



7705 Service Aggregation Router Gen 2

Release 25.3.R2

Services Overview Guide

3HE 21582 AAAA TQZZA 01

Edition: 01

April 2025

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.

This document includes Nokia proprietary and confidential information, which may not be distributed or disclosed to any third parties without the prior written consent of Nokia.

This document is intended for use by Nokia's customers ("You"/"Your") in connection with a product purchased or licensed from any company within Nokia Group of Companies. Use this document as agreed. You agree to notify Nokia of any errors you may find in this document; however, should you elect to use this document for any purpose(s) for which it is not intended, You understand and warrant that any determinations You may make or actions You may take will be based upon Your independent judgment and analysis of the content of this document.

Nokia reserves the right to make changes to this document without notice. At all times, the controlling version is the one available on Nokia's site.

No part of this document may be modified.

NO WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF AVAILABILITY, ACCURACY, RELIABILITY, TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, IS MADE IN RELATION TO THE CONTENT OF THIS DOCUMENT. IN NO EVENT WILL NOKIA BE LIABLE FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, DIRECT, INDIRECT, INCIDENTAL OR CONSEQUENTIAL OR ANY LOSSES, SUCH AS BUT NOT LIMITED TO LOSS OF PROFIT, REVENUE, BUSINESS INTERRUPTION, BUSINESS OPPORTUNITY OR DATA THAT MAY ARISE FROM THE USE OF THIS DOCUMENT OR THE INFORMATION IN IT, EVEN IN THE CASE OF ERRORS IN OR OMISSIONS FROM THIS DOCUMENT OR ITS CONTENT.

Copyright and trademark: Nokia is a registered trademark of Nokia Corporation. Other product names mentioned in this document may be trademarks of their respective owners.

© 2025 Nokia.

Table of contents

List of tables.....	8
List of figures.....	9
1 Getting started.....	11
1.1 About this guide.....	11
1.2 Platforms and terminology.....	11
1.3 Conventions.....	12
1.3.1 Precautionary and information messages.....	12
1.3.2 Options or substeps in procedures and sequential workflows.....	12
2 Introduction.....	14
2.1 Service types.....	14
2.2 Service policies.....	15
2.2.1 Multipoint shared queuing.....	15
2.2.1.1 Ingress queuing modes of operation.....	16
2.2.1.2 Ingress service queuing.....	16
2.2.1.3 Ingress shared queuing.....	17
2.2.1.4 Ingress multipoint shared queuing.....	19
3 Nokia service model.....	21
4 Service entities.....	22
4.1 Customers.....	22
4.2 SAPs.....	22
4.2.1 SAP encapsulation types and identifiers.....	23
4.2.2 Ethernet encapsulations.....	23
4.2.3 Default SAP on a dot1q port.....	24
4.2.4 QinQ SAPs.....	25
4.2.5 Services and SAP encapsulations.....	29
4.2.6 SAP configuration considerations.....	30
4.2.7 SAP bandwidth CAC.....	30
4.2.7.1 CAC enforcement.....	31
4.3 Connection profile VLAN SAPs.....	33

4.3.1	Using connection profile VLAN SAPs in dot1q ports.....	35
4.3.2	Using connection profile VLAN SAPs in QinQ ports.....	35
4.4	Service destination points.....	36
4.4.1	SDP binding.....	37
4.4.2	Spoke and mesh SDPs.....	37
4.4.3	SDP using BGP route tunnel.....	38
4.4.4	SDP keepalives.....	38
4.4.5	SDP selection rules.....	39
4.4.6	Class-based forwarding.....	39
4.4.6.1	Application of class-based forwarding over RSVP LSPs.....	39
4.4.6.2	Operation of class-based forwarding over RSVP LSPs.....	40
4.4.7	Source IPv4 address configuration in GRE SDP and GRE tunnel.....	41
4.4.7.1	Introduction and feature configuration.....	41
4.4.7.2	Feature operation with T-LDP and BGP service label signaling.....	43
4.4.8	GRE SDP tunnel fragmentation and reassembly.....	45
4.4.8.1	GRE SDP tunnel fragmentation.....	46
4.4.8.2	GRE SDP tunnel reassembly.....	46
4.4.8.3	NGE considerations.....	49
4.4.9	GRE SDP termination on router interface IP address.....	49
4.5	SAP and MPLS binding loopback with MAC swap.....	50
4.6	Promiscuous ETH-LBM mode of operation.....	55
5	Multiservice sites.....	57
6	Internal objects created for L2TP and NAT.....	58
7	Ethernet unnumbered interfaces.....	59
8	ECMP and weighted ECMP for services using RSVP and SR-TE LSPs.....	60
9	NGE.....	61
9.1	NGE overview.....	61
9.1.1	NGE key groups and encryption partitions.....	62
9.1.2	Network services platform management.....	63
9.2	Key groups.....	64
9.2.1	Key group algorithms.....	65
9.2.1.1	Encapsulating security payload.....	66

9.2.2	Security associations.....	66
9.2.2.1	Active outbound SA.....	66
9.3	Services encryption.....	67
9.3.1	Services encryption overview.....	67
9.3.2	Assigning key groups to services.....	69
9.3.3	VPRN Layer 3 spoke SDP encryption and MP-BGP-based VPRN encryption interaction.....	70
9.3.4	L2 service encryption using PW templates.....	71
9.3.5	Pseudowire switching for NGE traffic.....	71
9.3.6	Pseudowire control word for NGE traffic.....	71
9.3.7	NGE and RFC 8277.....	72
9.3.8	NGE for NG-MVPN.....	72
9.4	NGE packet overhead and MTU considerations.....	72
9.5	Statistics.....	74
9.6	Remote network monitoring support.....	75
9.7	Configuration notes.....	75
9.7.1	Enabling NGE for an SDP or VPRN service.....	75
9.7.2	Enabling NGE for a router interface.....	76
9.7.3	Changing NGE from one key group to another key group for an SDP or VPRN service..	76
9.7.4	Changing NGE from one key group to another key group for a router interface.....	76
9.7.5	Disabling NGE for an SDP or VPRN service.....	76
9.7.6	Disabling NGE for a router interface.....	76
10	Service creation process overview.....	77
11	Deploying and provisioning services.....	78
11.1	Building the core network.....	78
11.2	Performing service administration.....	78
11.3	Provisioning services.....	79
12	General configuration notes.....	80
13	Configuring global service entities with CLI.....	81
13.1	Service model entities.....	81
13.2	Basic configuration.....	81
13.3	Common configuration tasks.....	82
13.3.1	Configuring customers.....	82

13.3.1.1	Customer information.....	82
13.3.1.2	Configuring multiservice-sites.....	83
13.3.2	Configuring an SDP.....	83
13.3.2.1	SDP configuration tasks.....	83
13.3.2.2	Configuring a mixed-LSP SDP.....	84
14	Configuring NGE with CLI.....	86
14.1	Basic NGE configuration overview.....	86
14.2	Configuring NGE components.....	86
14.2.1	Configuring the global encryption label.....	86
14.2.2	Configuring a key group.....	87
14.2.3	Assigning a key group to an SDP, VPRN service, or PW template.....	88
15	Global service entity management tasks.....	90
15.1	Modifying customer accounts.....	90
15.2	Deleting customers.....	90
15.3	Modifying SDPs.....	90
15.4	Deleting SDPs.....	91
16	NGE management tasks.....	92
16.1	Modifying a key group.....	92
16.2	Removing a key group.....	93
16.2.1	Removing a key group from an SDP, VPRN service, or PW template.....	93
16.3	Changing key groups.....	93
16.3.1	Changing the key group for an SDP, VPRN service, PW template, or WLAN-GW group interface.....	94
16.4	Deleting a key group from an NGE node.....	94
17	Standards and protocol support.....	95
17.1	Bidirectional Forwarding Detection (BFD).....	95
17.2	Border Gateway Protocol (BGP).....	95
17.3	Bridging and management.....	96
17.4	Certificate management.....	97
17.5	Ethernet VPN (EVPN).....	97
17.6	gRPC Remote Procedure Calls (gRPC).....	97
17.7	Intermediate System to Intermediate System (IS-IS).....	98
17.8	Internet Protocol (IP) general.....	99

17.9	Internet Protocol (IP) multicast.....	99
17.10	Internet Protocol (IP) version 4.....	100
17.11	Internet Protocol (IP) version 6.....	100
17.12	Internet Protocol Security (IPsec).....	101
17.13	Label Distribution Protocol (LDP).....	102
17.14	Multiprotocol Label Switching (MPLS).....	103
17.15	Network Address Translation (NAT).....	103
17.16	Open Shortest Path First (OSPF).....	103
17.17	Path Computation Element Protocol (PCEP).....	104
17.18	Pseudowire (PW).....	104
17.19	Quality of Service (QoS).....	105
17.20	Remote Authentication Dial In User Service (RADIUS).....	105
17.21	Resource Reservation Protocol - Traffic Engineering (RSVP-TE).....	105
17.22	Routing Information Protocol (RIP).....	106
17.23	Segment Routing (SR).....	106
17.24	Simple Network Management Protocol (SNMP).....	106
17.25	Timing.....	108
17.26	Two-Way Active Measurement Protocol (TWAMP).....	108
17.27	Virtual Private LAN Service (VPLS).....	109
17.28	Yet Another Next Generation (YANG).....	109

List of tables

Table 1: Platforms and terminology.....11

Table 2: System/port settings (QinQ x.0 access SAP control enabled).....26

Table 3: Service and SAP encapsulations.....30

Table 4: SAP lookup matching order for dot1q ports.....35

Table 5: SAP lookup matching order for QinQ ports (QinQ x.0 access SAP control enabled).....35

Table 6: MAC-Swap configuration and options.....52

Table 7: NGE overhead for MPLS.....72

Table 8: NGE overhead for router interface.....73

Table 9: Accounting for NGE overhead SDP and service MTU — calculation examples.....73

List of figures

Figure 1: Unicast service queue mapping to multiple destination based hardware queues.....	17
Figure 2: Unicast service queuing with shared queuing enabled.....	18
Figure 3: Multipoint queue behavior with shared queuing enabled.....	19
Figure 4: Multipoint shared queuing using first pass unicast queues.....	20
Figure 5: Service entities.....	22
Figure 6: 7705 SAR Gen 2 service access point (SAP).....	23
Figure 7: Multiple SAPs on a single port/channel.....	24
Figure 8: Example 1 SAP classification QinQ ports.....	27
Figure 9: Example 2 SAP classification QinQ ports.....	28
Figure 10: VLAN tag handling.....	34
Figure 11: GRE SDP pointing from ALA-A to ALA-B.....	37
Figure 12: Class-based forwarding over SDP LSPs.....	39
Figure 13: Mismatched T-LDP control plane configuration.....	44
Figure 14: Proper setting of T-LDP control plane configuration.....	44
Figure 15: Source address mismatch between control and data planes.....	45
Figure 16: Ingress loopback packet processing.....	51
Figure 17: Egress loopback packet processing.....	52
Figure 18: Active loopback mode.....	54
Figure 19: NGE network with NSP NFM-P management.....	61
Figure 20: Key group partitioning.....	63
Figure 21: Key groups and a typical NGE packet.....	65

Figure 22: NGE MPLS/GRE/MPLSoUDP label stack.....	68
Figure 23: NGE and packet formats.....	69
Figure 24: Inbound and outbound key group assignments.....	70
Figure 25: Service creation and implementation flow.....	77

1 Getting started


1.1 About this guide

This guide describes subscriber services, and mirroring support provided by the 7705 SAR Gen 2 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.

Unless otherwise indicated, the topics and commands described in this guide apply only to the 7705 SAR Gen 2 platforms listed in [Platforms and terminology](#).


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



Note: Unless otherwise indicated, CLI commands, contexts, and configuration examples in this guide apply for both the classic CLI and the MD-CLI.


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

- *7705 SAR Gen 2 Classic CLI Command Reference Guide*
- *7705 SAR Gen 2 Clear, Monitor, Show, Tools CLI Command Reference Guide* (for both the MD-CLI and classic CLI)
- *7705 SAR Gen 2 MD-CLI Command Reference Guide*



Note: This guide generically covers Release 25.x.Rx content and may contain some content that will be released in later maintenance loads. See the *SR OS R25.x.Rx Software Release Notes*, part number 3HE 21562 000x TQZZA, for information about features supported in each load of the Release 25.x.Rx software. For a list of features and CLI commands that are present in SR OS but not supported on the 7705 SAR Gen 2 platforms, see "SR OS Features not Supported on SAR Gen 2" in the *SR OS R25.x.Rx Software Release Notes*.

1.2 Platforms and terminology



Note: Unless explicitly noted otherwise, this guide uses the terminology defined in the following table to collectively designate the specified platforms.

Table 1: Platforms and terminology

Platform	Collective platform designation
7705 SAR-1	7705 SAR Gen 2

1.3 Conventions

This section describes the general conventions used in this guide.

1.3.1 Precautionary and information messages

The following information symbols are 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 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 destination 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 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) service**

- **Ethernet pipe (Epipe)**

This is a Layer 2 point-to-point VLL service for Ethernet frames.

See the *7705 SAR Gen 2 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 *7705 SAR Gen 2 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 *7705 SAR Gen 2 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 *7705 SAR Gen 2 Layer 3 Services Guide: IES and VPRN* for more information about VPRN services.

2.2 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 options 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.

2.2.1 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.

2.2.1.1 Ingress queuing modes of operation

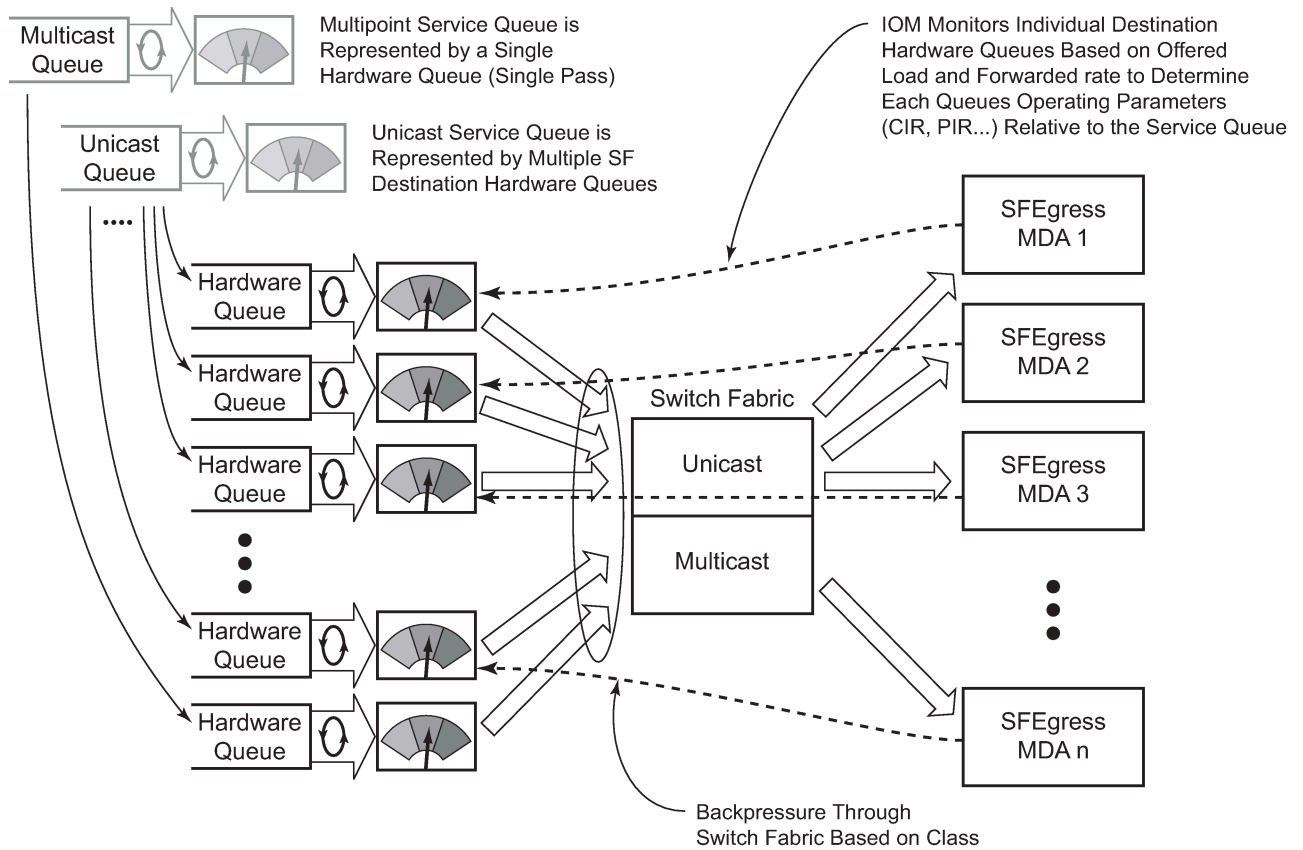
Three modes of ingress SAP queuing are supported for multipoint services (IES, VPLS and VPRN); service, shared, and multipoint shared. The same ingress queuing options are available for IES and VPLS subscriber SLA profile instance queuing.

2.2.1.2 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



OSSG225

2.2.1.3 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 command options 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

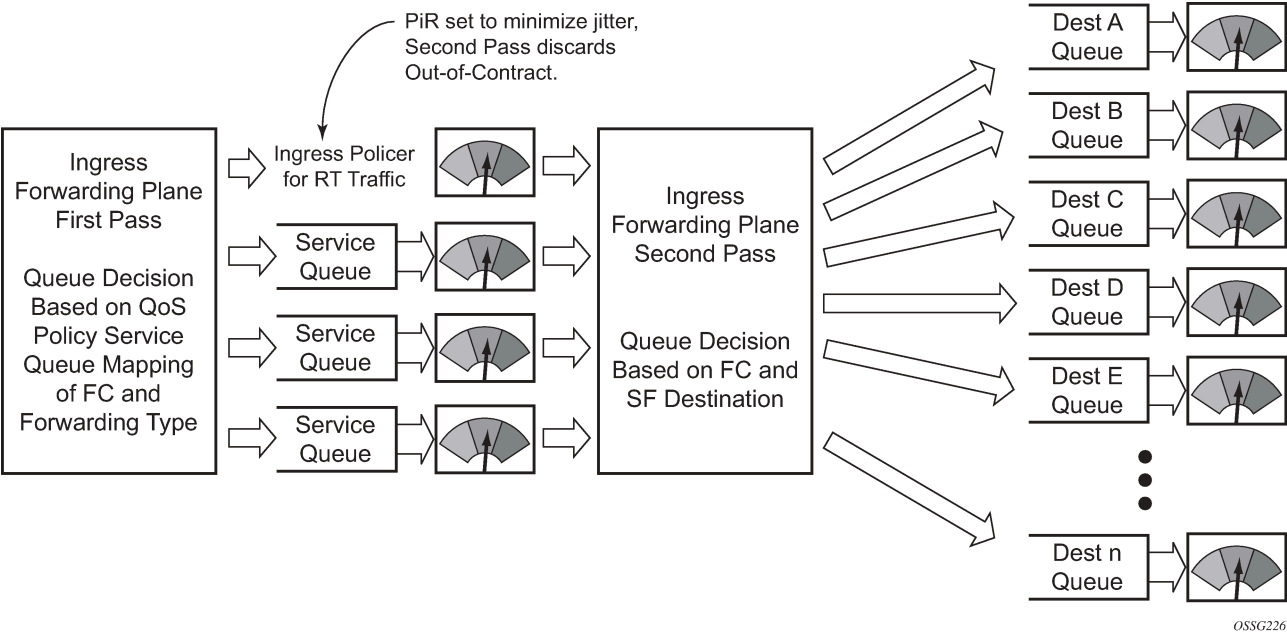
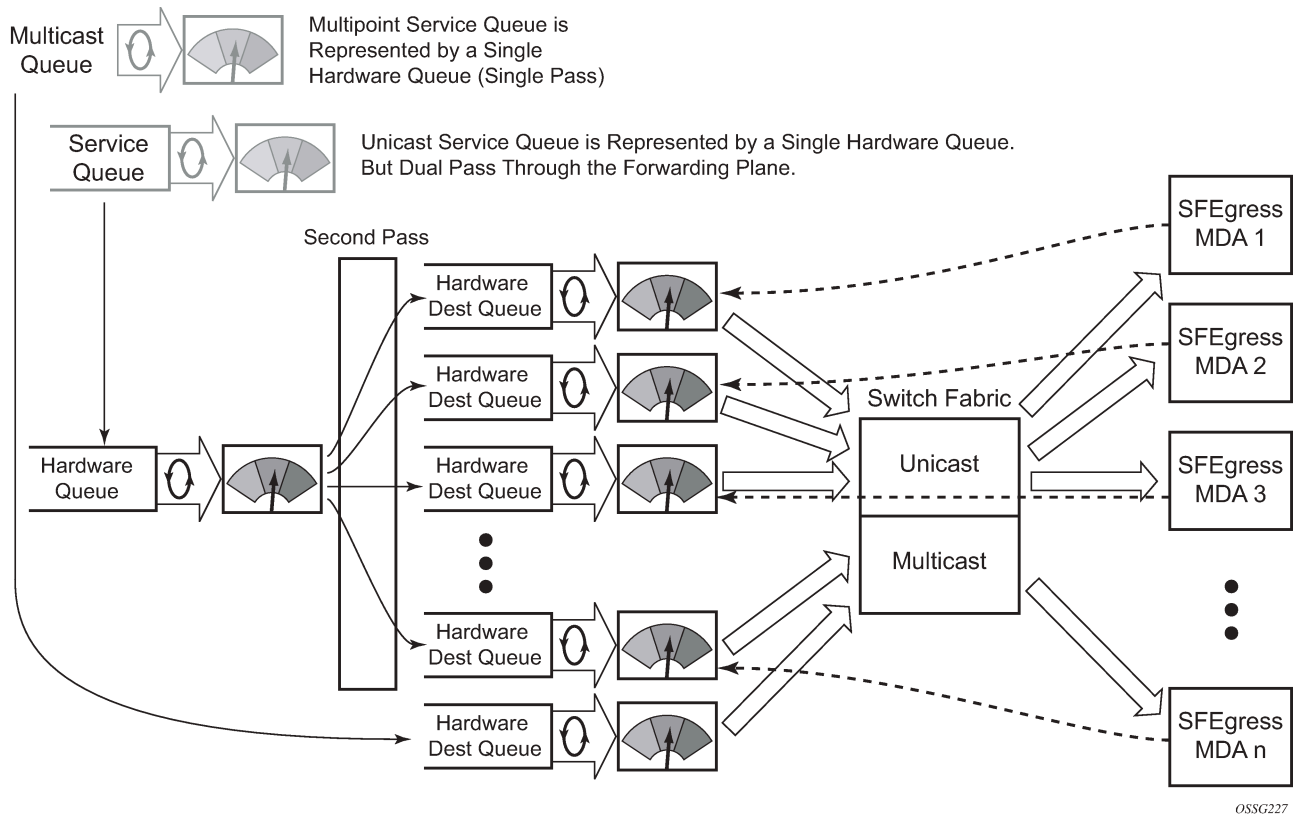


Figure 3: Multipoint queue behavior with shared queuing enabled



2.2.1.4 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 command options 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



OSSG228

3 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 or IP/MPLS provider core network, or both, 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 7705 SAR Gen 2 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 command options and statistics from ports to customers to services.

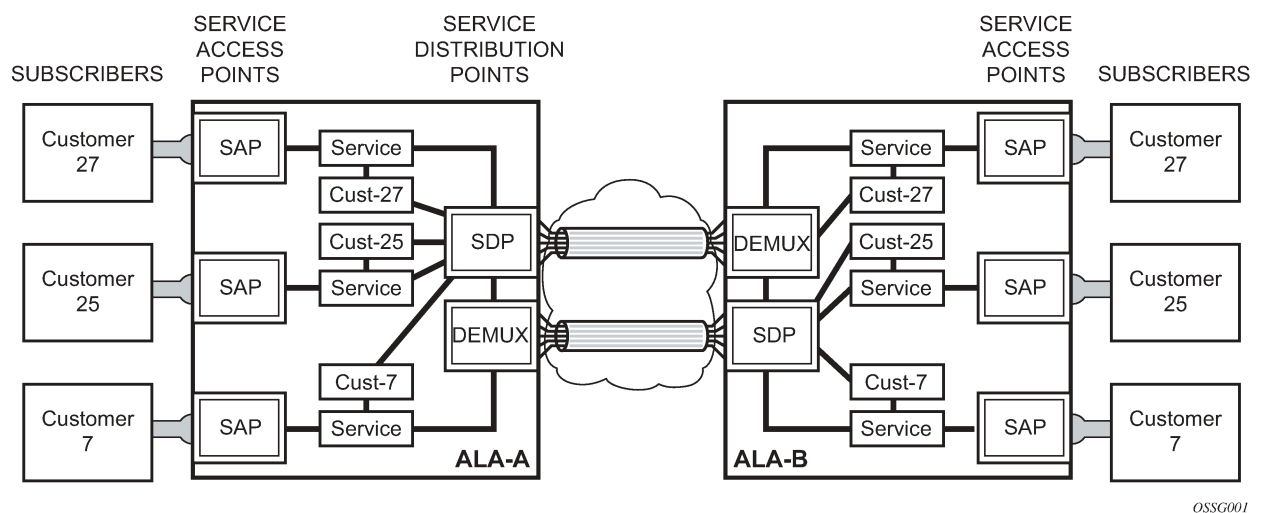
Service provisioning uses logical entities to provision a service where additional properties can be configured for bandwidth provisioning, QoS, security filtering, accounting/billing to the appropriate entity.

4 Service entities

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

- **Customers**
- **SAPs**
- **Service destination points** (for distributed services only)

Figure 5: Service entities



4.1 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.

4.2 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: 7705 SAR Gen 2 service access point \(SAP\)](#)). The SAP configuration requires that slot, MDA, and port/channel information be specified. The slot, MDA, and port/channel must be configured before provisioning a service (see the "Cards, MDAs, and Ports" sections of the *7705 SAR Gen 2 Interface Configuration Guide*).

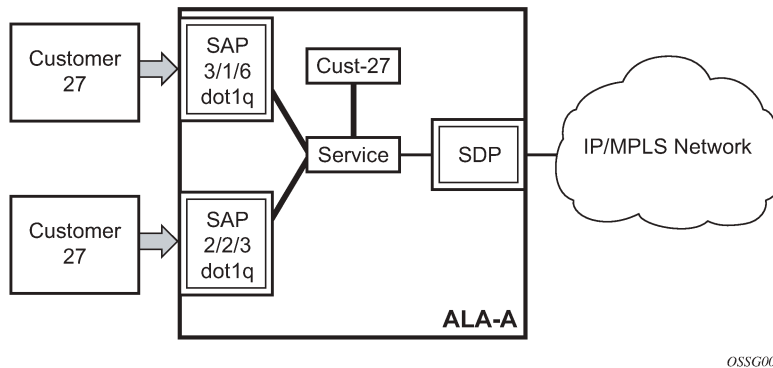
A SAP is a local entity to the router and is uniquely identified by:

- 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: 7705 SAR Gen 2 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.

4.2.1 SAP encapsulation types and identifiers

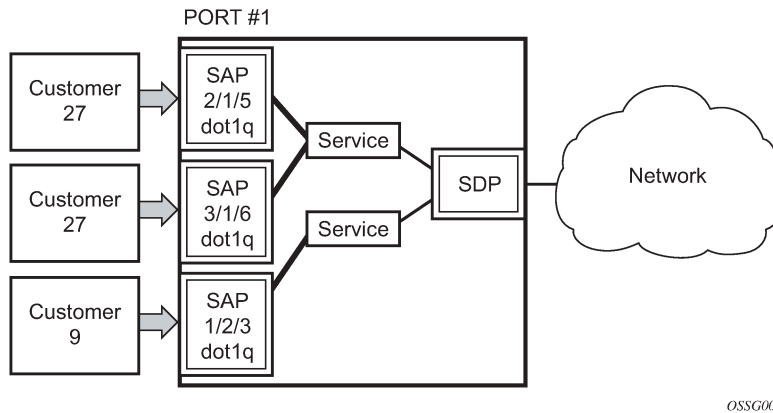
The encapsulation type is an access property of a service Ethernet port or SONET/SDH or TDM channel. The appropriate encapsulation type for the port or channel depends on the requirements to support multiple services on a single port or channel on the associated SAP and the capabilities of the downstream 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.

4.2.2 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. For example, the port is connected to a multitenant 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.

Figure 7: Multiple SAPs on a single port/channel



OSSG003

4.2.3 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 Q-tag 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 are 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.

4.2.4 QinQ SAPs

A QinQ SAP has the following format:

qinq <port-id | lag-id>:<qtag1 | cp-conn-prof-id>

Where:

- *qtag1* is the outer Q-tag value - [* , null, 0 to 4094]
- *qtag2* is the inner Q-tag value - [* , null, 0 to 4094]
- *cp*: keyword
- *conn-prof-id*: 1 to 8000

Regular QinQ SAPs have qtag1 and qtag2 values between 1 and 4094. In addition, QinQ Ethernet and LAG ports support the additional "default" SAPs. Use the following command to enable default SAPs for QinQ Ethernet and LAG ports:

- **MD-CLI**

```
configure service system extended-default-qinq-sap-lookup
```

- **classic CLI**

```
configure system ethernet new-qinq-untagged-sap
```

QinQ Ethernet and LAG ports support these additional "default" SAPs:

- **'*.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'** or **'*.*'**, or both, 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 Q-tag 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 1/1/1:10.0 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.

You can control the forwarding of packets on a QinQ X.0 access SAP to only accept frames with a single tag matching the SAP outer tag or frames with double tags where the outer tag matches the SAP outer tag and the inner tag is set to 0. [Table 2: System/port settings \(QinQ x.0 access SAP control enabled\)](#) shows the SAP lookup precedence order for incoming frames with <qtag1.qtag2> Q-tag values when this option is enabled.

Table 2: System/port settings (QinQ x.0 access SAP control enabled)

Incoming frame qtag1.qtag2	SAP lookup precedence order						
	:X.Y	:X.0	:X.*	:0.*	:null.null	:* null	:*.*
x.y	1st		2nd				3rd
x.0		1st	2nd				3rd
0.y				1st			2nd
0.0				1st			2nd
x		1st	2nd			3rd	4th
0				1st		2nd	3rd
<untagged>				1st	2nd	3rd	4th

The following considerations apply to the information described in [Table 2: System/port settings \(QinQ x.0 access SAP control enabled\)](#):

- 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: System/port settings \(QinQ x.0 access SAP control enabled\)](#) 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:

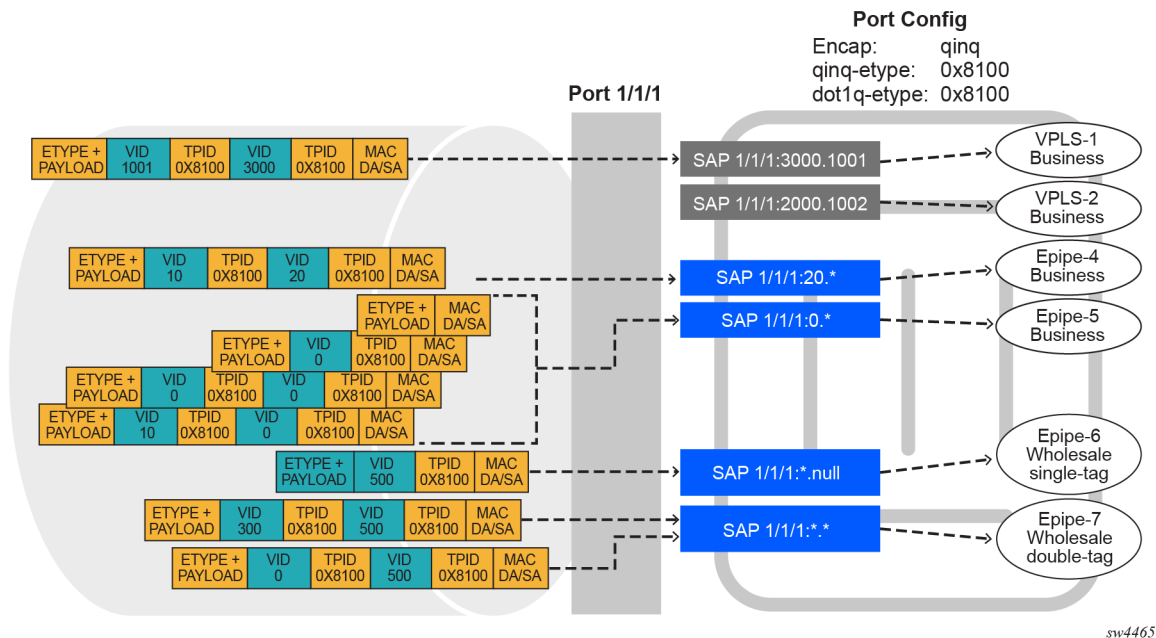
- In a port with SAPs 1/1/1:0.* and 1/1/1:*. * defined, the SAP lookup for a specific frame with VIDs (0, 300) yields SAP 1/1/1:0.* regardless of its operational status.
- The frame only matches SAP 1/1/1:*. * when the 0.* SAP is removed from the configuration.
- The following applies to VLAN tag handling:
 - 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 QinQ x.0 access SAP control is enabled.

As shown in [Figure 8: Example 1 SAP classification QinQ ports](#), assuming that QinQ x.0 access SAP control 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.* - 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



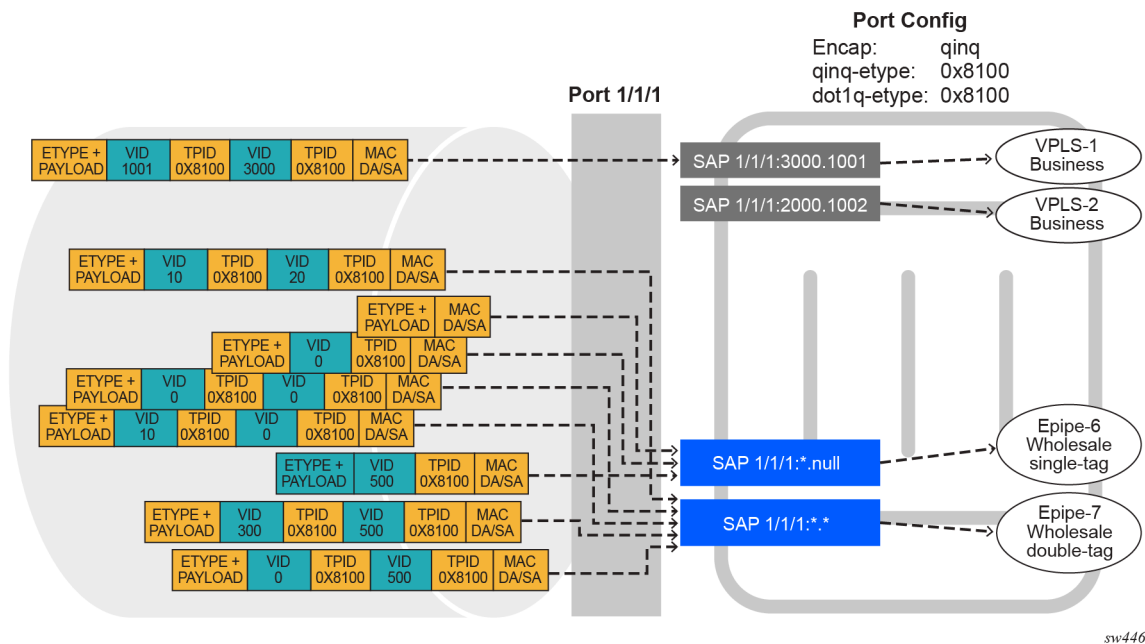
Based on the SAPs configuration described above, the incoming traffic is classified in the following way - notation (outer-VID, inner-VID):

- (3000, 1001) goes to vpls-1

- (20, 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 QinQ x.0 access SAP control 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:*.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, 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: System/port settings \(QinQ x.0 access SAP control enabled\)](#) 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
 - BPDU-translation
 - Eth-tunnels
 - IGMP-snooping
 - MLD-snooping

4.2.5 Services and SAP encapsulations

The following table lists the services and SAP encapsulations.

Table 3: Service and SAP encapsulations

Port type	Encapsulation
Ethernet	Null
Ethernet	Dot1q
Ethernet	QinQ

4.2.6 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 for the SAP is 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 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

4.2.7 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 under booking 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 commands to configure the administrator bandwidth and booking factor for the port or LAG.

```
configure lag access bandwidth
configure lag access booking-factor

configure port ethernet access bandwidth
configure port ethernet access booking-factor

configure service {epipe | ipipe | vpls} sap bandwidth
configure service {ies | vprn} interface sap bandwidth
```

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.

4.2.7.1 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.

Output 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.

```
=====
Ethernet Interface
=====
...

Access Bandwidth      : 500000 Booking Factor :100
Access Available BW   : 490000
Access Booked BW      : 10000

...
=====
```

If an increase of the SAP administrator bandwidth to 600 Mb/s is attempted; the operation fails with the following error because of insufficient available administrator bandwidth on the port with the following error:

```
MINOR: SVCMGR #2664 Insufficient bandwidth available
```

If the booking factor for the port is increased to 200%, 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.

```
=====
Ethernet Interface
```



```

=====
...
Access Bandwidth   : 500000                      Booking Factor : 200
Access Available BW: 400000
Access Booked BW   : 600000
...
=====

```

4.3 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 following commands define the range of customer VIDs to be matched when the CP SAP is associated with a dot1q or QinQ SAP:

- **MD-CLI**

```
configure connection-profile vlan qtag-range
```

- **classic CLI**

```
configure connection-profile-vlan vlan-range
```

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 command:

```
tools dump resource-usage system
```

See the *7705 SAR Gen 2 Clear, Monitor, Show, Tools CLI Command Reference Guide* for a complete description of the **tools dump resource-usage** commands:

Output example

```

=====
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
=====

```

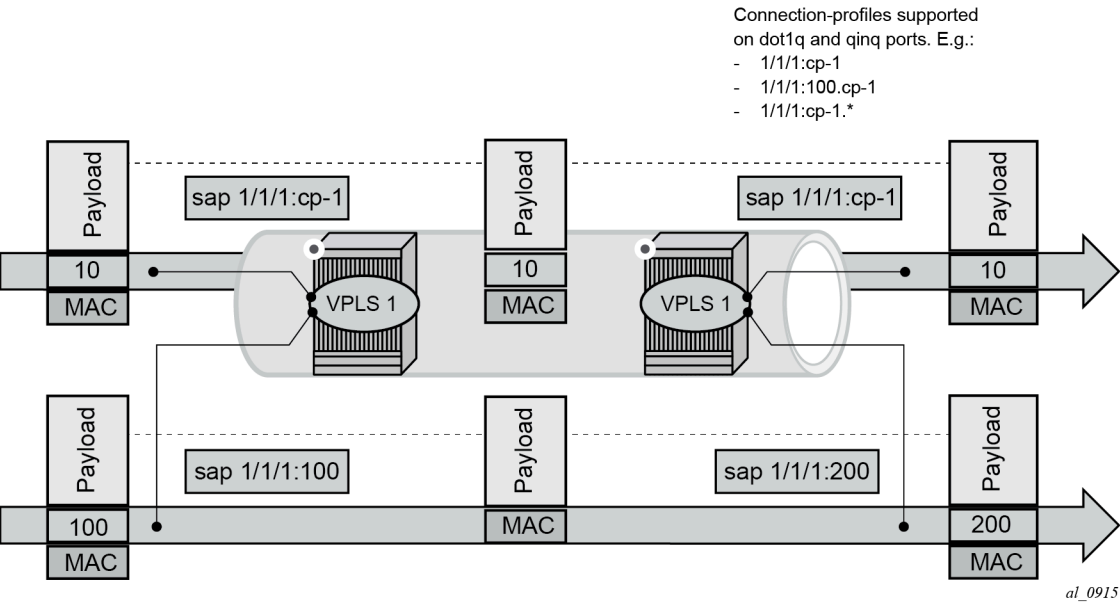
	Total	Allocated	Free

<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 VID of 1 to 50, 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 CP SAP allows the configuration of VLAN ranges with the following characteristics:

- A VLAN range can be defined as a single VID (for example, 101), or two VIDs delimiting the beginning and the end of the range (for example, 105 to 107).
- Discontinuous ranges are allowed.
- Overlapping ranges are not allowed within the same CP SAP 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 CP SAP must have at least one range (with a single or multiple VIDs) before it can be associated with a SAP.
- A CP SAP 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 CP 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.
-
- In the classic CLI, the **mac-ping**, **mac-trace**, **mac-purge**, and **mac-populate** OAM commands are not supported for CP SAPs.

4.3.1 Using connection profile VLAN SAPs in dot1q ports

[Table 4: SAP lookup matching order for dot1q ports](#) describes the SAP lookup matching order that is applied when a CP SAP is used in dot1q ports.

Table 4: SAP lookup matching order for dot1q ports

Incoming frame Q-tag 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

4.3.2 Using connection profile VLAN SAPs in QinQ ports

[Table 5: SAP lookup matching order for QinQ ports \(QinQ x.0 access SAP control enabled\)](#) describes the SAP lookup matching order that is applied when CP SAPs is used in QinQ ports.

Table 5: SAP lookup matching order for QinQ ports (QinQ x.0 access SAP control enabled)

Incoming frame qtag1qtag2 .	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

Incoming frame <i>qtag1qtag2</i> .	SAP lookup precedence order (assumption: X and Y are defined in CP ranges)							
	:X.Y	:X.0	:X.CP	:CP.*	:X.*	:0.*	:.null	:.*
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.*.

4.4 Service destination points

A service destination 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 or MPLS).


- An SDP is a management entity. Even though the SDP configuration and the services carried within are independent, they are related objects. Operations on the SDP affect all the services associated with the SDP. For example, the operational and administrative state of an SDP controls the state of services bound to the SDP.

An SDP from the local device to a far-end 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.

4.4.1 SDP binding

To configure a distributed service from ALA-A to ALA-B, the SDP ID (1) (shown in [Figure 11: GRE 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 SDP pointing from ALA-A to ALA-B



4.4.2 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.

4.4.3 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 command 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.

4.4.4 SDP keepalives

SDP keepalives actively monitor the SDP operational state using periodic Nokia SDP ping echo requests and echo reply messages. Nokia SDP ping is a part of the Nokia suite of service diagnostics built on a Nokia service-level OAM protocol. When the SDP ping is used in the SDP keepalive application, the SDP echo requests and echo reply messages are a mechanism for exchanging the far-end SDP status.

Configuring SDP keepalives on an SDP is optional. SDP keepalives 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 are 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. Message lengths for echo requests are configurable. If max drop count echo request messages do not receive an echo reply, the SDP immediately becomes operationally down.

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

4.4.5 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.

4.4.6 Class-based forwarding

4.4.6.1 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.

4.4.6.2 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.

4.4.7 Source IPv4 address configuration in GRE SDP and GRE tunnel

4.4.7.1 Introduction and feature configuration

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

```
configure service sdp local-end
```

The default value of the **local-end** command option 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 when configuring the local end:

- You can configure a maximum of 15 distinct address values under the following contexts:
 - **GRE SDPs**

```
configure service sdp local-end
```

- **L2oGRE SDPs**

```
configure service system gre-eth-bridged tunnel-termination
```

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
- VPLS with a provisioned spoke SDP
- BGP-AD VPLS with the provisioned SDP configured to use or prefer
- BGP-VPLS with the provisioned SDP configured to use or prefer

- Epipe with a provisioned spoke SDP
- Epipe with BGP-VPWS with the provisioned SDP configured to use or prefer

For services that support auto-binding to a GRE tunnel, the following command configures a single alternate source address per system.

```
configure service system vpn-gre-source-ip
```

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

A change to the value of the **vpn-gre-source-ip** command option 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** command option value:

- This single source address counts toward the maximum of 15 distinct address values per system that are used under the following contexts:

- **GRE SDPs**

```
configure service sdp local-end
```

- **L2oGRE SDPs**

```
configure service system gre-eth-bridged tunnel-termination
```

- The same source address can be used in both **vpn-gre-source-ip** and **configure service sdp local-end** contexts.
- The same source address cannot be used in both **vpn-gre-source-ip** and **configure 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
- 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, as shown in the following command. The command provisioned SDP can also be set to prefer:

- **MD-CLI**

```
configure service provisioned-sdp sdp-include
```

- **classic CLI**

```
configure service pw-template use-provisioned-sdp sdp-include
```

4.4.7.2 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 following command:


```
configure service sdp local-end
```

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 configuration](#) 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 configuration



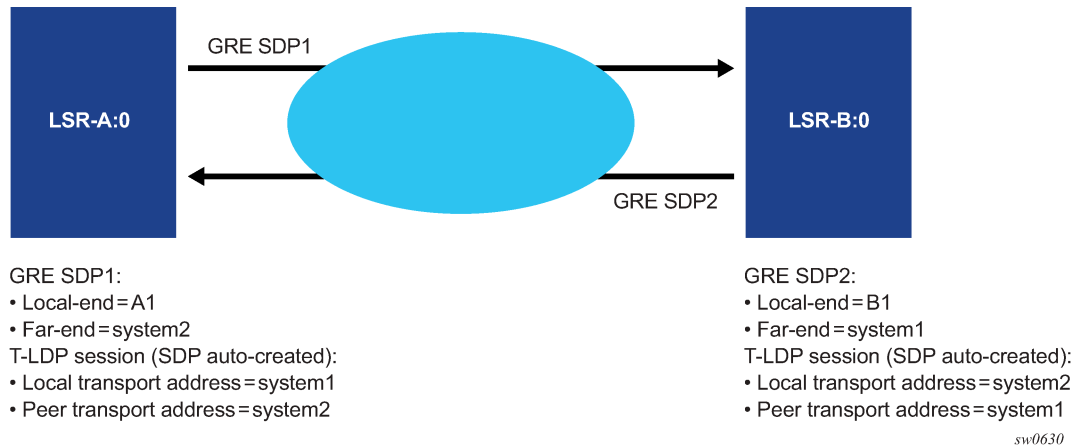
The user must match the local transport address of the T-LDP session to the local-end address of the GRE SDP in both the local and remote PE routers. This can be achieved by manually configuring a T-LDP session to the peer, or by auto-creating a T-LDP session with the targeted peer template feature, and setting the local LSR ID to the address configured in the local-end address 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 configuration](#) shows an example of a correct configuration.

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



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



Note: The LDP uses a priority mechanism to select which provisioned options 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 provisioned options. As long as the corresponding change is performed at the far-end router to match the local-end 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 under the BGP group context or under the neighbor context to the local-end address configured in 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.

4.4.8 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.

4.4.8.1 GRE SDP tunnel fragmentation

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

- SDP with type GRE (MPLS tunnels are not supported)

```
configure service sdp allow-fragmentation
```

- PW template with the **auto-gre-sdp** command option set

```
configure service pw-template allow-fragmentation
```

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

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

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 the **auto-gre-sdp** command option configured, 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.

4.4.8.2 GRE SDP tunnel reassembly

To reassemble GRE SDP tunnel fragments on the node terminating the service carried in the GRE SDP tunnel, the vISA-BB 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. Use the commands in the **configure card mda** context to enable the ISA-BB application on the node.

See the *7705 SAR Gen 2 Multiservice ISA and ESA Guide* for more information.

2. Use the following commands to configure an IP filter:

- **MD-CLI**

```
configure filter ip-filter default-action accept
configure filter ip-filter entry match protocol gre
configure filter ip-filter entry match fragment true
configure filter ip-filter entry action reassemble
```

- **classic CLI**

```
configure filter ip-filter default-action forward
configure filter ip-filter entry match protocol gre
```

```
configure filter ip-filter entry match fragment true
configure filter ip-filter entry action reassemble
```

3. Configure a NAT group using the following commands:

- **MD-CLI**

```
configure isa nat-group redundancy active-mda-limit 1
configure isa nat-group mda mda-BB-ISA
```

- **classic CLI**

```
configure isa nat-group active-mda-limit 1
configure isa nat-group mda mda-BB-ISA
```

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 using the following command:

- **MD-CLI**

```
configure router reassembly to-base-network
```

- **classic CLI**

```
configure router reassembly-group to-base-network
```

The following example displays a GRE SDP tunnel reassembly configuration.

Example: GRE SDP tunnel reassembly configuration (MD-CLI)

```
[ex:/configure card 1]
A:admin@node-2# info
  card-type imm-2pac-fp3
  mda 1 {
    admin-state enable
    mda-type isa2-bb
  }
```

```
[ex:/configure filter]
A:admin@node-2# info
  ip-filter "10" {
    default-action accept
    entry 1 {
      match {
        protocol gre
        fragment true
      }
      action {
        reassemble
      }
    }
  }
```

```
[ex:/configure isa]
A:admin@node-2# info
  nat-group 1 {
    admin-state enable
    redundancy {
```

```

        active-mda-limit 1
    }
    mda 1/2
}

[ex:/configure router "Base"]
A:admin@node-2# info
    interface "system" {
        admin-state enable
        ipv4 {
            primary {
                address 91.91.91.91
                prefix-length 32
            }
        }
    }
    interface "to-PE2" {
        port 1/1/1
        gre-termination true
        ingress {
            filter {
                ip "10"
            }
        }
        ipv4 {
            primary {
                address 10.1.1.10
                prefix-length 24
            }
        }
    }
}

```

Example: GRE SDP tunnel reassembly configuration (classic CLI)

```

A:node-2>config>card# info
-----
    card-type imm-2pac-fp3
    mda 1
        mda-type isa2-bb
        no shutdown
    exit
-----

```

```

A:node-2>config>filter# info
-----
...      ip-filter 10 name "10" create
        default-action forward
        entry 1 create
            match protocol gre
                fragment true
            exit
        action
            reassemble
        exit
    exit
...
-----

```

```

A:node-2>config>filter# info
-----
    ip-filter 10 name "10" create

```



```

    default-action forward
    entry 1 create
        match protocol gre
        fragment true
    exit
    action
        reassemble
    exit
exit
-----

```

```
A:node-2>config>isa# info
```

```

-----
...
nat-group 1 create
    no shutdown
    mda 1/2
    active-mda-limit 1
exit
-----

```

```
A:node-2>config>router# info
```

```

-----
interface "system"
    address 91.91.91.91/32
    no shutdown
exit
interface "to-PE2"
    address 10.1.1.10/24
    port 1/2/1
    ingress
        filter ip 10
    exit
    gre-termination
    no shutdown
exit
-----

```

4.4.8.3 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.

4.4.9 GRE SDP termination on router interface IP address

Use the following command to terminate GRE SDP tunnel packets directly on a router interface IP address:

- **MD-CLI**

```
configure router interface gre-termination
```

- **classic CLI**

```
configure router interface address gre-termination
```

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 command option must be set to the /31 value of the IP address of the router interface used to terminate the GRE SDP tunnel packets.

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 to terminate the GRE-SDP tunnel packets.

4.5 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.

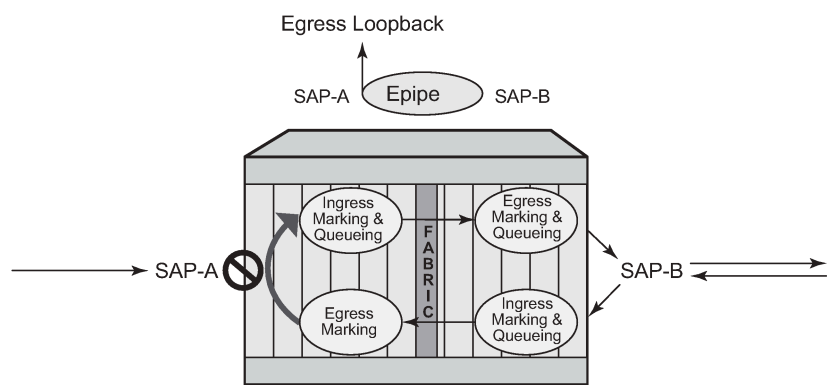
In the case of the egress loopback, the SAP or MPLS SDP binding is not isolated from the rest of the service it remains part of the service and reflects traffic back into the service.



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

Configuration		Reflection with inbound DA			
Action	Options	Unicast (learned)	Unicast (unknown)	Broadcast	Multicast
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 following command is configured under the non-operational SAP.

```
configure service epipe sap ignore-oper-down
```

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. Use the following command under the VPLS SAP over a LAG to allow the LAG SAP to be reached even with a non-operational SAP.

```
configure service vpls sap process-cpm-traffic-on-sap-down
```

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 commands in the **tools** context 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 example show output of the active loopback mode based on [Figure 18: Active loopback mode](#).

Output example: Active loopback mode configuration

```
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

CPM Ingress          : 491790      Packets
Octets               : 46721290

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

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

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

Use a **stop** command to stop a loopback.

Example

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

4.6 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** command 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** command 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** command option 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 option.

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 option. 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.

5 Multiservice sites

A customer site can be designated a multiservice 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 multiservice site can be on a single port or on a single slot. The SAPs in a multiservice site cannot span slots.

Multiservice 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 configuration defined in the policy. Multiservice 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 multiservice 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.

6 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** or **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.

7 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.

8 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. Weighted ECMP is supported where the service resolves directly to an ECMP set of RSVP or SR-TE LSPs with a configured load balancing weight, or where it resolves to a BGP tunnel and then uses an ECMP set of RSVP or SR-TE LSPs with a configured load balancing weight. Use the following commands to configure the weight for an LSP or LSP template.

```
configure router mpls lsp load-balancing-weight
configure router mpls lsp-template load-balancing-weight
```

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 *7705 SAR Gen 2 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, use the following command to configure weighted ECMP under the SDP used by the service.

```
configure service sdp weighted-ecmp
```

By default, weighted ECMP is disabled.

For VLL and VPLS services that use 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 configured in the following context.

```
configure service vpls load-balancing
configure service epipe load-balancing
```

Weighted ECMP is also supported for EVPN services. See the *7705 SAR Gen 2 Layer 2 Services and EVPN Guide* for more information.

9 NGE

This section provides information to configure Network Group Encryption (NGE).

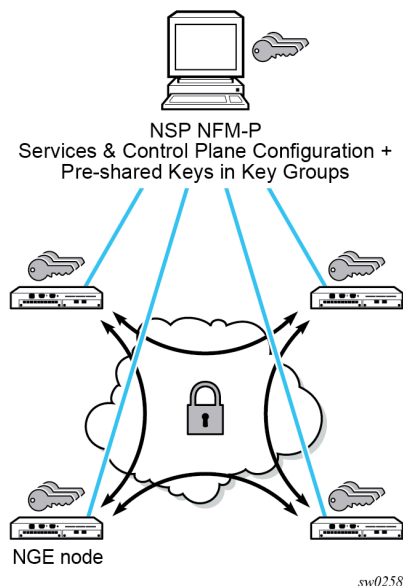
9.1 NGE overview

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 19: NGE network with NSP NFM-P management](#) shows an NGE network with NSP NFM-P services, control plane configuration, and key management.

Figure 19: 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 *7705 SAR Gen 2 Router Configuration Guide* for information about configuring NGE on router interfaces.

9.1.1 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 20: 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 20: Key group partitioning



sw0250

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.

9.1.2 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*.

9.2 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 21: 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 21: Key groups and a typical NGE packet



sw0251

9.2.1 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.

9.2.1.1 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)

9.2.2 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 21: Key groups and a typical NGE packet](#) illustrates these relationships.

As shown in [Figure 21: 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 you run an **info** command. 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).

9.2.2.1 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.

9.3 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
- 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).

9.3.1 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 22: NGE MPLS/GRE/MPLSoUDP label stack](#) illustrates the NGE MPLS, GRE, or MPLSoUDP label stack.

Figure 22: NGE MPLS/GRE/MPLSoUDP label stack

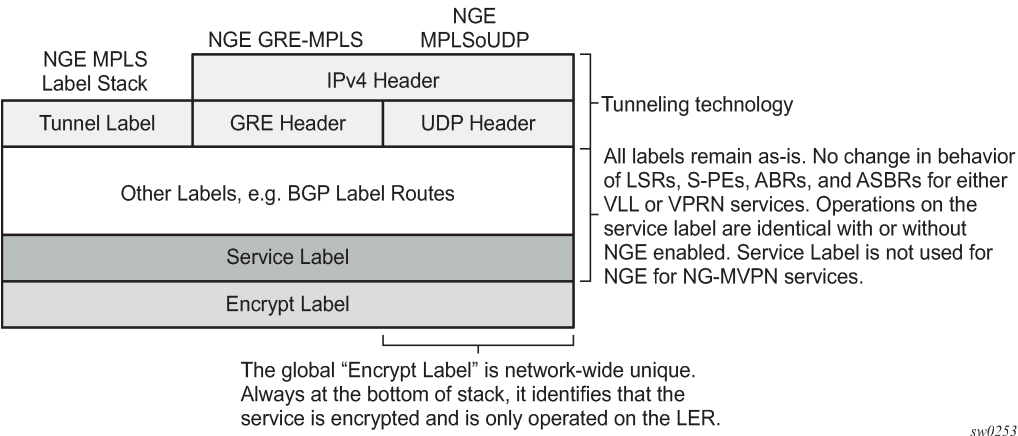
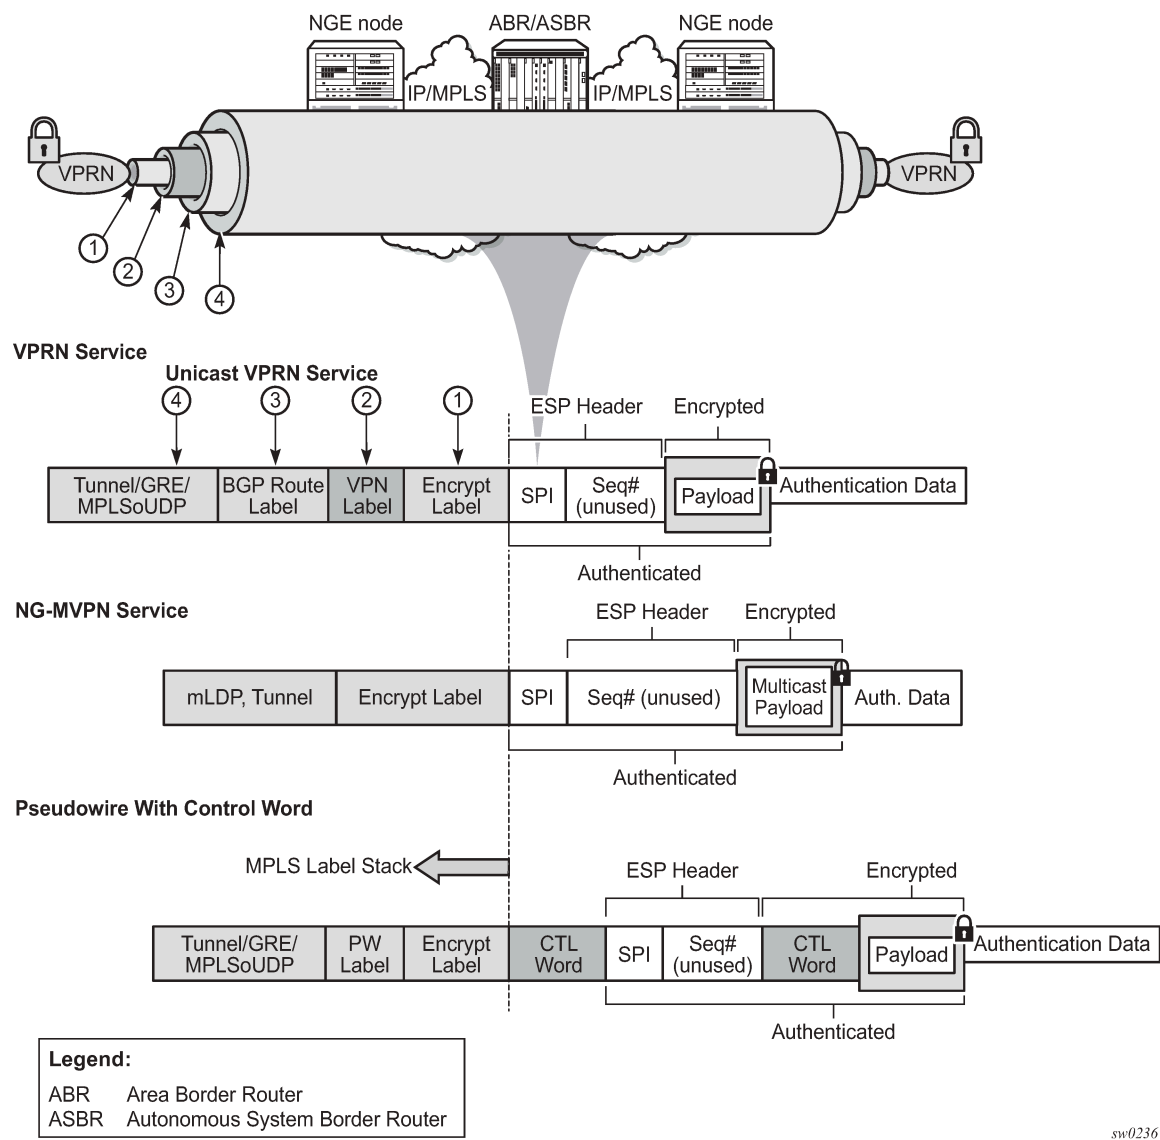


Figure 23: NGE and packet formats illustrates VPRN and PW (with control word) packet formats using NGE.

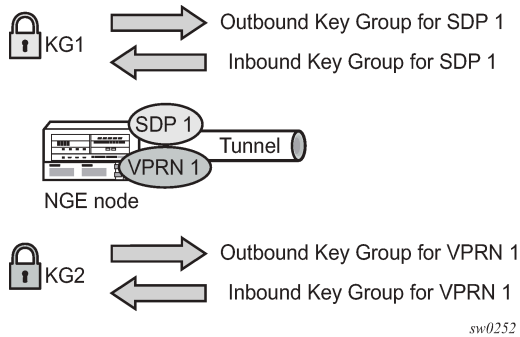
Figure 23: NGE and packet formats



9.3.2 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 24: Inbound and outbound key group assignments](#)).

Figure 24: 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.

9.3.3 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 with spoke SDPs (**spoke-sdp**) or auto-bind SDPs (**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 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.

9.3.4 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. Use the commands in the following context to bind a key group to a PW template for inbound or outbound packet processing.

```
configure service pw-template encryption-keygroup
```

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

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

9.3.5 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 *7705 SAR Gen 2 OAM and Diagnostics Guide* for more information.

9.3.6 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 the **control-word** command 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.

9.3.7 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).

9.3.8 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 *7705 SAR Gen 2 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.

9.4 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

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

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

Table 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 *7705 SAR Gen 2 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:

Service type	MTU values with and without NGE enabled
	<ul style="list-style-type: none"> Layer 3 spoke IP MTU (MPLS) <ul style="list-style-type: none"> = 1469 – 14 (inner Ethernet header) = 1455 bytes PW spoke SDP MTU (MPLS) <ul style="list-style-type: none"> = SDP MTU = 1469 bytes
GRE-based services	SDP MTU (GRE): = 1572 (network port MTU) – 14 (Ethernet header) – 20 (IP header) – 4 (GRE header) – 4 (inner label) = 1530 bytes (without NGE enabled) => 1449 bytes (with NGE enabled)
	Service MTU (GRE) considerations with NGE enabled: <ul style="list-style-type: none"> Layer 3 Spoke IP MTU (GRE) <ul style="list-style-type: none"> = 1449 – 14 (inner Ethernet header) = 1435 bytes PW Spoke MTU (GRE) <ul style="list-style-type: none"> = 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.

9.5 Statistics

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

- key group
- SPI
- MDA
- service

9.6 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.

9.7 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 *7705 SAR Gen 2 Layer 3 Services Guide: IES and VPRN* for more information.
- The NSP NFM-P provides command options that are not configurable using the CLI. See [Network services platform management](#) for more information.

9.7.1 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.

9.7.2 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.

9.7.3 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.

9.7.4 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.

9.7.5 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.

9.7.6 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.

10 Service creation process overview

The following figure shows the overall process to provision core and subscriber services.

Figure 25: Service creation and implementation flow



sw0148

11 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.

11.1 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](#)

11.2 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, user 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](#)

11.3 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](#).

12 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 Epipe, IES, VPLS, or VPRN services on the 7705 SAR Gen 2.
2. Bind SDPs.
3. Configure interfaces (where required) and SAPs.
4. Create exclusive QoS and filter policies.

13 Configuring global service entities with CLI

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

13.1 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 (customers)
- SDPs
- services
 - **Ethernet Pipe (Epipe) services**
See the *7705 SAR Gen 2 Layer 2 Services and EVPN Guide* for more information.
 - **Virtual Private LAN Service (VPLS)**
See the *7705 SAR Gen 2 Layer 2 Services and EVPN Guide* for more information.
 - **Internet Enhanced Service (IES)**
See the *7705 SAR Gen 2 Layer 3 Services Guide: IES and VPRN* for more information.
 - **Virtual Private Routed Network (VPRN) service**
See the *7705 SAR Gen 2 Layer 3 Services Guide: IES and VPRN* for more information.
- Service Access Points (SAPs)
 - **Ethernet Pipe (Epipe) services**
See the *7705 SAR Gen 2 Layer 2 Services and EVPN Guide* for more information.
 - **VPLS SAP**
See the *7705 SAR Gen 2 Layer 2 Services and EVPN Guide* for more information.
 - **IES SAP**
See the *7705 SAR Gen 2 Layer 3 Services Guide: IES and VPRN* for more information.
 - **VPRN Interface SAP**
See the *7705 SAR Gen 2 Layer 3 Services Guide: IES and VPRN* for more information.

13.2 Basic configuration

The most basic service configuration must have the following:

- a customer ID
- a service type

- a service ID

This is an optional service name that can 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 associated SDP for distributed services
- an interface (where required) identifying an IP address, IP subnet, and broadcast address
- for distributed services, an associated SDP

The commands in the following contexts provide nodes for basic configurations.

```
configure service customer
configure service sdp
configure service service type
```

13.3 Common configuration tasks

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

13.3.1 Configuring customers

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

- description
- contact name
- telephone number
- multiservice site

13.3.1.1 Customer information

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

```
A:node-2config>service# info
-----
...
    customer 5
      description "Nokia Customer"
      contact "Technical Support"
      phone "650 555-1212"
    exit
...
-----
```

13.3.1.2 Configuring multiservice-sites

Multiservice 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 multiservice customer site association must be removed from the SAP before local scheduler policies may be applied.

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

Use the following CLI command to configure customer multiservice site information.

```
configure service customer multi-service-site
```

13.3.2 Configuring an SDP

The most basic SDP must have the following:

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

13.3.2.1 SDP configuration tasks

About this task

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

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.

13.3.2.2 Configuring a mixed-LSP SDP

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

- **MD-CLI**

```
configure service sdp delivery-type mpls mixed-lsp-mode
```

- **classic CLI**

```
configure 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 the time 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:

- **MD-CLI**

```
configure service sdp delivery-type mpls mixed-lsp-mode revert-time
```

- **classic CLI**

```
configure service sdp mpls mixed-lsp-mode revert-time
```

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

- **MD-CLI**

```
configure service sdp delivery-type mpls mixed-lsp-mode revert-time never
```

- **classic CLI**

```
configure service sdp mpls mixed-lsp-mode revert-time infinite
```

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

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:

With 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 reprograms the line card with the LDP LSP if available. If not, the SDP goes operationally down.

When a higher priority type LSP becomes available, the service manager reverts back to this LSP at the expiry of the revert time timer or the failure of the currently active LSP, whichever comes first. The service manager then re-programs the line card accordingly.

The SDP reverts to the highest priority type LSP only if the currently active LSP failed when the following commands were configured:

- **MD-CLI**

```
configure service sdp delivery-type mpls mixed-lsp-mode revert-time never
```

- **classic CLI**

```
configure service sdp mpls mixed-lsp-mode revert-time infinite
```



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 **configure router ldp tunnel-down-damp-time** command.

If the value of the **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.

With LDP/BGP SDP, the service manager prefers the LDP LSP type over the BGP LSP type. The service manager re-programs the line card with the BGP LSP if available otherwise it brings down the SDP operationally.

The following are differences in behavior between the LDP/BGP SDP and the 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 labeled 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.

14 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.

14.1 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.

14.2 Configuring NGE components

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

14.2.1 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 command to configure the global encryption label.

```
configure group-encryption group-encryption-label
```

14.2.2 Configuring a key group

To configure a key group, set the following command options:

- 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 cleartext using the following command:

```
configure group-encryption encryption-keygroup security-association
```

Once entered, keys 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 commands in the following context to configure key group options.

```
configure group-encryption encryption-keygroup
```

The following example displays the key group configuration.

Example: MD-CLI

```
[ex:/configure group-encryption]
A:admin@node-2# info detail
  group-encryption-label 34
  encryption-keygroup 2 {
    description "Main_secure_KG"
    keygroup-name "KG1_secure"
    authentication-algorithm sha256
    encryption-algorithm aes128
    active-outbound-security-association 6
    security-association 2 {
      authentication-key "0XLyKVjy88fjyz0FGgpoklHAPB8344vN42vv6LMY5Zy1e08aiZe2CLa
LstrqXQaw" hash2
      encryption-key "bxYkRG2enIPs85zNMSDhX1BzGMaro8TAIFrwcysTRf8= hash2"
    }
    security-association 6 {
      authentication-key "RmzyeCJNICozfGXXQ4jfbQ1zRbW6nf5GcjTuCYSjQCAri1ufVhABj9No
Zqcmtwb8" hash2
      encryption-key "jWYIDRE0Td3jeViBBprxGQ4Dixn87UypaM1dNosk7Iw= hash2"
    }
  }
```

Example: classic CLI

```
A:node-2>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
4723a7d1d5f7b6ffa6fc encryption-
key 0x8d51db8f826239f672457442cecc73665f52cbe00aedfb4eda6166001247b4eb crypto
security-association spi 6 authentication-key 0x7fb9fc5553630924ee29973f
7b0a48f801b0ae1cb38b7666045274476a268e8d694ab6aa7ea050b7a43cdf8d80977625 encryption-
key 0x72bd9b87841dbecb2d114031367ab5d9153a41b7c79c8f889ac56b950d8fffa crypto
active-outbound-sa 6
exit
-----

```

14.2.3 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



Note: Key groups can only be assigned to SDPs or VPRNs using the classic CLI commands.

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 an auto-created 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 or auto-bind SDPs is supported using LDP, GRE, RSVP-TE, or segment routing (SR-ISIS, SR-OSPF, or SR-TE).

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

- **MD-CLI**

```

configure service pw-template encryption-keygroup inbound
configure service pw-template encryption-keygroup outbound
configure service vprn subscriber-interface group-interface wlan-gw group-encryption
encryption-keygroup-inbound

```


- **classic CLI**

```
configure service sdp encryption-keygroup direction {inbound | outbound}  
configure service vprn encryption-keygroup direction {inbound | outbound}  
configure service pw-template encryption-keygroup direction {inbound | outbound}
```



Note: After assigning a key group to the PW template, execute the following command:

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

15 Global service entity management tasks

This section describes global service entity management tasks.

15.1 Modifying customer accounts

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

Use the following command to display a list of customer IDs.

```
show service customer
```

Use the commands in the following context to edit customer information, such as description, contact, or phone number.

```
configure service customer
```

15.2 Deleting customers

Deleting a customer 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.

Use the following command to delete a customer:

- **MD-CLI**

```
configure service delete customer
```

- **classic CLI**

```
configure service no customer
```

15.3 Modifying SDPs

To access a specific SDP, specify the SDP ID. Use the following command to display a list of SDPs.

```
show service sdp
```

Modify the SDP command options, such as **description** or **far-end** by entering the information in the following context:



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

```
configure service sdp
```

15.4 Deleting SDPs

Deleting an SDP 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.

Use the following command to delete an SDP:

- **MD-CLI**

```
configure service delete sdp
```

- **classic CLI**

```
configure service no sdp
```

16 NGE management tasks

This section describes NGE management tasks.

16.1 Modifying a key group



Note: The following conditions apply for the classic CLI.

When modifying a key group, the user must adhere to 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 uses the key group.

The following example displays the modification of a key group using the following steps:

1. In the classic CLI, the active outgoing SA is deconfigured.
2. In the classic CLI, the SAs are removed.
3. The encryption algorithm is changed.
4. The SAs are reconfigured.
5. The active outgoing SA is reconfigured.

Example: MD-CLI

```
*[ex:/configure group-encryption]
A:admin@node-2# encryption-keygroup 1

*[ex:/configure group-encryption encryption-keygroup 1]
A:admin@node-2# encryption-algorithm aes256

*[ex:/configure group-encryption encryption-keygroup 1]
A:admin@node-2# security-association 6 authentication-key
0x666666660000000000666666660000000066666666000000006666666600000000

*[ex:/configure group-encryption encryption-keygroup 1]
A:admin@node-2# security-association 6 encryption-key
0x666666660000000000666666660000000066666666000000006666666600000000

*[ex:/configure group-encryption encryption-keygroup 1]
A:admin@node-2# active-outbound-security-association 6
```

Example: classic CLI

```
*A:node-2>config>grp-encryp# encryption-keygroup 1
*A:node-2>config>grp-encryp>encryp-keygrp# no active-outbound-sa
```

```
*A:node-2>config>grp-encryp>encryp-keygrp# no security-association spi 6
*A:node-2>config>grp-encryp>encryp-keygrp# esp-encryption-algorithm aes256
*A:node-2>config>grp-encryp>encryp-keygrp# security-association spi 6 authentication-
key 0x6666666600000000000000000000000000000000000000000000000000000000 encryption-key
0x6666666600000000000000000000000000000000000000000000000000000000
*A:node-2>config>grp-encryp>encryp-keygrp# active-outbound-sa 6
```

16.2 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.

16.2.1 Removing a key group from an SDP, VPRN service, or PW template

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



Note: Key groups can only be assigned to SDPs or VPRNs using the classic CLI commands.

- **MD-CLI**

```
configure service pw-template delete encryption-keygroup inbound
configure service pw-template delete encryption-keygroup outbound
```

- **classic CLI**

```
configure service sdp no encryption-keygroup direction {inbound | outbound}
configure service vprn no encryption-keygroup direction {inbound | outbound}
configure service pw-template no encryption-keygroup direction {inbound | outbound}
```



Note: After removing a key group to the PW template, the following command must be executed.

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

16.3 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.

16.3.1 Changing the key group for an SDP, VPRN service, PW template, or WLAN-GW group interface

Changing key groups for an SDP, VPRN service, PW template, or WLAN-GW group interface must be performed on all nodes for the service.

To change the key group on an SDP, VPRN service, PW template, or WLAN-GW group interface, perform the task as described in: [Changing key groups](#).



Note: Key groups can only be changed on SDPs and VPRNs using the classic CLI commands.



Note: For PW template changes, the following command must be executed after the changes are made.

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

16.4 Deleting a key group from an NGE node

To delete a key group from an NGE node, the key group must be removed (unbound) from all SDPs, VPRN services, PW templates, and router interfaces that use it.



Note: When deleting a key group from a PW template, the following command must be executed after the encryption keygroup changes are made.

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

Use the following command to locate the key group bindings.

```
show group-encryption encryption-keygroup
```

Use the following command to delete a key group:

- **MD-CLI**

```
configure group-encryption delete encryption-keygroup
```

- **classic CLI**

```
configure group-encryption no encryption-keygroup
```

17 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.

17.1 Bidirectional Forwarding Detection (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)*

17.2 Border Gateway Protocol (BGP)

draft-hares-idr-update-attr-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-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-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-link-bandwidth-03, *BGP Link Bandwidth Extended Community*

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 6368, *Internal BGP as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*

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 7752, *North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP*

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 9494, *Long-Lived Graceful Restart for BGP*

17.3 Bridging and management

IEEE 802.1AB, *Station and Media Access Control Connectivity Discovery*

IEEE 802.1ad, *Provider Bridges*

IEEE 802.1AX, *Link Aggregation*

IEEE 802.1Q, *Virtual LANs*

17.4 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*

17.5 Ethernet VPN (EVPN)

draft-ietf-bess-bgp-srv6-args-00, *SRv6 Argument Signaling for BGP Services*
draft-ietf-bess-evpn-ipvpn-interworking-06, *EVPN Interworking with IPVPN*
draft-sr-bess-evpn-vpws-gateway-03, *Ethernet VPN Virtual Private Wire Services Gateway Solution*
RFC 7432, *BGP MPLS-Based Ethernet VPN*
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 9047, *Propagation of ARP/ND Flags in an Ethernet Virtual Private Network (EVPN)*
RFC 9135, *Integrated Routing and Bridging in Ethernet VPN (EVPN)*
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)*

17.6 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 0.8.0, *gRPC Network Management Interface (gNMI) Service Specification*
gnmi_ext.proto version 0.1.0, *gNMI Commit Confirmed Extension*
system.proto version 1.0.0, *gRPC Network Operations Interface (gNOI) System Service*
PROTOCOL-HTTP2, *gRPC over HTTP2*

17.7 Intermediate System to Intermediate System (IS-IS)

draft-ietf-isis-mi-02, *IS-IS Multi-Instance*

draft-ietf-lsr-igp-ureach-prefix-announce-01, *IGP Unreachable Prefix Announcement* – without U-Flag and UP-Flag

draft-kaplan-isis-ext-eth-02, *Extended Ethernet Frame Size Support*

ISO/IEC 10589:2002 Second Edition, *Intermediate system to Intermediate system intra-domain routing 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 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 5308, *Routing IPv6 with IS-IS*

RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*

RFC 5310, *IS-IS Generic Cryptographic Authentication*

RFC 6213, *IS-IS BFD-Enabled TLV*

RFC 6232, *Purge Originator Identification TLV for IS-IS*

RFC 6233, *IS-IS Registry Extension for Purges*

RFC 7775, *IS-IS Route Preference for Extended IP and IPv6 Reachability*

RFC 7794, *IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability* – sections 2.1 and 2.3

RFC 7981, *IS-IS Extensions for Advertising Router Information*

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*

17.8 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 2784, *Generic Routing Encapsulation (GRE)*

RFC 3164, *The BSD syslog Protocol*

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

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

17.9 Internet Protocol (IP) multicast

RFC 1112, *Host Extensions for IP Multicasting*

RFC 2236, *Internet Group Management Protocol, Version 2*

RFC 2365, *Administratively Scoped IP Multicast*

RFC 2375, *IPv6 Multicast Address Assignments*

RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*

RFC 3376, *Internet Group Management Protocol, Version 3*

RFC 3446, *Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*

RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*

RFC 3618, *Multicast Source Discovery Protocol (MSDP)*

RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*

RFC 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 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 7761, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*
RFC 8487, *Mtrace Version 2: Traceroute Facility for IP Multicast*

17.10 Internet Protocol (IP) version 4

RFC 791, *Internet Protocol*
RFC 792, *Internet Control Message Protocol*
RFC 826, *An Ethernet Address Resolution Protocol*
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 1812, *Requirements for IPv4 Routers*
RFC 1918, *Address Allocation for Private Internets*
RFC 2131, *Dynamic Host Configuration Protocol*; Relay only
RFC 2132, *DHCP Options and BOOTP Vendor Extensions* – DHCP
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

17.11 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 3736, *Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6*
RFC 3971, *SEcure Neighbor Discovery (SEND)*
RFC 4007, *IPv6 Scoped Address Architecture*
RFC 4191, *Default Router Preferences and More-Specific Routes* – Default Router Preference
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 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 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*
RFC 8021, *Generation of IPv6 Atomic Fragments Considered Harmful*
RFC 8200, *Internet Protocol, Version 6 (IPv6) Specification*

17.12 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)*
RFC 8784, *Mixing Preshared Keys in the Internet Key Exchange Protocol Version 2 (IKEv2) for Post-quantum Security*

17.13 Label Distribution Protocol (LDP)

draft-pdutta-mpls-ldp-adj-capability-00, *LDP Adjacency Capabilities*
draft-pdutta-mpls-ldp-v2-00, *LDP Version 2*
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*

17.14 Multiprotocol Label Switching (MPLS)

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 5332, *MPLS Multicast Encapsulations*

RFC 5884, *Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)*

RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*

RFC 7308, *Extended Administrative Groups in MPLS Traffic Engineering (MPLS-TE)*

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

17.15 Network Address Translation (NAT)

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*

17.16 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 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 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 5642, *Dynamic Hostname Exchange Mechanism for OSPF*

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 8920, *OSPF Application-Specific Link Attributes*

17.17 Path Computation Element Protocol (PCEP)

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

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

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

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

RFC 8253, *PCEPS: Usage of TLS to Provide a Secure Transport for the Path Computation Element Communication Protocol (PCEP)*

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*

17.18 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*

17.19 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)*

17.20 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 3162, *RADIUS and IPv6*
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*

17.21 Resource Reservation Protocol - Traffic Engineering (RSVP-TE)

RFC 2702, *Requirements for Traffic Engineering over MPLS*
RFC 2747, *RSVP Cryptographic Authentication*
RFC 2961, *RSVP Refresh Overhead Reduction Extensions*
RFC 3097, *RSVP Cryptographic Authentication -- Updated Message Type Value*

RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*
RFC 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 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks*

17.22 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*

17.23 Segment Routing (SR)

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 8426, *Recommendations for RSVP-TE and Segment Routing (SR) Label Switched Path (LSP) Coexistence*
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 8665, *OSPF Extensions for Segment Routing*
RFC 8667, *IS-IS Extensions for Segment Routing*
RFC 8669, *Segment Routing Prefix Segment Identifier Extensions for BGP*
RFC 9256, *Segment Routing Policy Architecture*
RFC 9350, *IGP Flexible Algorithm*

17.24 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-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

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 2206, RSVP Management Information Base using SMIv2

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 3273, Remote Network Monitoring Management Information Base for High Capacity Networks

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

RFC 3430, Simple Network Management Protocol (SNMP) over Transmission Control Protocol (TCP) Transport Mapping

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 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework*

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 3826, *The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model*

RFC 3877, *Alarm Management Information Base (MIB)*

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 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 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*

17.25 Timing

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

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

RFC 8573, *Message Authentication Code for the Network Time Protocol*

17.26 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*

RFC 9503, *Simple Two-Way Active Measurement Protocol (STAMP) Extensions for Segment Routing Networks – excluding Sections 3, 4.1.2 and 4.1.3*

17.27 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 7117, *Multicast in Virtual Private LAN Service (VPLS)*

17.28 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*

Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)