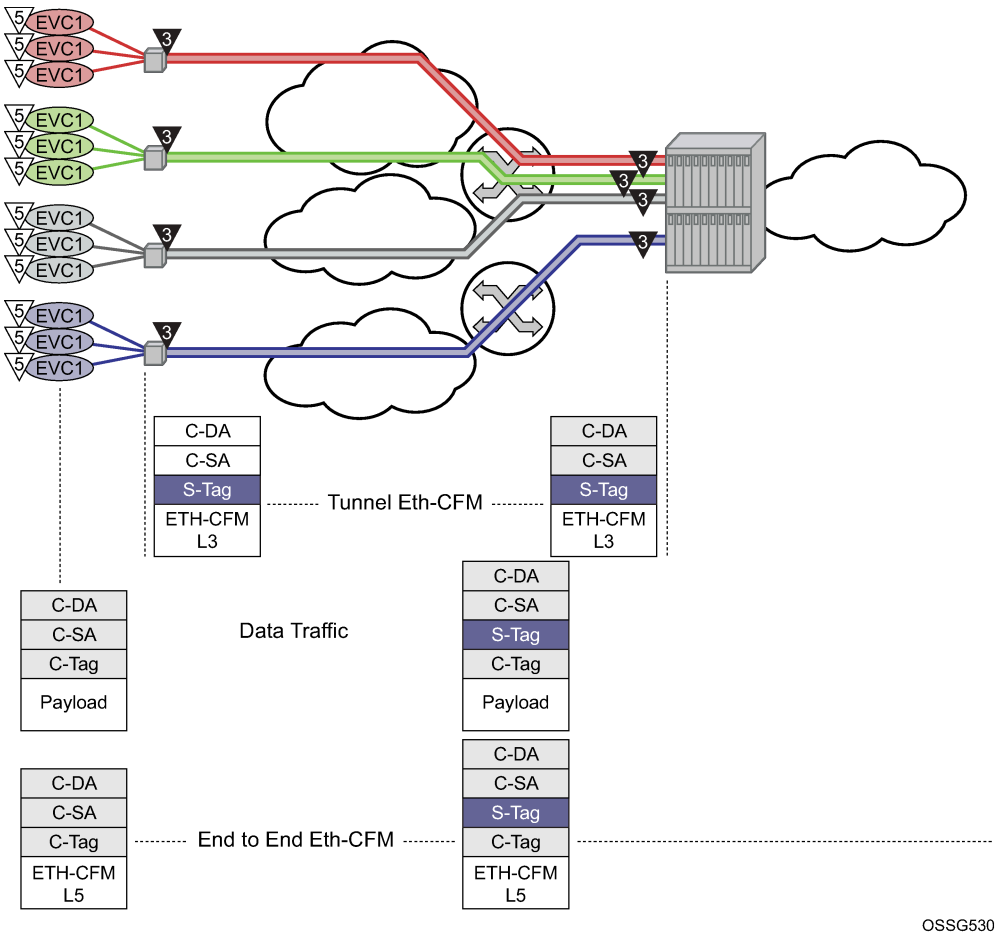

This is only one example of a tagged service. The principles of a tunnel MEP may be applied to other service as applicable. Remember that tunnel MEPs are only supported on LAGs that are configured with QinQ encapsulation and must have an outer VLAN.

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 correct 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.

Because 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. Because 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. Because 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 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 are 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 is only allowed three operational status changes in a three second window before holding the state for the remaining time in that window. Messages are 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 are not able to be sent and failures occur. As an example, a physical port of 1Gb/s that has negotiated an operational speed of 100 Mb/s with a peer is allowed to send up to a maximum of 5 Mb/s of Tunnel MEP control traffic.

Figure 42: Tunnel MEP example shows how fate can be shared between the Tunnel MEP and the services configured on the same LAG and outer VLAN.

Figure 42: Tunnel MEP example

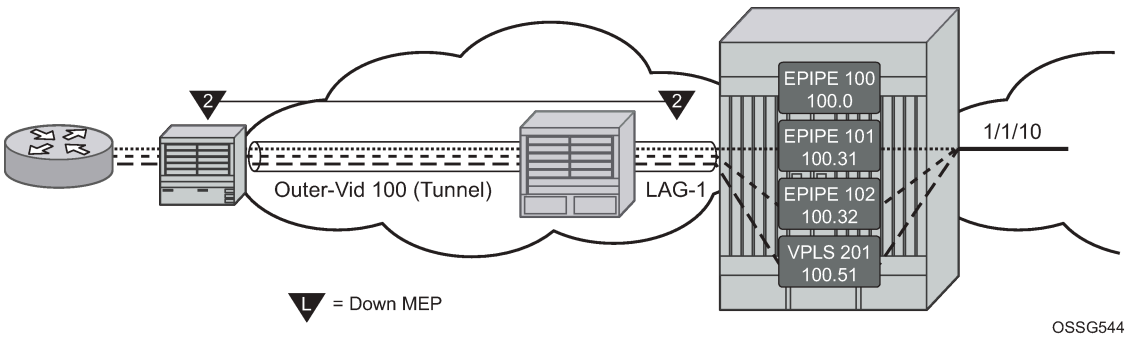

In this example, a single Tunnel, LAG-1 outer VLAN 100, carries three services. Epipe 101, Epipe 102 and VPLS 201 are the service extraction points on the aggregation node. Epipe 100 is the extraction point for the Tunnel MEP eth-cfm traffic. This is a single SAP Epipe that is operationally shutdown. 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. Because 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:

1. Epipe 101 starts to generate AIS out the mate sap.
2. Epipe 102 SAP flag is set.
3. VPLS 201 SAP goes down.

Output from one of the nodes is included below. Because both 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).

The following illustration, Figure 43: Router MEP example, shows how a Router Facility MEP can be configured on a routed interface in the base router instance.

Figure 43: Router MEP example

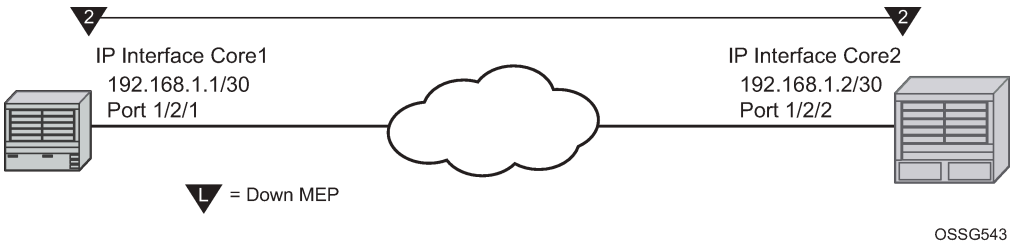

ETH-CFM tools for proactive management (ETH-CC), troubleshooting (Loopback, Linktrace, and so on) and profiling (Delay Measurement, and so on) are supported. The configuration and some ETH-CFM test commands are shown for Node1 (left). Following the on-demand test output, the configuration for Node 2 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 12: Facility MEP support overview provides an overview of Facility MEP support.

Table 12: Facility MEP support overview

	Port MEPs	Tunnel MEPs		LAG MEPs	Router MEPs
		Port	LAG		
Sub Second	Yes	Yes	Yes	Yes	Yes
Port:					
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	—	Yes	Yes	—	—
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 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. Because the MP relies heavily on the underlying MC-LAG construct, consideration must be given for the correct 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 44: Independent processing UP MEP example 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 44: Independent processing UP MEP example

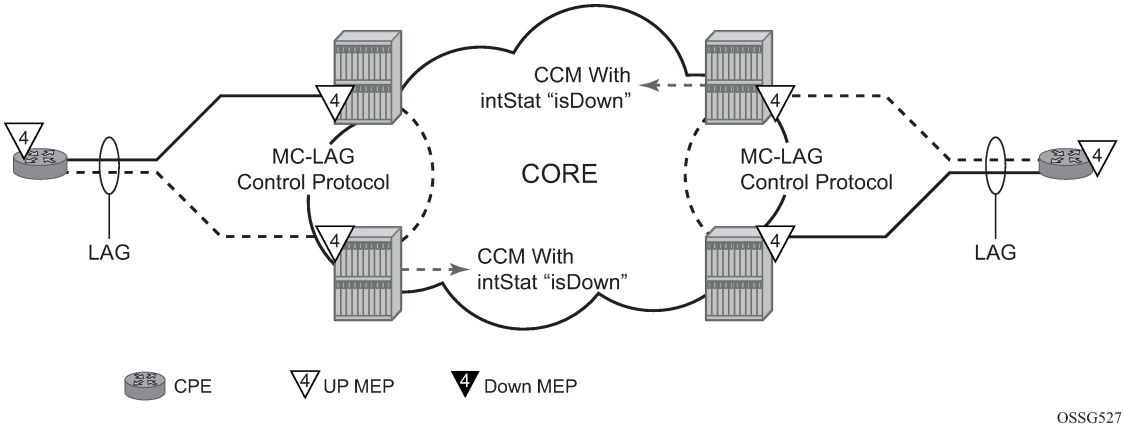

A similar condition exists if down MEPs are configured on the SAPs that are operationally down. Figure 45: Independent processing down MEP example shows how the same service configured with down MEPs would generate AIS, if enabled, toward the remote client at the configured client-meg-level, in the reverse direction of the MEP. This is also the correct behavior from the perspective ETH-CFM.

Figure 45: Independent processing down MEP example

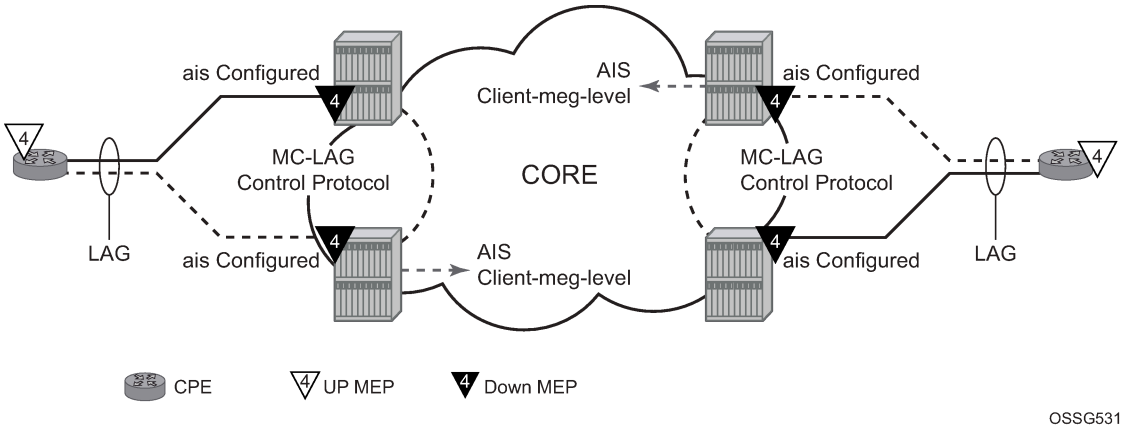

2.15.2.2 Linking ETH-CFM to MC-LAG state

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. 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 does 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 attempt to execute. The tests record failures as a result of the MEP state. These failures are not representative of the network.

This feature relies on the correct 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 registers 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 delays 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 is held for the fixed 1s interval and the last Interface-Status-TLV bits are 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 is 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 assumes 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 does not include any faults. It assumes that the transitioning SAPs are stabilizing as the switchover proceeds. The fixed 1s timer starts and a second CCM PDU based on the UP MEPs interval is 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 represents the state of the SAP.

In scaled environments it is important to configure the propagation-hold-time and the CCM intervals to achieve the needed 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 is not 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.

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 46: ETH-CFM and MC-LAG example

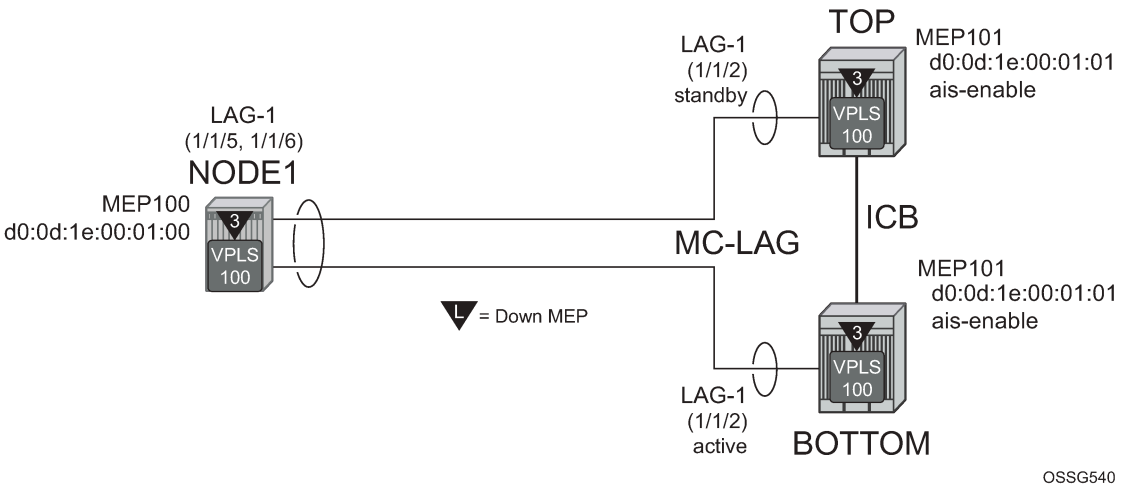

Only one of the MEPs on the MC-LAG nodes is active for VPLS service 100. The other MEP is in a shutdown mode, so that even when the MC-LAG is in standby and the port state is **Link Up**, the MEP is in a pseudo shutdown state.

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 feature: 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. 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, maintain their existing behavior of transitioning the MEP to a failed state and raising the correct 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 continue to use the existing method. Ethernet-tunnels are 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 changes the values without having to delete the command before 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 is 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 before allowing the transition from sub second to non-sub second CCM interval.

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. After 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. See 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 instead of 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 determines whether 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. Correct domain nesting must be followed to avoid ETH-CFM error conditions.

vMEPs support VPLS, M-VPLS, 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 are raised, assuming CCM is enabled, and leaking of ETH-CFM packets occurs (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 are automatically deleted when they coexist on the same SAP or spoke SDP as the **vmep-filter**. Because 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 redirects the ETH-CFM packets to the CPU (consider using CPU protection). Consider 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 because 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 example 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 correct 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.

2.16.1 Basic NGE configuration overview

About this task

This procedure configures NGE for an MPLS service or router interface.

Procedure

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.

- **a.** Configure the encryption and authentication algorithms for the group.
- **b.** Configure a security association (SA) that contains the encryption and authentication keys.
- **c.** Configure the active outbound SA for the group.

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

- **a.** For each SDP, VPRN service, or router interface, configure the outbound direction key group.
- **b.** For each SDP, VPRN service, or router interface, configure the inbound direction key group.

2.16.2 Configuring NGE components

Use the CLI syntax in the subsequent sections to configure NGE parameters.

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:

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

The following example displays global encryption label usage:

```
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:

```
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:

```
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
 0x88433A6DB4FA4F8A490EF661CBE69F010BFAE9C2784BED7059E5ADAAB1A225C6 encryption-key
 0x63DCDD501B66F85441E4A55B597DA617
    config>grp-encryp>encryp-keygrp# security-association spi 6 authentication-key
 0x88433A6DB4FA4F8A490EF661CBE69F010BFAE9C2784BED7059E5ADAAB1A225C5 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, VPRN service, or PW template

A key group can be assigned to the following entities:

- SDPs
- VPRNs
- PW templates

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

- 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
- routed VPLS service into a VPRN or IES
- MP-BGP-based VPRNs

- BGP-VPLS and BGP-VPWS with **auto-gre-sdp**

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, resolving routes using spoke SDPs (**spoke-sdp**) or auto-bind SDPs (**auto-bind-tunnel**) is supported using LDP, GRE, RSVP-TE, or segment routing (SR-ISIS, SR-OSPF, or SR-TE).

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

Note: After assigning a key group to the PW template, the following **tools** command must be executed:
tools>perform>service>eval-pw-template>allow-service-impact

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

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

```
config>service# pw-template policy-id auto-gre-sdp [create]
        – encryption-keygroup keygroup-id direction {inbound | outbound}
```

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

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

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

```
config>service# pw-template 12 auto-gre-sdp create
    config>service>pw-template# encryption-keygroup 4 direction inbound
    config>service>pw-template# encryption-keygroup 4 direction outbound
    config>service>pw-template# exit all
    tools>perform>service>eval-pw-template>allow-service-impact
```

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 describes 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.

```
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
```

```
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.

```
config>service# no customer customer-id
```

```
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 , and then enter the new information.

Note: When created, the SDP encapsulation type cannot be modified.

```
config>service# sdp sdp-id
```

```
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 shut down and removed (unbound) from all customer services where it is applied.

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

```
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 describes NGE management tasks.

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.

```
config# group-encryption
        – encryption-keygroup keygroup-id
            – no active-outbound-sa
            – no security-association spi spi
        – exit
```

```
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
```

```
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
```

```
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
 0x0123456789012345678901234567890123456789012345678901234567890123 encryption-key
 0x0123456789012345678901234567890123456789012345678901234567890123
    config>grp-encryp>encryp-keygrp# security-association spi 6 authentication-key
 0x0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF encryption-key
 0x0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF [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 0x01234567890123456789012345678901234567890123456789012345678901234567890123 encryption-
key 0x01234567890123456789012345678901234567890123456789012345678901234567890123
            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, VPRN service, or PW template

Use the following CLI syntax to remove a key group from an SDP, VPRN service, or PW template:

Note: After removing a key group to the PW template, the following tools command must be executed:
tools>perform>service>eval-pw-template>allow-service-impact

```
config>service# sdp sdp-id
        – no encryption-keygroup direction {inbound | outbound}
```

```
config>service# vprn service-id
        – no encryption-keygroup direction {inbound | outbound}
```

```
config>service# pw-template policy-id auto-gre-sdp
        – no encryption-keygroup direction {inbound | outbound}
```

The following examples display a key group removed from an SDP, VPRN service, or PW template:

```
config>service# sdp 61
    config>service>sdp# no encryption-keygroup direction inbound
    config>service>sdp# no encryption-keygroup direction outbound
```

```
config>service# vprn 22
    config>service>vprn# no encryption-keygroup direction inbound
    config>service>vprn# no encryption-keygroup direction outbound
```

```
config>service# pw-template 12
    config>service>pw-template# no encryption-keygroup direction inbound
    config>service>pw-template# no encryption-keygroup direction outbound
    tools>perform>service>eval-pw-template>allow-service-impact
```

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

About this task

To change a key group requires a removal, a change, and an installation of the key group.

Procedure

Step 1. Remove the inbound direction key group.

Step 2. Change the outbound direction key group.

Step 3. Install the new inbound direction key group.

2.18.4 Changing the key group for an SDP, VPRN service, or PW template

Changing key groups for an SDP, VPRN service, or PW template 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 or PW template is similar.

Note: For PW template changes, the following **tools** command must be executed after the **encryption-keygroup** changes are made:
tools>perform>service>eval-pw-template>allow-service-impact

In the example below, the inbound and outbound key groups are changed from key group 4 to key group 6.

```
config>service# sdp sdp-id
       – no encryption-keygroup direction {inbound|outbound}
```

```
config>service# sdp 61
    config>service>sdp# no encryption-keygroup 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
...
----------------------------------------------
domain1>config>service# info
```

2.18.5 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, PW templates, and router interfaces that use it.

Note: When deleting a key group from a PW template, the following **tools** command must be executed after the **encryption-keygroup** changes are made:
tools>perform>service>eval-pw-template>allow-service-impact

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:

```
config# group-encryption
       – no encryption-keygroup keygroup-id
```

```
config>grp-encryp# no encryption-keygroup 8
```

3 Standards and protocol support

Note:
The information provided in this chapter is subject to change without notice and may not apply to all platforms.

Nokia assumes no responsibility for inaccuracies.

3.1 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*

3.2 Application Assurance (AA)

3GPP Release 12, *ADC rules over Gx interfaces*

RFC 3507, *Internet Content Adaptation Protocol (ICAP)*

3.3 Bidirectional Forwarding Detection (BFD)

draft-ietf-idr-bgp-ls-sbfd-extensions-01, *BGP Link-State Extensions for Seamless BFD*

RFC 5880, *Bidirectional Forwarding Detection (BFD)*

RFC 5881, *Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop)*

RFC 5882, *Generic Application of Bidirectional Forwarding Detection (BFD)*

RFC 5883, *Bidirectional Forwarding Detection (BFD) for Multihop Paths*

RFC 7130, *Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces*

RFC 7880, *Seamless Bidirectional Forwarding Detection (S-BFD)*

RFC 7881, *Seamless Bidirectional Forwarding Detection (S-BFD) for IPv4, IPv6, and MPLS*

RFC 7883, *Advertising Seamless Bidirectional Forwarding Detection (S-BFD) Discriminators in IS-IS*

RFC 7884, *OSPF Extensions to Advertise Seamless Bidirectional Forwarding Detection (S-BFD) Target Discriminators*

3.4 Border Gateway Protocol (BGP)

draft-gredler-idr-bgplu-epe-14, *Egress Peer Engineering using BGP-LU*

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-ls-app-specific-attr-16, *Application-Specific Attributes Advertisement with BGP Link-State*

draft-ietf-idr-bgp-ls-flex-algo-06, *Flexible Algorithm Definition Advertisement with BGP Link-State*

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-idr-long-lived-gr-00, *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 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 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 5492, *Capabilities Advertisement with BGP-4*

RFC 5668, *4-Octet AS Specific BGP Extended Community*

RFC 6286, *Autonomous-System-Wide Unique BGP Identifier for BGP-4*

RFC 6793, *BGP Support for Four-Octet Autonomous System (AS) Number Space*

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 7606, *Revised Error Handling for BGP UPDATE Messages*

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*

RFC 8097, *BGP Prefix Origin Validation State Extended Community*

RFC 8212, *Default External BGP (EBGP) Route Propagation Behavior without Policies*

RFC 8277, *Using BGP to Bind MPLS Labels to Address Prefixes*

RFC 8571, *BGP - Link State (BGP-LS) Advertisement of IGP Traffic Engineering Performance Metric Extensions*

RFC 8950, *Advertising IPv4 Network Layer Reachability Information (NLRI) with an IPv6 Next Hop*

RFC 8955, *Dissemination of Flow Specification Rules*

RFC 8956, *Dissemination of Flow Specification Rules for IPv6*

RFC 9086, *Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing BGP Egress Peer Engineering*

3.5 Broadband Network Gateway (BNG) Control and User Plane Separation (CUPS)

3GPP 23.007, *Restoration procedures*

3GPP 29.244, *Interface between the Control Plane and the User Plane nodes*

3GPP 29.281, *General Packet Radio System (GPRS) Tunnelling Protocol User Plane (GTPv1-U)*

BBF TR-459, *Control and User Plane Separation for a Disaggregated BNG*

RFC 8300, *Network Service Header (NSH)*

3.6 Certificate management

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 5280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*

RFC 6712, *Internet X.509 Public Key Infrastructure -- HTTP Transfer for the Certificate Management Protocol (CMP)*

RFC 7030, *Enrollment over Secure Transport*

RFC 7468, *Textual Encodings of PKIX, PKCS, and CMS Structures*

3.7 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*

3.8 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 802.3ac, *VLAN Tag*

IEEE 802.3ad, *Link Aggregation*

IEEE 802.3ah, *Ethernet in the First Mile*

IEEE 802.3x, *Ethernet Flow Control*

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*

3.9 Ethernet VPN (EVPN)

draft-ietf-bess-evpn-ipvpn-interworking-06, *EVPN Interworking with IPVPN*

draft-ietf-bess-evpn-irb-mcast-04, *EVPN Optimized Inter-Subnet Multicast (OISM) Forwarding* – ingress replication

draft-ietf-bess-evpn-pref-df-06, *Preference-based EVPN DF Election*

draft-ietf-bess-evpn-unequal-lb-16, *Weighted Multi-Path Procedures for EVPN Multi-Homing* – section 9

draft-ietf-bess-evpn-virtual-eth-segment-06, *EVPN Virtual Ethernet Segment*

draft-ietf-bess-pbb-evpn-isid-cmacflush-00, *PBB-EVPN ISID-based CMAC-Flush*

draft-sajassi-bess-evpn-ip-aliasing-05, *EVPN Support for L3 Fast Convergence and Aliasing/Backup Path* – IP Prefix routes

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)*

RFC 8560, *Seamless Integration of Ethernet VPN (EVPN) with Virtual Private LAN Service (VPLS) and Their Provider Backbone Bridge (PBB) Equivalents*

RFC 8584, *DF Election and AC-influenced DF Election*

RFC 9136, *IP Prefix Advertisement in Ethernet VPN (EVPN)*

RFC 9161, *Operational Aspects of Proxy ARP/ND in Ethernet Virtual Private Networks*

RFC 9251, *Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Proxies for Ethernet VPN (EVPN)*

3.10 gRPC Remote Procedure Calls (gRPC)

cert.proto version 0.1.0, *gRPC Network Operations Interface (gNOI) Certificate Management Service*

file.proto version 0.1.0, *gRPC Network Operations Interface (gNOI) File Service*

gnmi.proto version 0.7.0, *gRPC Network Management Interface (gNMI) Service Specification*

PROTOCOL-HTTP2, *gRPC over HTTP2*

system.proto Version 1.0.0, *gRPC Network Operations Interface (gNOI) System Service*

3.11 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, *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 6119, *IPv6 Traffic Engineering in IS-IS*

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 7987, *IS-IS Minimum Remaining Lifetime*

RFC 8202, *IS-IS Multi-Instance* – single topology

RFC 8570, *IS-IS Traffic Engineering (TE) Metric Extensions* – Min/Max Unidirectional Link Delay metric for flex-algo, RSVP, SR-TE

RFC 8919, *IS-IS Application-Specific Link Attributes*

3.12 Internet Protocol (IP) Fast Reroute (FRR)

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)*

RFC 8518, *Selection of Loop-Free Alternates for Multi-Homed Prefixes*

3.13 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 2617, *HTTP Authentication: Basic and Digest Access Authentication*

RFC 2784, *Generic Routing Encapsulation (GRE)*

RFC 2818, *HTTP Over TLS*

RFC 2890, *Key and Sequence Number Extensions to GRE*

RFC 3164, *The BSD syslog Protocol*

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 4511, *Lightweight Directory Access Protocol (LDAP): The Protocol*

RFC 4513, *Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms* – TLS

RFC 4632, *Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*

RFC 5082, *The Generalized TTL Security Mechanism (GTSM)*

RFC 5246, *The Transport Layer Security (TLS) Protocol Version 1.2* – TLS client, RSA public key

RFC 5425, *Transport Layer Security (TLS) Transport Mapping for Syslog* – RFC 3164 with TLS

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*

RFC 7011, *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information*

RFC 7012, *Information Model for IP Flow Information Export*

RFC 7230, *Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing*

RFC 7231, *Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content*

RFC 7232, *Hypertext Transfer Protocol (HTTP/1.1): Conditional Requests*

RFC 7301, *Transport Layer Security (TLS) Application Layer Protocol Negotiation Extension*

RFC 7616, *HTTP Digest Access Authentication*

RFC 8446, *The Transport Layer Security (TLS) Protocol Version 1.3*

3.14 Internet Protocol (IP) multicast

cisco-ipmulticast/pim-autorp-spec01, *Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast* – version 1

draft-ietf-bier-pim-signaling-08, *PIM Signaling Through BIER Core*

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)*

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*

RFC 8444, *OSPFv2 Extensions for Bit Index Explicit Replication (BIER)*

RFC 8487, *Mtrace Version 2: Traceroute Facility for IP Multicast*

RFC 8534, *Explicit Tracking with Wildcard Routes in Multicast VPN* – (C-*,C-*) wildcard

RFC 8556, *Multicast VPN Using Bit Index Explicit Replication (BIER)*

3.15 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)* – relay

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

3.16 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 6437, *IPv6 Flow Label Specification*

RFC 6603, *Prefix Exclude Option for DHCPv6-based Prefix Delegation*

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*

3.17 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 2560, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*

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 4109, *Algorithms for Internet Key Exchange version 1 (IKEv1)*

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 4754, *IKE and IKEv2 Authentication Using the Elliptic Curve Digital Signature Algorithm (ECDSA)*

RFC 4835, *Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (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 5282, *Using Authenticated Encryption Algorithms with the Encrypted Payload of the IKEv2 Protocol*

RFC 5903, *ECP Groups for IKE and IKEv2*

RFC 5996, *Internet Key Exchange Protocol Version 2 (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 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)*

3.18 Label Distribution Protocol (LDP)

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 7473, *Controlling State Advertisements of Non-negotiated LDP Applications*

RFC 7552, *Updates to LDP for IPv6*

3.19 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"*

3.20 Multiprotocol Label Switching (MPLS)

draft-ietf-mpls-lsp-ping-ospfv3-codepoint-02, *OSPFv3 CodePoint for MPLS LSP Ping*

RFC 3031, *Multiprotocol Label Switching Architecture*

RFC 3032, *MPLS Label Stack Encoding*

RFC 3270, *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services* – E-LSP

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 6374, *Packet Loss and Delay Measurement for MPLS Networks* – Delay Measurement, Channel Type 0x000C

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 6790, *The Use of Entropy Labels in MPLS Forwarding*

RFC 7510, *Encapsulating MPLS in UDP*

RFC 7746, *Label Switched Path (LSP) Self-Ping*

RFC 7876, *UDP Return Path for Packet Loss and Delay Measurement for MPLS Networks* – Delay Measurement

RFC 8029, *Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures*

3.21 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*

3.22 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 7753, *Port Control Protocol (PCP) Extension for Port-Set Allocation*

RFC 7915, *IP/ICMP Translation Algorithm*

3.23 Network Configuration Protocol (NETCONF)

RFC 5277, *NETCONF Event Notifications*

RFC 6020, *YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)*

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 8342, *Network Management Datastore Architecture (NMDA)* – Startup, Candidate, Running and Intended datastores

RFC 8525, *YANG Library*

RFC 8526, *NETCONF Extensions to Support the Network Management Datastore Architecture* – <get-data> operation

3.24 Open Shortest Path First (OSPF)

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 5642, *Dynamic Hostname Exchange Mechanism for OSPF*

RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*

RFC 5838, *Support of Address Families in OSPFv3*

RFC 6549, *OSPFv2 Multi-Instance Extensions*

RFC 6987, *OSPF Stub Router Advertisement*

RFC 7471, *OSPF Traffic Engineering (TE) Metric Extensions* – Min/Max Unidirectional Link Delay metric for flex-algo, RSVP, SR-TE

RFC 7684, *OSPFv2 Prefix/Link Attribute Advertisement*

RFC 7770, *Extensions to OSPF for Advertising Optional Router Capabilities*

RFC 8362, *OSPFv3 Link State Advertisement (LSA) Extensibility*

RFC 8920, *OSPF Application-Specific Link Attributes*

3.25 OpenFlow

TS-007 Version 1.3.1, *OpenFlow Switch Specification* – OpenFlow-hybrid switches

3.26 Path Computation Element Protocol (PCEP)

draft-ietf-pce-binding-label-sid-15, *Carrying Binding Label/Segment Identifier (SID) in PCE-based Networks.* – MPLS binding SIDs

draft-alvarez-pce-path-profiles-04, *PCE Path Profiles*

draft-dhs-spring-pce-sr-p2mp-policy-00, *PCEP extensions for p2mp sr policy*

RFC 5440, *Path Computation Element (PCE) Communication Protocol (PCEP)*

RFC 8231, *Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE*

RFC 8281, *PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model*

RFC 8408, *Conveying Path Setup Type in PCE Communication Protocol (PCEP) Messages*

RFC 8664, *Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing*

3.27 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*

3.28 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 4006, *Diameter Credit-Control Application*

RFC 6733, *Diameter Base Protocol*

3.29 Pseudowire (PW)

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*

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 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*

RFC 7392, *Explicit Path Routing for Dynamic Multi-Segment Pseudowires* – ER-TLV and ER-HOP IPv4 Prefix

RFC 8395, *Extensions to BGP-Signaled Pseudowires to Support Flow-Aware Transport Labels*

3.30 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 2597, *Assured Forwarding PHB Group*

RFC 3140, *Per Hop Behavior Identification Codes*

RFC 3246, *An Expedited Forwarding PHB (Per-Hop Behavior)*

3.31 Remote Authentication Dial In User Service (RADIUS)

draft-oscca-cfrg-sm3-02, *The SM3 Cryptographic Hash Function*

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 6613, *RADIUS over TCP* – with TLS

RFC 6614, *Transport Layer Security (TLS) Encryption for RADIUS*

RFC 6929, *Remote Authentication Dial-In User Service (RADIUS) Protocol Extensions*

RFC 6911, *RADIUS attributes for IPv6 Access Networks*

3.32 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*

3.33 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*

3.34 Segment Routing (SR)

draft-bashandy-rtgwg-segment-routing-uloop-06, *Loop avoidance using Segment Routing*

draft-filsfils-spring-net-pgm-extension-srv6-usid-13, *Network Programming extension: SRv6 uSID instruction*

draft-filsfils-spring-srv6-net-pgm-insertion-04, *SRv6 NET-PGM extension: Insertion*

draft-ietf-6man-spring-srv6-oam-10, *Operations, Administration, and Maintenance (OAM) in Segment Routing Networks with IPv6 Data plane (SRv6)*

draft-ietf-idr-bgp-ls-segment-routing-ext-16, *BGP Link-State extensions for Segment Routing*

draft-ietf-idr-segment-routing-te-policy-11, *Advertising Segment Routing Policies in BGP*

draft-ietf-isis-mpls-elc-10, *Signaling Entropy Label Capability and Entropy Readable Label Depth Using IS-IS* – advertising ELC

draft-ietf-lsr-flex-algo-16, *IGP Flexible Algorithm*

draft-ietf-lsr-isis-srv6-extensions-14, *IS-IS Extension to Support Segment Routing over IPv6 Dataplane*

draft-ietf-ospf-mpls-elc-12, *Signaling Entropy Label Capability and Entropy Readable Label-stack Depth Using OSPF* – advertising ELC

draft-ietf-rtgwg-segment-routing-ti-lfa-01, *Topology Independent Fast Reroute using Segment Routing*

draft-ietf-spring-conflict-resolution-05, *Segment Routing MPLS Conflict Resolution*

draft-ietf-spring-segment-routing-policy-08, *Segment Routing Policy Architecture*

draft-ietf-teas-sr-rsvp-coexistence-rec-02, *Recommendations for RSVP-TE and Segment Routing LSP co-existence*

draft-voyer-6man-extension-header-insertion-10, *Deployments With Insertion of IPv6 Segment Routing Headers*

draft-voyer-pim-sr-p2mp-policy-02, *Segment Routing Point-to-Multipoint Policy*

draft-voyer-spring-sr-p2mp-policy-03, *SR Replication Policy for P2MP Service Delivery*

RFC 8287, *Label Switched Path (LSP) Ping/Traceroute for Segment Routing (SR) IGP-Prefix and IGP-Adjacency Segment Identifiers (SIDs) with MPLS Data Planes*

RFC 8476, *Signaling Maximum SID Depth (MSD) Using OSPF* – node MSD

RFC 8491, *Signaling Maximum SID Depth (MSD) Using IS-IS* – node MSD

RFC 8660, *Segment Routing with the MPLS Data Plane*

RFC 8661, *Segment Routing MPLS Interworking with LDP*

RFC 8663, *MPLS Segment Routing over IP* – BGP SR with SR-MPLS-over-UDP/IP

RFC 8665, *OSPF Extensions for Segment Routing*

RFC 8666, *OSPFv3 Extensions for Segment Routing*

RFC 8667, *IS-IS Extensions for Segment Routing*

RFC 8669, *Segment Routing Prefix Segment Identifier Extensions for BGP*

RFC 8754, *IPv6 Segment Routing Header (SRH)*

RFC 8814, *Signaling Maximum SID Depth (MSD) Using the Border Gateway Protocol - Link State*

RFC 8986, *Segment Routing over IPv6 (SRv6) Network Programming*

RFC 9252, *BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)*

3.35 Simple Network Management Protocol (SNMP)

draft-blumenthal-aes-usm-04, *The AES Cipher Algorithm in the SNMP's User-based Security Model* – CFB128-AES-192 and CFB128-AES-256

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

ESO-CONSORTIUM-MIB revision 200406230000Z, *esoConsortiumMIB*

IANA-ADDRESS-FAMILY-NUMBERS-MIB revision 200203140000Z, *ianaAddressFamilyNumbers*

IANAifType-MIB revision 200505270000Z, *ianaifType*

IANA-RTPROTO-MIB revision 200009260000Z, *ianaRtProtoMIB*

IEEE8021-CFM-MIB revision 200706100000Z, *ieee8021CfmMib*

IEEE8021-PAE-MIB revision 200101160000Z, *ieee8021paeMIB*

IEEE8023-LAG-MIB revision 200006270000Z, *lagMIB*

LLDP-MIB revision 200505060000Z, *lldpMIB*

RFC 1157, *A Simple Network Management Protocol (SNMP)*

RFC 1212, *Concise MIB Definitions*

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 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 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 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 3410, *Introduction and Applicability Statements for Internet Standard Management Framework*

RFC 3411, *An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks*

RFC 3412, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*

RFC 3413, *Simple Network Management Protocol (SNMP) Applications*

RFC 3414, *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*

RFC 3415, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*

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 3418, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*

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 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 7420, *Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module*

RFC 7630, *HMAC-SHA-2 Authentication Protocols in the User-based Security Model (USM) for SNMPv3*

SFLOW-MIB revision 200309240000Z, *sFlowMIB*

3.36 Timing

GR-1244-CORE Issue 3, *Clocks for the Synchronized Network: Common Generic Criteria*

GR-253-CORE Issue 3, *SONET Transport Systems: Common Generic Criteria*

IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*

ITU-T G.781, *Synchronization layer functions*

ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC)*

ITU-T G.8261, *Timing and synchronization aspects in packet networks*

ITU-T G.8262, *Timing characteristics of synchronous Ethernet equipment slave clock (EEC)*

ITU-T G.8262.1, *Timing characteristics of an enhanced synchronous Ethernet equipment slave clock (eEEC)*

ITU-T G.8264, *Distribution of timing information through packet networks*

ITU-T G.8265.1, *Precision time protocol telecom profile for frequency synchronization*

ITU-T G.8275.1, *Precision time protocol telecom profile for phase/time synchronization with full timing support from the network*

RFC 3339, *Date and Time on the Internet: Timestamps*

RFC 5905, *Network Time Protocol Version 4: Protocol and Algorithms Specification*

3.37 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*

RFC 8545, *Well-Known Port Assignments for the One-Way Active Measurement Protocol (OWAMP) and the Two-Way Active Measurement Protocol (TWAMP)* – TWAMP

RFC 8762, *Simple Two-Way Active Measurement Protocol* – unauthenticated

RFC 8972, *Simple Two-Way Active Measurement Protocol Optional Extensions* – unauthenticated

3.38 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)*

3.39 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, *RTP: A Transport Protocol for Real-Time Applications* – Appendix A.8

RFC 4585, *Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)*

RFC 4588, *RTP Retransmission Payload Format*

3.40 Wireless Local Area Network (WLAN) gateway

3GPP TS 23.402, *Architecture enhancements for non-3GPP accesses* – S2a roaming based on GPRS

3.41 Yet Another Next Generation (YANG)

RFC 6991, *Common YANG Data Types*

RFC 7950, *The YANG 1.1 Data Modeling Language*

RFC 7951, *JSON Encoding of Data Modeled with YANG*

3.42 Yet Another Next Generation (YANG) OpenConfig Modules

openconfig-aaa.yang version 0.4.0, *OpenConfig AAA Module*

openconfig-aaa-radius.yang version 0.3.0, *OpenConfig AAA RADIUS Module*

openconfig-aaa-tacacs.yang version 0.3.0, *OpenConfig AAA TACACS+ Module*

openconfig-acl.yang version 1.0.0, *OpenConfig ACL Module*

openconfig-bfd.yang version 0.1.0, *OpenConfig BFD Module*

openconfig-bgp.yang version 3.0.1, *OpenConfig BGP Module*

openconfig-bgp-common.yang version 3.0.1, *OpenConfig BGP Common Module*

openconfig-bgp-common-multiprotocol.yang version 3.0.1, *OpenConfig BGP Common Multiprotocol Module*

openconfig-bgp-common-structure.yang version 3.0.1, *OpenConfig BGP Common Structure Module*

openconfig-bgp-global.yang version 3.0.1, *OpenConfig BGP Global Module*

openconfig-bgp-neighbor.yang version 3.0.1, *OpenConfig BGP Neighbor Module*

openconfig-bgp-peer-group.yang version 3.0.1, *OpenConfig BGP Peer Group Module*

openconfig-bgp-policy.yang version 4.0.1, *OpenConfig BGP Policy Module*

openconfig-if-aggregate.yang version 2.0.0, *OpenConfig Interfaces Aggregated Module*

openconfig-if-ethernet.yang version 2.0.0, *OpenConfig Interfaces Ethernet Module*

openconfig-if-ip.yang version 2.0.0, *OpenConfig Interfaces IP Module*

openconfig-if-ip-ext.yang version 2.0.0, *OpenConfig Interfaces IP Extensions Module*

openconfig-igmp.yang version 0.2.0, *OpenConfig IGMP Module*

openconfig-interfaces.yang version 2.0.0, *OpenConfig Interfaces Module*

openconfig-isis.yang version 0.3.2, *OpenConfig IS-IS Module*

openconfig-isis-policy.yang version 0.3.2, *OpenConfig IS-IS Policy Module*

openconfig-isis-routing.yang version 0.3.2, *OpenConfig IS-IS Routing Module*

openconfig-lacp.yang version 1.1.0, *OpenConfig LACP Module*

openconfig-lldp.yang version 0.1.0, *OpenConfig LLDP Module*

openconfig-local-routing.yang version 1.0.1, *OpenConfig Local Routing Module*

openconfig-mpls.yang version 2.3.0, *OpenConfig MPLS Module*

openconfig-mpls-ldp.yang version 3.0.2, *OpenConfig MPLS LDP Module*

openconfig-mpls-rsvp.yang version 2.3.0, *OpenConfig MPLS RSVP Module*

openconfig-mpls-te.yang version 2.3.0, *OpenConfig MPLS TE Module*

openconfig-network-instance.yang version 1.1.0, *OpenConfig Network Instance Module*

openconfig-network-instance-l3.yang version 0.11.1, *OpenConfig L3 Network Instance Module* – static routes

openconfig-packet-match.yang version 1.0.0, *OpenConfig Packet Match Module*

openconfig-pim.yang version 0.2.0 *OpenConfig PIM Module*

openconfig-platform.yang version 0.12.2, *OpenConfig Platform Module*

openconfig-platform-fan.yang version 0.1.1, *OpenConfig Platform Fan Module*

openconfig-platform-linecard.yang version 0.1.2, *OpenConfig Platform Linecard Module*

openconfig-platform-port.yang version 0.3.3, *OpenConfig Port Module*

openconfig-platform-transceiver.yang version 0.7.1, *OpenConfig Transceiver Module*

openconfig-procmon.yang version 0.4.0, *OpenConfig Process Monitoring Module*

openconfig-relay-agent.yang version 0.1.0, *OpenConfig Relay Agent Module*

openconfig-routing-policy.yang version 3.0.0, *OpenConfig Routing Policy Module*

openconfig-rsvp-sr-ext.yang version 0.1.0, *OpenConfig RSVP-TE and SR Extensions Module*

openconfig-system.yang version 0.9.1, *OpenConfig System Module*

openconfig-system-logging.yang version 0.3.1, *OpenConfig System Logging Module*

openconfig-system-terminal.yang version 0.3.0, *OpenConfig System Terminal Module*

openconfig-telemetry.yang version 0.5.0, *OpenConfig Telemetry Module*

openconfig-terminal-device.yang version 1.7.3, *OpenConfig Terminal Optics Device Module*

openconfig-vlan.yang version 2.0.0, *OpenConfig VLAN Module*

Customer document and product support

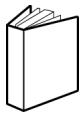

Customer documentation

Customer documentation welcome page

Technical support

Product support portal

Documentation feedback

Customer documentation feedback