



7750 Service Router Virtualized Service Router Release 25.3.R1

BNG CUPS User Plane Function Guide

3HE 21195 AAAA TQZZA 01
Edition: 01
March 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

1	Getting started.....	6
1.1	About this guide.....	6
1.2	Conventions.....	6
1.2.1	Precautionary and information messages.....	7
1.2.2	Options or substeps in procedures and sequential workflows.....	7
2	CUPS UP configuration quick start.....	8
2.1	Configuration overview.....	8
2.2	Configuring the CUPS UP for subscriber management.....	9
3	PFCP association.....	14
3.1	BNG-UP PFCP association.....	14
3.2	Multiple PFCP associations for FWA.....	16
3.3	PFCP heartbeats and headless mode.....	16
3.4	Default IBCP session.....	17
3.5	Overload and prioritization.....	18
3.6	Operational commands and debugging.....	19
4	Session management.....	20
4.1	Subscribers, QoS, and filters.....	20
4.2	IBCP.....	20
4.3	IP gateway, services, and routing.....	21
4.4	Statistics reporting.....	22
4.4.1	Selective aggregate statistics.....	23
4.4.2	Custom statistics reporting.....	26
4.5	Threshold and quota monitoring.....	27
4.6	Operational commands.....	28
4.7	SAP and group interface templates.....	28
4.7.1	Mixing different encapsulation sessions on the same port.....	29
4.8	Fixed access sessions.....	30
4.9	Fixed wireless access sessions.....	31
4.9.1	Configuring FWA sessions.....	31
4.9.2	Dynamic QoS based on PCC rules.....	32
4.9.2.1	Overview of dynamic QoS based on PCC rules.....	32

4.9.2.2	Using alternate QoS profiles in combination with dynamic PCC rules.....	33
4.9.2.3	Configuring profiles for dynamic object creation.....	34
4.9.2.4	Examples of profile configuration for dynamic object creation.....	36
4.10	Routed subscriber sessions.....	39
4.11	Call trace.....	40
5	Network Address Translation.....	42
5.1	Residential NAT for BNG CUPS.....	42
5.2	UP NAT policy template.....	42
5.3	Guidelines for configuring extended port blocks.....	43
5.4	Guidelines for configuring NAT subscribers in the sub-profile.....	43
5.5	Guidelines for configuring NAT groups.....	44
5.6	Guidelines for configuring accounting and logging.....	44
5.7	Guidelines for configuring watermarks.....	44
5.8	Guidelines for configuring intra-chassis redundancy.....	45
5.9	Provisioning residential NAT for BNG CUPS.....	46
6	BNG-UP resiliency.....	47
6.1	Resiliency based on Fate Sharing Group.....	47
6.2	BNG-UP health reporting.....	49
6.3	Interaction with headless mode.....	51
7	Layer 2 Tunneling Protocol.....	53
7.1	BNG-UP-triggered L2TP access concentrator.....	53
8	Lawful intercept.....	55
8.1	Overview of the LI implementation on the BNG-UP.....	55
8.2	Provisioning SNMPv3 and LI subscribers for the BNG-UP.....	56
9	Standards and protocol support.....	57
9.1	Access Node Control Protocol (ANCP).....	57
9.2	Bidirectional Forwarding Detection (BFD).....	57
9.3	Border Gateway Protocol (BGP).....	57
9.4	Bridging and management.....	59
9.5	Broadband Network Gateway (BNG) Control and User Plane Separation (CUPS).....	60
9.6	Certificate management.....	60
9.7	Circuit emulation.....	61

9.8	Ethernet.....	61
9.9	Ethernet VPN (EVPN).....	61
9.10	gRPC Remote Procedure Calls (gRPC).....	62
9.11	Intermediate System to Intermediate System (IS-IS).....	62
9.12	Internet Protocol (IP) Fast Reroute (FRR).....	63
9.13	Internet Protocol (IP) general.....	63
9.14	Internet Protocol (IP) multicast.....	65
9.15	Internet Protocol (IP) version 4.....	66
9.16	Internet Protocol (IP) version 6.....	67
9.17	Internet Protocol Security (IPsec).....	68
9.18	Label Distribution Protocol (LDP).....	69
9.19	Layer Two Tunneling Protocol (L2TP) Network Server (LNS).....	70
9.20	Multiprotocol Label Switching (MPLS).....	70
9.21	Multiprotocol Label Switching - Transport Profile (MPLS-TP).....	71
9.22	Network Address Translation (NAT).....	71
9.23	Network Configuration Protocol (NETCONF).....	72
9.24	Open Shortest Path First (OSPF).....	72
9.25	OpenFlow.....	73
9.26	Path Computation Element Protocol (PCEP).....	73
9.27	Point-to-Point Protocol (PPP).....	74
9.28	Policy management and credit control.....	74
9.29	Pseudowire (PW).....	74
9.30	Quality of Service (QoS).....	75
9.31	Remote Authentication Dial In User Service (RADIUS).....	75
9.32	Resource Reservation Protocol - Traffic Engineering (RSVP-TE).....	76
9.33	Routing Information Protocol (RIP).....	76
9.34	Segment Routing (SR).....	76
9.35	Simple Network Management Protocol (SNMP).....	78
9.36	Timing.....	80
9.37	Two-Way Active Measurement Protocol (TWAMP).....	81
9.38	Virtual Private LAN Service (VPLS).....	81
9.39	Voice and video.....	81
9.40	Yet Another Next Generation (YANG).....	82
9.41	Yet Another Next Generation (YANG) OpenConfig Models.....	82

1 Getting started

1.1 About this guide

This guide describes the Nokia SR OS User Plane Function (UPF) in a Broadband Network Gateway (BNG) Control and User Plane Separation (CUPS) system. The BNG-UP is based on the fundamentals of the BNG and reuses fundamental QoS and forwarding concepts, while moving the control plane handling to the MAG-c.

See the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide* for more information about BNG.

See the guides listed in *MAG-c Guide to Documentation* for more information about the BNG CUPS solution and the MAG-c.

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

The topics and commands described in this document apply to the following SR OS products:

- 7750 SR
- Virtualized Service Router

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



Note: The BNG-UP supports configuration using the classic CLI and the MD-CLI. This guide provides configuration examples based on the MD-CLI syntax only.

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

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



Note: This guide generically covers Release 25.x.Rx content and may contain some content that will be released in later maintenance loads. For information about features supported in each load of the Release 25.x.Rx software or for a list of unsupported features by platform and chassis, see the *SR OS R25.x.Rx Software Release Notes*, part number 3HE 21562 000x TQZZA.

1.2 Conventions

This section describes the general conventions used in this guide.

1.2.1 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.

1.2.2 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 CUPS UP configuration quick start

For a CUPS deployment scenario, configure the CUPS user plane (UP) for subscriber management using fixed access or fixed wireless access (FWA) sessions, in conjunction with the MAG-c control plane (CP).

See the following for information about the configuration tasks:

- [Configuration overview](#) provides a summary of the CUPS UP configuration tasks.
- The *MAG-c Control Plane Function Guide* provides information about configuring the MAG-c CP.

2.1 Configuration overview

The following table lists the main tasks required to configure the CUPS UP for subscriber management.

Table 1: Tasks for configuring the CUPS UP

Configuration component	Configuration task
Configure the basic system components, including cards, ports, users, and configuration modes.	See the following guides: <i>7450 ESS, 7750 SR, 7950 XRS, and VSR Basic System Configuration Guide</i> <i>7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide</i>
Configure the IP connectivity from the CUPS UP to the MAG-c.	See Configuring the CUPS UP for subscriber management , step 1.
Configure the PFCP association.	See Configuring the CUPS UP for subscriber management , step 2.
Configure the fixed access or FWA services.	See Configuring the CUPS UP for subscriber management , step 3.
Configure the network services.	See Configuring the CUPS UP for subscriber management , step 4.
Configure the profiles and templates for subscribers, SLAs, SAPs, and group interfaces.	See Configuring the CUPS UP for subscriber management , step 5.

2.2 Configuring the CUPS UP for subscriber management

Prerequisites

- Ensure the basic system components are configured, including the cards, ports, users, and configuration modes. See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Basic System Configuration Guide*.
- Familiarize yourself with the steps and examples for configuring IP connectivity in the base router and VPRNs. See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide* for more information.
- Familiarize yourself with the steps and examples for configuring the MAG-c CP. See the *MAG-c Control Plane Function Guide* for more information.

About this task

To configure the CUPS UP for subscriber management (in conjunction with the MAG-c CP), using fixed access or FWA sessions, you must configure IP connectivity, PFCP association, access and network services, subscriber and SLA profiles, and SAP and group interface templates.

Perform all the steps in the order specified. These steps include links to additional topics, if applicable, for users who require more detailed information.

Procedure

Step 1. Configure IP connectivity to the MAG-c function.

The CUPS UP requires an interface (preferably a loopback interface), in the base router instance or a VPRN, to terminate and originate PFCP traffic with the MAG-c.

The following example displays a loopback interface named "pfcpc_endpoint" in the base router configuration that uses the IPv4 address 192.0.2.1 for PFCP communication.

Example

Loopback interface configuration with IPv4 address for PFCP communication

```
[ex:/configure router "Base"]
A:admin@UP-West# info
  interface "pfcpc_endpoint" {
    loopback
    ipv4 {
      primary {
        address 192.0.2.1
        prefix-length 32
      }
    }
  }
}
```

Step 2. Configure the PFCP association.

The PFCP association manages all communication-related subscriber states between the UP and the CP.

The following example displays the connection to a MAG-c listening on IP 1.1.1.1, using the interface configured in step 1. The **node-id** command is configured to identify this UP on the MAG-c.

See [PFCP association](#) for more information about the PFCP configuration.

Example**PFCP association configuration for connection to the MAG-c**

```
[ex:/configure subscriber-mgmt pfcpl association "MAG-c"]
A:admin@UP-West# info
  admin-state enable
  node-id {
    fqdn "up-west"
  }
  interface {
    router-instance "Base"
    name "pfcpl_endpoint"
  }
  peer {
    ip-address 1.1.1.1
  }
}
```

Step 3. Configure the access services.

- **Fixed access sessions**

The access service is a VPLS with capture SAPs configured for the port, and the VLAN contexts in which subscribers are terminated.

See [Fixed access sessions](#) for more information.

- **FWA sessions**

The access service is a VPRN or the base router. In the VPRN, an interface (preferably a loopback interface) terminates GTP-u traffic. GTP-u traffic is sent with a source IP address of this interface, and follows regular routing. You must also configure a forwarding path extensions (FPE) construct in the same access service to terminate FWA subscribers.

See [Fixed wireless access sessions](#) for more information.

Example**Access services configuration for a fixed access session**

This example displays the following configurations:

- A single capture SAP terminates subscribers on port 1/1/c1/1.
- The capture SAP uses wildcards for the S-tag and C-tag to allow any S-tag and C-tag.
- To start CUPS functionality and avoid treating the sessions like integrated BNG sessions, the capture SAP is linked to the PFCP association created in step 2 .
- The trigger-packet configuration determines which packet is sent to the CP for initial session setup.

```
[ex:/configure service vpls "fixed_access"]
A:admin@UP-West# info
  admin-state enable
  customer "1"
  capture-sap 1/1/c1/1:*.* {
    pfcpl {
      association "MAG-c"
    }
    trigger-packet {
      dhcp true
      dhcp6 true
      pppoe true
      rtr-solicit true
    }
  }
}
```

```
}
}
```

Example

Access services configuration for an FWA session

This example displays the following configurations:

- GTP-u traffic terminates on an FPE linked to an internal PXC on MAC 1 and MAC 4 of card and MDA 1/1.
- Listening is configured on IPv4 172.20.1.2 or IPv6 2001:db8:2020::2 in a VPRN named "to_ran".

Other VPRN configurations, such as the interfaces to the RAN network, are not shown.



Note: The UP automatically balances FWA sessions over both PXC.

See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide* for more information about the following topics:

- See "Port cross-connect", for information about how to configure port cross-connects.
- See "FPE", for information about configuring FPEs.

```
[ex:/configure card 1 mda 1]
A:admin@UP-West# info
  xconnect {
    mac 1 {
      loopback 1 {
      }
    }
    mac 4 {
      loopback 1 {
      }
    }
  }
}

[ex:/configure port 1/1/m1/1]
A:admin@UP-West# info
  admin-state enable

[ex:/configure port 1/1/m4/1]
A:admin@UP-West# info
  admin-state enable

[ex:/configure port-xc]
A:admin@UP-West# info
  pxc 1 {
    admin-state enable
    port-id 1/1/m1/1
  }
  pxc 2 {
    admin-state enable
    port-id 1/1/m4/1
  }
}

[ex:/configure fwd-path-ext fpe 1]
A:admin@UP-West# info
  multi-path {
    path 1 {
      pxc 1
    }
  }
}
```

```

    path 2 {
        pxc 2
    }
}
application {
    sub-mgmt-extension true
}

[ex:/configure service vprn "to_ran"]
A:admin@UP-West# info
    gtp {
        upf-data-endpoint {
            interface "gtp_u_endpoint"
            fpe 1
        }
    }
    interface "gtp_u_endpoint" {
        loopback true
        ipv4 {
            primary {
                address 172.20.1.2
                prefix-length 32
            }
        }
        ipv6 {
            address 2001:db8:2020::2 {
                prefix-length 128
            }
        }
    }
}

```

Step 4. Configure the network services.

Fixed access and FWA sessions use IES or VPRN services to provide VPRN connectivity. You do not require a specific subscriber-management configuration; the service must only exist and be referenced by the MAG-c when installing the session.

See [IP gateway, services, and routing](#) for more information about how to use these services.

Step 5. Configure default profiles.

CUPS sessions require the following explicitly-configured profiles and templates:

- SLA profiles
- subscriber profiles
- SAP templates
- group-interface templates

If the CP does not specify values for these profiles or templates, the UP automatically uses any profile or template configured with the name “default”. The following example displays default configuration settings for these profiles and templates.

See [SAP and group interface templates](#) and [Subscribers, QoS, and filters](#) for more information.



Note: The default SLA and subscriber profiles provide the minimal functionality to set up a subscriber. For actual deployments, tailor these profiles to the required use cases. See [Overview of dynamic QoS based on PCC rules](#) for information about profile configurations for FWA sessions with dynamic PCC rules.

Example**Configuration of SLA and subscriber profiles and SAP and group-interface templates**

```
[ex:/configure subscriber-mgmt]
A:admin@UP-West# info
  group-interface-template "default" {
  }
  sap-template "default" {
  }
  sub-profile "default" {
    control {
      cups true
    }
  }
  sla-profile "default" {
    control {
      cups true
    }
  }
}
```

3 PFCP association

This chapter provides an overview of the BNG-UP Packet Forwarding Control Protocol (PFCP) association configuration, heartbeat and headless mode operation, and operational management and debugging options.

3.1 BNG-UP PFCP association

A BNG-UP requires an active PFCP association with the MAG-c. Through the PFCP association, the MAG-c installs the rules that determine how the BNG-UP forwards subscriber traffic.

The BNG-UP requires the PFCP association to be preconfigured using the **configure subscriber-mgmt pfcp association** command.

Each PFCP association is linked to a specific peer, interface, and router instance. The loopback interface is the recommended interface because it allows resiliency over multiple physical interfaces.

As soon as the PFCP association is configured and administratively enabled, the BNG-UP probes the MAG-c by sending a PFCP Heartbeat Request message. When the BNG-UP receives a corresponding PFCP Heartbeat Response message, it establishes a permanent PFCP association by sending a PFCP Association Setup Request message. If either the Heartbeat or the Association Setup message fails (for example, times out) the BNG-UP periodically retransmits the Heartbeat until this procedure succeeds or the PFCP association is administratively disabled.

The BNG-UP can also set up the PFCP association when it receives a PFCP Association Setup Request from the configured peer. When the peer node ID is configured, the BNG-UP does not accept an association setup from a non-configured peer.

**Note:**


- Nokia recommends always configuring a peer node ID.
- When configuring a single PFCP association, the peer node ID must be a fully qualified domain name (FQDN). The system ignores the option to configure an IP address for a single association.

Use the **fqdn** or the **use-ip-address** command option in the following context to configure the peer node ID.

```
configure subscriber-mgmt pfcp association peer node-id
```

The following figure shows the PFCP association setup flow for the BNG-UP.

Figure 1: BNG-UP PFCP association setup flow



The PFCP protocol supports node identification using node IDs. The node ID can be either an IP address or a FQDN. By default, the IP address of the linked interface is chosen; however, this can be overridden using the **node-id** command. All the BNG-UPs and MAG-Cs in a deployment require different node IDs.

PFCP messages sent for an active association use the following configuration under the PFCP association **tx** command:

- The **ttl** command defines the outgoing TTL.
- The **timeout** command defines the request message timeout (T1).
- The **retries** command defines the number of times a request message is retried (N1).



Note: For heartbeat messages, N1 and T1 are configured separately. For correct operation, N1 and T1 must be configured identically on both the MAG-c and BNG-UP. See the *MAG-c Control Plane Function Guide* and the *MAG-c CLI Reference Guide* for more information about the MAG-c configuration.

The QoS of the outgoing PFCP messages is managed through **sgt-qos** command configuration in the routing instance used by the PFCP. In the **application** list command, a **pfcp** keyword can be mapped to its own DSCP value (default NC2), after which that DSCP value can be mapped to a specific Forwarding Class (FC).

See "QoS for Self-Generated (CPU) Traffic on Network Interfaces" in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Quality of Service Guide*, for more information about QoS configuration.

When a PFCP association is administratively disabled, it is not immediately brought down. The BNG-UP requests a graceful MAG-c release and keeps the PFCP association up until the MAG-c removes it, or the association release timeout expires.

To configure the PFCP association release timeout, use the **association-release-timeout** command.

To force an immediate PFCP association removal, configure the **association-release-timeout** to **none**.

The following example shows a BNG-UP PFCP association configuration.

Example: BNG-UP PFCP association configuration

```
[ex:configure subscriber-mgmt pfcp association "MAG-c"]
A:admin@DUT-B# info detail
  admin-state enable
  ## description
  association-setup-retry 1
  association-release-timeout 3600
```

```

path-restoration-time 180
node-id {
  ## fqdn
  ## use-ip-address
}
interface {
  router-instance "to_magc"
  name "endpoint"
}
peer {
  ip-address 17.17.17.10
}
heartbeat {
  interval 60
  timeout 5
  retries 4
}
tx {
  timeout 5
  retries 3
  ttl 255
}

```

3.2 Multiple PFPCP associations for FWA

The FWA-UP supports configuration of multiple PFPCP associations to allow multiple MAG-c instances to control the FWA-UP. A peer Node ID must be configured for the association, to enable the FWA-UP to correctly match incoming PFPCP requests to one of the configured associations, based on the signaled Node ID. Any incoming packet containing a Node ID that does not match the configured Node ID is dropped.

Use one of the following commands to configure the peer Node ID to match incoming PFPCP messages to one of the configured associations based on the signaled Node ID.

- Use the FQDN to identify the peer Node ID.

```
configure subscriber-mgmt pfcpcp association peer node-id fully-qualified-domain-name
```

- Use the configured peer IP address to identify the peer Node ID. This is the default option.



Note: This configuration is ignored if only one association is configured.

```
configure subscriber-mgmt pfcpcp association peer node-id use-ip-address
```



Caution:

The assigned session IP addresses used to provision all connected MAG-c instances must not overlap. The FWA-UP rejects any session with overlapping IP addresses or aggregate routes.

3.3 PFPCP heartbeats and headless mode

The following concepts define the connectivity between the BNG-UP and the MAG-c:

- **PFCP association**
One PFCP association is allowed per BNG-UP and MAG-c. The identifiers of the association are the BNG-UP and MAG-c node IDs.
- **PFCP path**
Multiple PFCP paths are possible per association. The identifier of a PFCP path is the pair of IP addresses that are used to communicate between the BNG-UP and the MAG-c. Paths are not negotiated but are learned while using PFCP signaling.
See [BNG-UP PFCP association](#) for information about how PFCP associations are negotiated.



Note: The terms PFCP path and PFCP association are often used interchangeably because there is typically only one PFCP path per PFCP association.

The BNG-UP only uses one IP address, but starts heartbeats for each PFCP path it learns. The frequency, timeout, and retry values of the heartbeats are configured for the PFCP association using the **interval**, **retries**, and **timeout** commands in the **configure subscriber-mgmt pfcp association heartbeat** context.

If a heartbeat fails, the BNG-UP starts a timer based on the **path-restoration-time** command configuration under the PFCP association. If the timer expires or is not configured, all sessions associated with that path are removed. If the path recovers before the timer expires, the timer is canceled, and no sessions are removed.



Note: For correct operation, the heartbeat configuration must be identical on both the BNG-UP and MAG-c. Nokia strongly recommends configuring the **path-restoration-time** to at least twice the sum of the **heartbeat interval** plus the total timeout (**heartbeat retries** x **heartbeat timeout** = $N1 \times T1$).

To expedite the detection of path failures, enable BFD using the **bfd-expedited-path-down** command in the **configure subscriber-mgmt pfcp association** context. When enabled, the system starts a BFD session for each known PFCP path on the BNG-UP. If the system detects a BFD failure, it immediately brings down the associated path. BFD does not affect the path recovery detection, which requires the configuration of PFCP heartbeats. BFD-based path down detection requires the configuration of the **path-restoration-time** command under the PFCP association.

3.4 Default IBCP session

The BNG-UP always enables the IPoE and PPPoE BBF function features. A compatible MAG-c uses this as an indication to set up a default IBCP tunnel to send upstream data-trigger packets and control plane packets, such as DHCP discover or PADI, that do not match an existing session. The default tunnel is signaled as a special PFCP session without a PDN type. The special PFCP session can apply traffic-matching rules the same as any other PFCP session for which the BNG-UP applies rules.

Packets sent over the default In-band Control Plane (IBCP) tunnel include local BNG-UP command options that are inserted in a network service header. The network service header is inserted between the GTP-U header and the encapsulated control packet. These command options include the following:

- **MAC address**

This is the MAC address that is associated with the capture SAP on which the IBCP packet was received. The MAG-c uses the MAC address as a source MAC address when sending control packet responses. In case of inter-BNG-UP resiliency, the MAG-c can ignore this MAC address and use the FSG MAC address instead. For more information about inter-BNG-UP resiliency, see [BNG-UP resiliency](#).

- **Layer 2 access ID**

The Layer 2 access ID (also known as the logical port) identifies the local port, the LAG, or the pw-port that is used by the capture SAP. The MAG-c reflects this parameter as is when setting up a session so the BNG-UP installs the session in this context. To simplify provisioning on the MAG-c, you can provide an alias for the L2 access ID. The alias must be unique within the BNG-UP and cannot overlap with any identifier such as a port or a LAG name.

To configure the L2 access ID alias, use the **configure service vpls capture-sap pfcpl2-access-id-alias** command.

See [IBCP](#) for more information about IBCP.

3.5 Overload and prioritization

The BNG-UP prioritizes PFCP session messages based on the message priority (MP) bits signaled in the PFCP header. The MP is a 4-bit optional field that accepts a value from 0 to 15, where 0 denotes the highest priority and 15 the lowest priority. The MP field can be excluded by setting the MP bit in the header to 0. The BNG-UP treats this the same as if the MP field was included with the value 15.

The following table shows the MP field in the PFCP header.

Table 2: Message header for session-related PFCP messages

Bits								
Octets	8	7	6	5	4	3	2	1
1	Version			Spare	Spare	FO	MP	S=1
2	Message Type							
3	Message Length (1st Octet)							
4	Message Length (2nd Octet)							
5 to 12	Session Endpoint Identifier							
13 to 15	Sequence Number							
16	Message Priority				Spare			

When processing the message, the BNG-UP applies an absolute priority based on this MP at various stages in the system. High priority messages take precedence over lower priority messages, and may be processed in a different order from that in which they arrive.

The MAG-c determines the MP to use for session prioritization and as a base overload protection mechanism. For example, the MAG-c marks messages for an already-established session as higher priority than messages for a new session. This ensures that in overload conditions the BNG-UP drops new sessions in favor of maintaining existing sessions.

See the *MAG-c Control Plane Function Guide* for more information about overload and prioritization.

3.6 Operational commands and debugging

This section describes commands that can be used for operational and debugging purposes.

To display the number of PFCP associations and sessions, use the **show subscriber-mgmt pfcpl summary** command.

Use the following command to display PFCP message statistics, including information about packets that are received and transmitted and any transmission errors.

```
show subscriber-mgmt pfcpl statistics
```

Use the following command to reset the PFCP association statistics to zero.

```
clear subscriber-mgmt pfcpl-association statistics
```

Use the following command to display information including association-level options such as function features and node IDs.

```
show subscriber-mgmt pfcpl association
```

Use the following command to display path state information.

```
show subscriber-mgmt pfcpl peer
```

See [PFCP heartbeats and headless mode](#) for information about the distinction between PFCP association and path.

Use the following command to display default IBCP tunnel information.

```
show subscriber-mgmt pfcpl session default-tunnel
```

You can display an overview of the PFCP configuration for the session. Various filters can be applied to narrow a specific session or set of sessions. Use the following command to display details about a specific PFCP session.

```
show subscriber-mgmt pfcpl session detail
```

See [Operational commands](#) for more information about basic operational commands.

Use the following command to forcefully remove a PFCP session.

```
clear subscriber-mgmt pfcpl-session
```



Caution: Although it is possible to forcefully remove a PFCP session, Nokia does not recommend invoking this command in a live network because this may cause the MAG-c to be out of sync with the BNG-UP. If the BNG-UP contains a session that is not on the MAG-c, it is preferable to first run a manual PFCP audit.

To perform basic PFCP debugging, use the **debug subscriber-mgmt pfcpl** command. The output of this command allows inspection of PFCP packets received and transmitted. As well, any session-specific failure is reported in the PFCP response to the MAG-c, which can display the specific error as part of session debugging.

4 Session management

This chapter provides an overview of the common and fixed access session functionality.

4.1 Subscribers, QoS, and filters

The BNG CUPS automatically links sessions together into subscribers, based on the QoS Enforcement Rule (QER) Correlation ID it receives in the PFCP session. If the BNG-UP does not receive a QER Correlation ID, it assumes there is only one session per subscriber.

The usual subscriber-management processing applies, as described in the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide*, sections "QoS for subscribers and hosts" through "Configuring IP and IPv6 filter policies for subscriber hosts".

The PFCP may pass the following parameters:

- subscriber and SLA profile names, or the name "default" if no profiles are configured



Note: You must always configure the SLA and subscriber profiles for use with BNG CUPS. This disables any feature that is not supported within BNG CUPS. Use the commands in the following contexts to configure a SLA and subscriber profile.

```
configure subscriber-mgmt sla-profile control cups
configure subscriber-mgmt sub-profile control cups
```

- direct QoS overrides of QoS objects such as aggregate-rate, schedulers, arbiters, queues, and policers
- SLA filter overrides, either by name or ID
- intermediate destination ID

PFCP can also signal GBR and MBR values directly in the QER, which is not linked directly to a specific QoS object. If the QER applies to the entire session, use the following command to map the MBR and GBR to a direct QoS object PIR or CIR override.

```
configure subscriber-mgmt sla-profile name pfcp-mappings session-qer
```

The MAG-c signals the QER rate, for example, to install a session-AMBR (5G) or APN-AMBR (4G) for FWA sessions.

4.2 IBCP

Most BNG session types have one or more control plane messages that are sent in-band and therefore arrive directly on the BNG-UP. Because the BNG-UP cannot handle these messages, they are forwarded to the MAG-c. To accomplish this, the MAG-c installs specific Ethernet or IP filter rules that match these packets; for example, by matching UDP destination port 67 to extract DHCP. These packets are

encapsulated in GTP-U and sent to the MAG-c. Similarly, the MAG-c sends downstream In-Band Control Plane (IBCP) packets over GTP-U toward the BNG-UP.

For upstream traffic, the BNG-UP sends any control plane messages that do not match a session over a default tunnel. See [Default IBCP session](#) for information about how this tunnel is signaled. If the control plane messages do not match the default tunnel rules, the messages are dropped.

When a session is created, either out-of-band or via a trigger over the default tunnel, the MAG-c installs per-session control plane rules for both upstream and downstream. Packets that match the upstream rules are forwarded to the MAG-c using the signaled GTP-U parameters. For downstream rules, the BNG-UP allocates a TEID that the MAG-c can use to send packets. The BNG-UP does not support a default downstream IBCP tunnel.

The upstream IBCP (including default) follows the **sgt-qos dscp application** configuration, using the **ibcp** keyword on the router or VPRN. A specific DSCP value (default NC2) can be provisioned and mapped to a specific FC, as usual.

See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Quality of Service Guide*, section "QoS for Self-Generated (CPU) Traffic on Network Interfaces" for more information.

Downstream IBCP packets are handled directly in the datapath. Ingress QoS is applied based on the provisioning. Egress QoS depends on the following session types:

- **fixed access sessions**

Egress QoS packets bypass per-session processing, including egress QoS and filters. Egress QoS is instead based on the QoS configuration of the capture SAP that is linked to the session.

- **FWA sessions**

Egress QoS packets go through regular per-session processing and are subject to the QoS and filters provisioned in the SLA profile.

4.3 IP gateway, services, and routing

In many deployments, a BNG-UP acts as a direct IP gateway for sessions. The MAG-c provides all the IP addresses and framed routes using the PFCP protocol. To assist with forwarding, the MAG-c also signals the following information using the PFCP protocol:

- name of the preprovisioned IES or VPRN service in which forwarding must occur
- aggregate routes that the BNG-UP announces in routing protocols to attract traffic



Note: The MAG-c guarantees that no addresses from the aggregate routes are assigned to sessions on another BNG-UP. See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Unicast Routing Protocols Guide* for more information about route policies.

- IPv4 gateway address, which is typically a dedicated address within the aggregate route
- IPv6 gateway link-local address, only one of which is supported per VPRN or IES

The BNG-UP runs the appropriate routing protocols and responds to ARP/ND for the gateway addresses. For the purpose of exporting routes, use the following command option to distinguish all CUPS routes.

```
configure policy-options policy-statement entry from origin pfcp
```

Additionally, the user can further distinguish routes using the following options for the **protocol** command in the same context:

- **sub-mgmt** – for session IP addresses and prefixes, as signaled in the UE IP Address IE in PFCP
- **managed** – for framed routes, as signaled in the Framed-Route IE in PFCP



Note: When using the **pd-as-framed-route** option on the MAG-c, an IPv6 PD prefix uses the **managed** option (instead of the **sub-mgmt** option) because the BNG-UP cannot distinguish the PD framed route from a regular framed route.

- **direct** – for all other routes, such as aggregate routes and gateway addresses

4.4 Statistics reporting

Statistics reporting uses the PFCP Usage Reporting Rule (URR) mechanism. The BNG-UP supports a single URR to count all statistics that are related to a session and supports sending the following statistics for the URR:

- Aggregate octet counters are always signaled.
- Aggregate packet counters are signaled, if enabled by the MAG-c.
- Per-queue and per-policer statistics are signaled, if enabled by the MAG-c. The specific counters depend on the statistics mode (**stat-mode**) configured in either the QoS policy or SLA profile.

All counters, including aggregates, are based on QoS counters and are therefore affected by QoS modifiers, such as the **packet-byte-offset** command.

The BNG-UP sends reports for a URR in the following cases:

- The MAG-c explicitly queries the BNG-UP via a PFCP Session Modification Request.
- The periodic URR reporting is enabled and the BNG-UP sends unsolicited PFCP Session Report Request messages.
- A threshold or quota is reached; see [Threshold and quota monitoring](#) for more information.

PFCP statistics are reported in an incremental manner. This means that only new statistics after the last report are signaled. To achieve this, the BNG-UP baselines the counters on every report. Consequently, it is not possible to use the following command to manually clear statistics on the BNG-UP.

```
clear service statistics subscriber
```

Other operational commands such as the following only show the accumulated statistics on the BNG-UP.

```
show service active-subscribers detail
```

Because statistics are based on QoS counters, sessions sharing the same SLA Profile Instance (SPI) also share statistics, and a report for one session baselines the counters for the entire SPI. As a result, per-session statistics on the MAG-c are not correct when sharing an SPI; however, their aggregate counts are correct. The MAG-c must provide the appropriate aggregate level (for example, subscriber-level accounting). When an SPI changes, the BNG-UP reports the final SPI statistics in PFCP if instructed to do so by the MAG-c.

Hardware failures are automatically taken into account for statistics reporting. Statistics generated after the last report are irretrievably lost. However, as a result of the incremental reporting, the MAG-c does

not lose any older counters and does not see a sudden reset. That is, aggregate counters on the MAG-c never decrease as a result of a hardware failure. However, the BNG-UP local statistics as seen in **show** commands reset upon a hardware failure, and therefore a mismatch of MAG-c counters may result.

4.4.1 Selective aggregate statistics

Default and custom aggregate statistics calculation

The BNG-UP calculates aggregate statistics based on the statistics of policers and queues. By default, the BNG-UP accumulates all the following policers and queues for statistics calculation:

- **All egress queues**
All egress queues includes both dynamically created and statically configured queues.
- **All statically configured egress policers**
Dynamic egress policers are by default not included. The traffic that is subject to a dynamic egress policer is usually also subject to a static or dynamic egress queue, and therefore is already counted.
- **All ingress queues**
All ingress queues include both dynamically created and statically configured queues.
- **All ingress policers**
All ingress policers includes both dynamically created and statically configured policers.

You can change the default calculation by configuring a custom selection of policers and queues. Use the commands in the following context to change the default queues and policers that the BNG-UP accumulates for aggregate statistics calculations.

```
configure subscriber-mgmt sla-profile aggregate-qos-statistics
```

Example: custom selective aggregate statistics to only count policers

The following example shows a configuration that counts all dynamic policers and no queues.

```
[ex:/configure subscriber-mgmt sla-profile "example" aggregate-qos-statistics]
A:admin@UP-East# info
egress {
  queues {
    none
  }
  policers {
    all-dynamic
  }
}
```

Use case examples

The following are examples of use cases for which you may want to configure granular control of aggregate statistics calculation:

- to zero rate specific traffic
- to avoid double counting packets in QoS hierarchies that combine policers and queues

You can use this functionality to zero rate traffic, so it is not counted in aggregate statistics while maintaining full QoS enforcement. The following example shows a zero-rating QoS configuration for FWA with dynamic flows, where the following applies:

- The zero-rated traffic is identified by a specific DSCP value AF42, and is excluded from aggregated statistics but is still subject to the session AMBR.
- All traffic other than the zero-rated traffic is subject to a dynamic policer (SDF MBR), and either the session AMBR queue, or a GFBR or MFBR dynamic queue.
- The QoS profile is configured with the following:
 - A single static queue enforces the session AMBR. This is not visible in the configuration example because the default queue is used.
 - A single static policer for zero rating is configured with an infinite PIR or CIR because this policer is only used for counting, and not for rate enforcement.
 - All necessary configurations to enable dynamic flows are used. See [Dynamic QoS based on PCC rules](#) for more information about the configuration of dynamic objects.
 - A single static IP criteria entry is used, with a higher precedence value (lower entry ID) than any dynamic criteria. This entry matches the DSC accepts the traffic, and subjects the traffic to the static policer and static queue.
- The SLA profile is configured to use queue 1 as the session AMBR and only count traffic from any dynamic policers for aggregate statistics.

Example: Zero rating configuration with dynamic QoS



Note: Only the egress IPv4 configuration is shown, to reduce the example output.

```
[ex:/configure qos sap-egress "zero_rating"]
A:admin@UP-East# info
  subscriber-mgmt {
    pcc-rule-entry {
      range {
        start 10000
        end 65535
      }
    }
    dynamic-policer {
      policer-id-range {
        start 10
        end 63
      }
    }
    dynamic-queue {
      queue-id-range {
        start 2
        end 8
      }
    }
  }
  policer 1 {
    description "Policer used for zero-rated traffic"
    rate {
      pir max
      cir max
    }
  }
  ip-criteria {
    entry 10 {
      match {
        dscp af42
      }
    }
  }
}
```



```

        action {
            type accept
            policer 1
            queue 1
        }
    }
}

```

```

[ex:/configure subscriber-mgmt sla-profile "default"]
A:admin@UP-East# info
control {
    cups true
}
egress {
    qos {
        sap-egress {
            policy-name "zero_rating"
        }
    }
}
aggregate-qos-statistics {
    egress {
        queues {
            none
        }
        policers {
            all-dynamic
        }
    }
}
pfc-mappings {
    session-qer {
        downlink {
            queue 1
        }
    }
}
}

```

You can also use the selective aggregate statistics functionality to avoid double counting packets in more complex QoS hierarchies. The following example shows a configuration where all the traffic is subject to a common queue, but DNS (TCP/UDP port 53) traffic is additionally subject to a lower rate using IP criteria and policers. The aggregate statistics configuration shown in the example avoids counting the DNS traffic twice, as part of the static policer and the static queue.

Example: Configuration to avoid double-counting aggregate statistics as part of the static policer and static queue



Note: Only the egress IPv4 configuration is shown, to reduce the example output.

```

[ex:/configure qos sap-egress "rate_limit_dns"]
A:admin@UP-East# info
queue 1 {
    rate {
        pir 100000
        cir 10000
    }
}
policer 1 {
    description "Policer to restrict DNS traffic to 1Mbps"
}

```

```

    rate {
        pir 1000
        cir 1000
    }
}
ip-criteria {
    entry 10 {
        match {
            protocol tcp
            src-port {
                eq 53
            }
        }
        action {
            type accept
            policer 1
            queue 1
        }
    }
    entry 20 {
        match {
            protocol udp
            src-port {
                eq 53
            }
        }
        action {
            type accept
            policer 1
            queue 1
        }
    }
}
}

```

```

[ex:/configure subscriber-mgmt sla-profile "rate_limit_dns"]
A:admin@UP-East# info
control {
    cups true
}
egress {
    qos {
        sap-egress {
            policy-name "rate_limit_dns"
        }
    }
}
aggregate-qos-statistics {
    egress {
        policers {
            all-dynamic
        }
    }
}
}

```

4.4.2 Custom statistics reporting

You can provision the BNG-UP with a custom statistics group that reports statistics based on a subset of user-specified QoS objects. The MAG-c references this custom statistics group in the URRs it sends in

PFCP messages. The BNG-UP reflects the aggregated QoS object statistics in the URRs in PFCP usage reports.

Use the following command to provision a custom statistics group.

```
configure subscriber-mgmt sla-profile custom-statistics-group
```

As an example, consider the zero rating use case described in [Selective aggregate statistics](#). The following example shows how to configure a custom statistics group to report the volume of traffic that is zero rated.

Example: Custom statistics-group configuration to report zero-rated traffic

```
[ex:/configure subscriber-mgmt sla-profile "default" custom-statistics-group "zero-rated-traffic"]
A:admin@UP-East# info
  egress {
    policers {
      policer [1]
    }
  }
```

4.5 Threshold and quota monitoring

MAG-c threshold and quota monitoring builds upon the URR-based usage reporting, as described in [Statistics Reporting](#). The MAG-c signals a quota or threshold in the URR to monitor the total volume consumed. When a threshold or quota is reached, the UP generates a usage report to the MAG-c, including the consumed statistics. The UP indicates in the Usage Report Trigger field, whether the report was triggered by a quota or a threshold. This allows the MAG-c to take appropriate actions, such as the following:

- trigger a statistics report for offline charging
- request new quota or thresholds for online charging
- apply more complex out-of-credit actions such as custom PCC rules

When a threshold is reached, the UP reapplies the threshold immediately. This can be used to enable volume-based envelope reporting, where a report is generated based on the change in volume, for example every 1 GB, instead of periodically.

When a quota is reached, the UP immediately stops forwarding all traffic associated with the URR. For the session URR, it drops traffic until new quota are granted. Control plane traffic is still sent to the MAG-c corresponding to the IBCP rules.

Internally, the UP uses credit-control categories to perform monitoring. See the 7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide, "Volume and time-based accounting", for more information about credit-control categories. This has the following consequences:

- As with category maps, monitoring is done per SPI, not per session. To avoid unexpected behavior with per-session charging, it is important to use a single SPI per-session model.
- Line cards consume category resources. Specifically, each monitored unit (for example, total volume) per URR consumes one category.
- Checks for category exhaust limits are performed periodically and not on each packet. This means that a monitoring overshoot is possible. The amount of overshoot depends on the system type and system load.

Relationship between the UP and MAG-c threshold and quota

The MAG-c typically maps multiple applications to a URR threshold or quota. It is therefore normal that a threshold or quota installed on the UP does not correspond to a threshold or quota applied by a charging function such as the CHF.

The following are example scenarios of multiple applications mapped to a URR threshold or quota:

- A charging quota can map to both a UP quota, to immediately stop forwarding upon exhaustion, and to a UP threshold, to apply more complex actions (such as HTTP redirection) instead of stopping forwarding.
- An online charging threshold and offline charging volume limit can both use the threshold, based on whichever expires first. For example, if a volume limit of 300 MB and a threshold of 1 GB are configured, the initial URR threshold is 300 MB, but after 3 reports it changes to 100 MB to reflect the remaining online threshold of 1 GB – 3 * 300 MB.

4.6 Operational commands

Most of the traditional BNG operational commands, as described in the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide*, apply to the CUPS BNG-UP. The significant exceptions to this rule are operational commands related to specific protocols (such as DHCP, DHCPv6, RADIUS, and PPPoE), because a BNG-UP is not aware of these states.

The primary BNG-UP operational commands are as follows:

- **show service active-subscribers**

This command contains several sub-commands that provide details about a specific subscriber or session within a subscriber. These commands incorporate information about CUPS subscribers. Information that is only available on the MAG-c is not shown on the BNG-UP (for example, details on RADIUS and metadata such as **remote-id** and **circuit-id**).

- **show subscriber-mgmt statistics**

This command contains several sub-commands that provide a wide variety of statistics on various granularity levels. These commands are extended to incorporate BNG CUPS statistics.

IBCP statistics can be displayed via the PFCP statistics using the **show subscriber-mgmt pfcp statistics** command.

Operational commands that are specific to PFCP associations are described in [Operational commands and debugging](#).

4.7 SAP and group interface templates

The system auto-provisions any required objects, which means that subscriber interfaces, group interfaces, and SAPs do not need to be provisioned. These objects are hidden from configuration and are not modifiable. Aside from the capture SAP, the only required configuration is the VPRN or IES where IP forwarding occurs.

Use the commands in the following context to manage SAP creation using a SAP template.

```
configure subscriber-mgmt sap-template
```

The SAP template supports the configuration of the following:

- Use the **hold-time** command to delay the deletion of the SAP after the last PFCP session is removed.



Note: Although the **infinite** option can be configured for the **hold-time** command, it is not recommended.

Use the following command option to clear idle SAPs.

```
clear subscriber-mgmt sap-template idle-saps
```

- Use the **cpu-protection** and **dist-cpu-protection** commands to configure CPU protection; see the *7450 ESS, 7750 SR, 7950 XRS, and VSR System Management Guide*, section "Centralized CPU protection and distributed CPU protection" for more information about CPU protection.



Note: On platforms where CPU protection and distributed CPU protection are not supported, these commands are ignored.

Use the commands in the following context to configure a group interface template to manage the created group-interface.

```
configure subscriber-mgmt group-interface-template
```

When setting up a PFCP session, the PFCP passes a template name. If the template name is absent, the system falls back to the name "default".

Use the following commands to configure the group interface template:

- Use **ipv4 icmp** to configure ICMP.
- Use **ip-mtu** to configure IP MTU to outgoing packets.
- Use **ipv4 remote-proxy-arp** to configure the remote proxy ARP; see *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide*, section "Proxy ARP".
- Use **ipv4 urpf-check** and **ipv6 urpf-check** to configure URPF checks; see *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide*, section "Unicast reverse path forwarding check".
- Use **trigger-packet data** to enable data trigger.



Note: Configure the SAP and group interface templates on the BNG-UP (as well as the name "default") to ensure the session setup does not fail.

Changing the configuration of a template does not automatically change all created SAPs or group interfaces.

4.7.1 Mixing different encapsulation sessions on the same port

The BNG-UP supports the following mix of encapsulation sessions:

- *.* capture SAP – dot1q and null encapsulation sessions



Note: To support this, enable the following command on the systems where this is not the default.

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

- * capture SAP – null encapsulation sessions

The system internally creates SAPs with :N.0, :0.0, and :0 encapsulation to support the mix. The user can apply the SAP templates to these SAPs similarly to any other SAPs.

Traffic as seen on the wire typically does not include an S-tag or C-tag with a tag value explicitly set to zero because the absence of a tag is equal to a zero tag. Occasionally, an explicit zero tag is included; for example, when a dot1p or DEI bit needs to be set in the tag header. The internally created SAPs interact with such traffic as follows:

- Because generated traffic never includes an explicit zero tag, it is not possible to set a dot1p or DEI bit in the tag header.
- Traffic with a C-tag value explicitly set to zero is always dropped, even if it matches an internally created :N.0 or :0.0 SAP.
- Traffic with only the S-tag value explicitly set to zero is handled if it matches either a :0 or a :0.0 SAP. This kind of traffic matches the same session as if it was received without any S-tag.

4.8 Fixed access sessions

To enable fixed access sessions, you must provision a capture SAP under the VPLS service with appropriate values for trigger-packet and a link to the PFCP association. The triggers are mandatory and are not automatically derived from the default IBCP tunnel.

Sessions without any encapsulation are supported on a dot1q capture SAP. The system creates internal constructs to correctly handle sessions without encapsulation. These sessions can be combined with dot1q encapsulated sessions on the same capture SAP.

The following example shows trigger-packet provisioning in a PFCP association configuration.

Example

```
A:admin@node-2# info
  pfcf {
    association "BNG-CPF"
  }
  trigger-packet {
    pppoe true
  }
```

To identify sessions in the data plane, the MAG-c must provide the following parameters:

- The logical port (also known as the Layer 2 access ID) identifies the port, the LAG, or the pw-port where the session is terminated. The MAG-c knows the correct logical port because the BNG-UP includes this logical port in the IBCP packets sent over the default IBCP tunnel. See [Default IBCP session](#).
- The VLAN tags, along with the logical port, identify a SAP where the session is terminated, also known as a Layer 2 circuit (l2-circuit). The MAG-c must signal all the VLAN tags to match the encapsulation type provisioned for the port; for example, only signaling an S-tag for a QinQ port is not allowed. As an exception, the MAG-c can install sessions without any VLAN tags on a dot1q capture SAP, as described at the beginning of this section.
- The source MAC address is required.
- The PPPoE session ID is used for PPPoE only.

- The IP anti-spoofing IP address is optionally used to enable IP anti-spoofing. While this can be signaled per session, the BNG-UP only supports a single anti-spoof type per SAP. When a second session on the same SAP has a conflicting anti-spoof indication, the setup fails. IP anti-spoofing is not supported for framed routes.

For PPPoE, the BNG-UP can perform LCP keep-alive offload, if supported and signaled by the MAG-c. The BNG-UP automatically signals support for this feature when the PFCP association is created.

4.9 Fixed wireless access sessions

The BNG-UP supports both 4G and 5G FWA sessions.

4.9.1 Configuring FWA sessions

Prerequisites

Consider the following before configuring FWA sessions:

- The configuration for both 4G and 5G sessions is the same and can be shared. Both 4G and 5G use GTP-U tunnels; however, an additional GTP-U extension header is added for 5G that contains a 64-bit QoS flow indicator (QFI), and optionally a 1-bit reflective QoS indicator (RQI). The system decides whether to include this extension header based on dynamic signaling in the PFCP; no additional user configuration is required.
- FWA sessions do not use default IBCP tunnels because the initial session setup is out-of-band and does not involve the FWA-UP. Per-session IBCP is supported to forward deferred allocation (DHCP), ICMPv6 RS/RA, and DHCPv6 packets.
- The system automatically creates all the required constructs to correctly handle FWA sessions. Additionally, the system load balances sessions over each forwarding path extension (FPE), using the path that is the least loaded when the session is set up.
- See the *MAG-c Control Plane Function Guide*, "Fixed wireless access sessions", for more information about configuring FWA sessions.

About this task

FWA sessions require configuration of an FWA-UP data endpoint for the router or VPRN, an optional secondary interface to terminate GTP-U tunnels, and an FPE construct to enable the datapath function in the router.

Procedure

Step 1. Configure an FWA-UP data endpoint for the router or service VPRN.

The endpoint must reference an interface in the routing instance that is configured with GTP-U tunnels to the RAN. The MAG-c CP also signals this routing instance in the S1-u or N3 Network Instance (network realm) in PFCP in the format *service-name*, where the service name is either the VPRN service name or "Base" to indicate the base router.

```
configure router gtp upf-data-endpoint interface
configure service vprn gtp upf-data-endpoint interface
```

Step 2. Optional: Configure a secondary interface to terminate GTP-U tunnels.

The interface created in step 1 is used for the primary, and the interface created in this step is used for the secondary. You can use this, for example, to differentiate 4G from 5G traffic, by provisioning different network instances (network realms) for S1-u (4G) and N3 (5G).

```
configure router gtp upf-data-endpoint secondary-interface
configure service vprn gtp upf-data-endpoint secondary-interface
```

The MAG-c CP can choose the primary or secondary interface to use by applying the corresponding special S1-u or N3 Network Instance (network realm) configurations on the MAG-c:

- `service-name#%primary`
- `service-name#%secondary`

Step 3. Configure an FPE construct as type **multi-path** to enable the datapath functions in the router. See the *7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide* for more information.

Example: FWA-UP data-endpoint and multipath FPE configuration

```
[ex:/configure service vprn "to_ran" gtp upf-data-endpoint]
A:admin@FWA-UP# info
    interface "gtp_u_endpoint"
    fpe 1

[ex:/configure fwd-path-ext fpe 1]
A:admin@FWA-UP# info
    multi-path {
        path 1 {
            pxc 1
        }
    }
    application {
        sub-mgmt-extension true
    }
```

4.9.2 Dynamic QoS based on PCC rules

FWA sessions support dynamic QoS based on policy and charging control (PCC) rules, based on the configuration guidelines and examples provided in the topics in this section.

4.9.2.1 Overview of dynamic QoS based on PCC rules

FWA sessions support dynamic QoS based on policy and charging control (PCC) rules. An FWA session can have multiple PCC rules, each composed of a set of packet classifiers and specific QoS behavior. Acting as a service management function (SMF), the MAG-c receives these PCC rules from multiple sources, such as a policy control function (PCF). The MAG-c translates these PCC rules to common PFCP constructs, where each PCC rule is represented by a packet detection rule (PDR) and associated QoS enforcement rule (QER) or forwarding action rule (FAR).

The FWA-UP supports up to two QER rates per PDR, which it enforces using the following automatically-derived two-level QoS hierarchy:

- The first rate is enforced using policing, and represents the service data flow (SDF) maximum bit rate (MBR) as constructed by the MAG-c. This rate is optional in the 5G QoS model. For uplink traffic, this rate is not enforced if the session AMBR is mapped to a policer using the following command.

```
configure subscriber-mgmt sla-profile pfc-mappings session-qer uplink
```

- The second rate is enforced using shaping for downlink traffic and policing for uplink traffic. This rate represents either the session AMBR or the guaranteed flow bit rate (GFBR) or maximum flow bit rate (MFBR), depending on whether the PCC rule is mapped to a GBR or non-GBR QoS flow.

See the *MAG-c Control Plane Function Guide*, "Fixed Wireless Access QoS", for more information about SDF MBR, session-AMBR, GFBR, and MFBRs.

For each PCC rule, the FWA-UP correctly marks packets as follows:

- The FWA-UP marks the QoS flow identifier (QFI) and reflective QoS indicator (RQI) in the GTP-U header. The RAN and FWA RG use these values respectively to correctly handle QoS enforcement and classification.
- The FWA-UP marks the DSCP fields on both the inner data packet and the outer GTP-U header. This is important primarily in cases where a node that is not mobile-QoS aware is a bottleneck. These nodes cannot use QFI because it is dynamic and session dependent; for example, a specific QFI 1 may mean high-priority for one session and low-priority for the next. Therefore, you can also map traffic to the less flexible static DSCP, which provides good basic prioritization in these nodes.

The FWA-UP automatically configures the required resources to manage the per-flow QoS, based on the following PFCP rules:

- Configure a policer for each SDF MBR to be enforced.
- Configure a queue for each GFBR/MFBR to be enforced.



Note: The queue for the session-AMBR must be preprovisioned and identified by the session QER and SLA profile, as described in "[Subscribers, QoS, and Filters](#)".

- Configure IP criteria entries that include the signaled classifiers as match criteria, point to the applicable configured policer and queue, and enforce the correct QFI and RQI fields of the GTP-U headers.
- Configure IP filter entries that enforce drop and forward functionality and optionally enforce DSCP marking.

See [Configuring profiles for dynamic object creation](#) for information about configuring profiles for dynamic object creation and [Examples of profile configuration for dynamic object creation](#) for example configurations.

4.9.2.2 Using alternate QoS profiles in combination with dynamic PCC rules

At high system scale, there may not be enough QoS resources available to enforce all dynamic QoS using the configured model. For example, depending on the hardware used, the number of arbiters on a system may not suffice for full scale. The typical solution is to adapt to a different QoS model. In the arbiters example, it is an option to bypass arbiters by not enforcing uplink SDF MBR and only enforcing uplink session AMBR.

To optimize this behavior, the FWA-UP supports QoS oversubscription using alternate QoS profiles. An alternate QoS profile is an SLA or subscriber profile that is used when the system reaches a specific threshold. This allows the system to use an optimal QoS profile as long as possible, and only switch to

a sub-optimal QoS profile when the system reaches that threshold. Continuing the arbiter example, in best-case scenarios, all UPs in a deployment have enough arbiters to handle all sessions with SDF MBR enforcement. However, if a single UP fails, the sessions from that UP push the remaining UPs to their resource limits. When the system uses alternate profiles, only those additional failed-over sessions use the alternate profiles without SDF MBR enforcement.

The system always applies alternate profiles to a whole subscriber, and determines when the first session of that subscriber is created. If you create a subscriber without alternate profiles, any subsequent sessions are created without alternate profiles, even if at that time new subscribers are created with alternate profiles. The opposite is true if you create a subscriber with alternate profiles. Any subsequent sessions are created with alternate profiles, even if at that time new subscribers are created without alternate profiles.

Perform the following steps to configure QoS oversubscription.

1. Configure the number of subscribers without alternate QoS profiles that the system can support.

```
configure subscriber-management alternate-profile-threshold subscriber-count
```



Note: It is important to know how many QoS resources a single subscriber consumes, to ensure that the system reaches this threshold before resource exhaustion. If the system exhausts the resources first, the creation of new subscribers fails even if alternate profiles are configured.

2. Configure an alternate SLA profile or subscriber profile using the following commands. The profiles in which these alternate profiles are referenced are the original profiles used to create the subscriber if no thresholds are hit.

```
configure subscriber-management sub-profile alternate-sub-profile
configure subscriber-management sla-profile alternate-sla-profile
```

4.9.2.3 Configuring profiles for dynamic object creation

Prerequisites

Consider the following when configuring dynamic resources:

- All FWA sessions using PCC rules can use the same SLA and SUB profiles. The system correctly creates the per-session QoS resources based on these. It is also possible to configure the SLA and SUB profile with the name **default**, so that the MAG-c does not need to provision a specific SLA and SUB profile for these sessions. See [Subscribers, QoS, and filters](#) for details.
- To manage future growth, Nokia recommends configuring the largest range possible for dynamic resources, to accommodate adding or removing PCC rules. The range does not preallocate system resources and there is no reason to minimize these.
- See [Examples of profile configuration for dynamic object creation](#) for examples of dynamic resource configuration.

About this task

To support automatic resource configuration, provision the session with the following IP filter, QoS, and subscriber-management profiles:

Procedure

- Step 1.** Use the commands in the following contexts to configure an IPv4 and IPv6 filter with an entry range to use for PCC rule-based entries.

```
configure filter ip-filter subscriber-mgmt shared-entry pcc-rule range
configure filter ipv6-filter subscriber-mgmt shared-entry pcc-rule range
```



Note: This is not mandatory if you do not require DSCP marking or explicit flow blocking. However, Nokia recommends enabling this to accommodate any future functionality extensions that may require dynamic IP filter entries.

- Step 2.** Use the following commands to configure a subscriber profile that points to a policer-control-policy.

```
configure subscriber-mgmt sub-profile
configure subscriber-mgmt sub-profile egress qos policer-control-policy policy-name
```

- Step 3.** Configure a SAP egress policy with the following:

- a.** Configure at least one single static queue, which may be the default.

```
configure qos sap-egress queue
```

- b.** Configure an IP criteria entry range for the PCC rule derived entries.

```
configure qos sap-egress subscriber-mgmt pcc-rule-entry range
```

- c.** Configure a dynamic policer ID range to use when configuring the policers for SDF MBR policing. You can also configure other optional dynamic policer settings using commands in the following context.

```
configure qos sap-egress subscriber-mgmt dynamic-policer policer-id-range
```

- d.** Create a dynamic queue ID range for the queues required for GFBR and MFBR shaping. You can also configure other optional dynamic queue settings using commands in the following context.

```
configure qos sap-egress subscriber-mgmt dynamic-queue queue-id-range
```

- Step 4.** Configure a SAP ingress policy with the following:

- a.** Use the commands in the following context to configure an IP criteria entry range to create the PCC rule derived entries.

```
configure qos sap-ingress subscriber-mgmt pcc-rule-entry range
```

- b.** If uplink SDF MBR enforcement is not required, configure a single static policer and map all FCs to this policer using the following commands:

```
configure qos sap-ingress policer
configure qos sap-ingress fc policer
```



Note: Even when SDF MBR is required, Nokia recommends to configure this policer and point it to an arbiter in the policer control policy referenced by the subscriber profile. You can use the default root arbiter for this by using the name **root**. In this case, mapping the FCs is not required for correct classification but it does disable the creation of an ingress queue to minimize QoS resource usage.

- c. Configure a dynamic policer ID range to use when configuring the policers required for SDF MBR policing. You can also configure other optional dynamic policer settings in this context.

```
configure qos sap-ingress subscriber-mgmt dynamic-policer policer-id-range
```

- Step 5.** Configure a SLA profile that points to the configured IP filter, IPv6 filter, SAP egress policy, and SAP ingress policy.

- a. Use the commands in the following context to configure and correctly apply session-AMBR.

```
configure subscriber-mgmt sla-profile pfc-mappings session-qer
```

- b. Perform one of the following.

- If uplink SDF MBR enforcement is required, configure an uplink arbiter that points to an arbiter in the policer-control policy that the subscriber profile references. Use the name **root** to specify the default root arbiter. This must be the same arbiter to which the dynamic-policers is parented by the aforementioned configuration.

```
configure subscriber-mgmt sla-profile pfc-mappings session-qer uplink arbiter
```

- If uplink SDF MBR enforcement is not required, configure the policer created in step 4.b.

```
configure subscriber-mgmt sla-profile pfc-mappings session-qer uplink policer
```

- c. Configure the downlink queue to the static queue that is configured in the SAP egress policy.

```
configure subscriber-mgmt sla-profile pfc-mappings session-qer downlink queue
```

- d. Configure the SLA Profile to have a single instance per session. This is required to correctly handle QoS and charging if multiple sessions are created for the same subscriber.

```
configure subscriber-mgmt sla-profile def-instance-sharing
```

4.9.2.4 Examples of profile configuration for dynamic object creation

The following examples show filter, subscriber-management, and QoS configurations that enable creation of all the dynamic objects described in [Overview of dynamic QoS based on PCC rules](#) and [Configuring profiles for dynamic object creation](#).

Example: Filter configuration for creation of dynamic objects

```
[ex:/configure filter]
A:admin@node-2# info
  ip-filter "cups_default_ipv4_filter" {
    default-action accept
```



```

    }
  }
  pfc-p mappings {
    session-qer {
      uplink {
        arbiter "root"
      }
      downlink {
        queue 1
      }
    }
  }
}

```

Example: QoS configuration for creation of dynamic objects

```

[ex:/configure qos ]
A:admin@node-2# info
  sap-ingress "cups_default_ingress_qos_policy" {
    subscriber-mgmt {
      pcc-rule-entry {
        range {
          start 1
          end 65535
        }
      }
      dynamic-policer {
        policer-id-range {
          start 2
          end 63
        }
      }
    }
  }
  policer 1 {
    arbiter-parent {
      arbiter-name "root"
    }
  }
  fc "af" {
    policer 1
  }
  fc "be" {
    policer 1
  }
  fc "ef" {
    policer 1
  }
  fc "h1" {
    policer 1
  }
  fc "h2" {
    policer 1
  }
  fc "l1" {
    policer 1
  }
  fc "l2" {
    policer 1
  }
  fc "nc" {
    policer 1
  }
}

```

```

policer-control-policy "cups_default_pcp" {
}
sap-egress "cups_default_egress_qos_policy" {
  subscriber-mgmt {
    pcc-rule-entry {
      range {
        start 1
        end 65535
      }
    }
    dynamic-policer {
      policer-id-range {
        start 1
        end 63
      }
    }
    dynamic-queue {
      queue-id-range {
        start 2
        end 8
      }
    }
  }
}
}

```

4.10 Routed subscriber sessions

Enterprise IP VPNs often use BGP to exchange routing information, for example, between headquarters and branch offices. Providing BGP connectivity to an enterprise router via a residential access connection requires BGP peering over a BNG CUPS subscriber session. To configure this on the BNG-UP, use a static BGP neighbor that has as its neighbor address the fixed IPv4 or IPv6 WAN address of the subscriber session.

The following example shows an auto-generated configuration based on a session setup.

Example: Routed subscriber session configuration

```

A:node-2/config>svc#
vprn 1000
  subscriber-interface "_tmnx_cups_1131076" fwd-service 2148278386
  /fwd-subscriber-interface "_tmnx_cups_1131075" wan-mode mode128
  description "default subscriber interface template"
  private-retail-subnets
  address 10.0.240.254/20
  ipv6
  exit
exit

--- the auto-generated configuration ends here ---

bgp
  group "enterprise-1"
  neighbor 10.0.0.1
  family ipv4
  local-address 10.0.240.254
  type external
  local-as 65536
  peer-as 65501
  exit

```

```

    exit
    no shutdown
  exit
  no shutdown

```

BNG CUPS subscriber sessions support static BGP neighbors for the following:

- BGPv4 neighbors with IPv4 and IPv6 address families
- BGPv6 neighbors with IPv4 and IPv6 address families
- IPoE and PPPoE sessions
- single-hop and multi-hop BGP neighbors

4.11 Call trace

CUPS sessions support the call-trace functionality for both PFCP messages and uplink IBCP messages.



Note: In a CUPS context, the **connectivity-management** application option covers uplink IBCP messages, and the **pfcp** application option covers PFCP messages. Protocol-specific applications, for example **radius-auth**, **radius-acct**, **gx**, and **ppp-event**, are not applicable because the BNG-UP does not handle these protocols.

For CUPS sessions, call trace is enabled on the MAG-c CP; it is not directly enabled on the BNG-UP. When the MAG-c initiates a call-trace session, it sends the name of the call-trace profile in a PFCP message to the BNG-UP. The BNG-UP starts a call-trace session of its own, if it is configured with a matching profile name or a default profile:

- **A matching profile name is configured on the BNG-UP**

The BNG-UP uses the configured matching profile to start a call-trace session. Use the following command to configure a call-trace profile for the BNG-UP.

```
configure call-trace trace-profile name
```

- **A matching profile name is not configured on the BNG-UP but a default profile is configured**

If there is no matching profile, but a default call-trace profile is configured, the BNG-UP uses the default profile to start a call-trace session and generate a log. Use the following command to configure a default call-trace profile for the BNG-UP.

```
configure subscriber-mgmt pfcp association default-calltrace-profile
```

- **A matching profile name and a default profile are not configured on the BNG-UP**

The BNG-UP does not start a call-trace session and does not log an error.

After a call-trace session starts on the BNG-UP, it continues to be active until any of the following events occur:

- The BNG-UP receives a PFCP Session Delete Request from the MAG-c.
- The MAG-c disables call trace for a session and excludes the call-trace profile from the PFCP Session Modification Request to the BNG-UP.
- The call-trace session times out or reaches the size limit configured for the BNG-UP.

See the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide* for more information about call trace.

5 Network Address Translation

This chapter provides an overview of Network Address Translation (NAT) functionality for BNG CUPS.

5.1 Residential NAT for BNG CUPS

For BNG CUPS, NAT responsibilities are divided between the MAG-c and BNG-UP.

The role of the MAG-c is to associate the subscriber session with NAT during the session authentication phase. This process consists mainly of allocating the outside IP address and port-block to the NAT subscriber session. These parameters are submitted to the BNG-UP through the PFCP association.

The BNG-UP performs NAT on the data traffic. On the BNG-UP, NAT runs on ISAs, ESA-VMs or vISAs. For the inside IP addresses, the incoming data traffic is sprayed across ISAs or ESA-VMs. This traffic spraying is based on the subscriber context, which typically represents a residence. For the outside IP addresses, the NAT prefix that is received from the MAG-c is segmented into smaller subnets and equally distributed across ISAs. This approach requires fair load distribution of traffic across service adapters in the upstream and downstream directions.



Note: In this document, all service adapter types are referred to as ISAs, except when it is necessary to identify a specific type. See the *7450 ESS, 7750 SR, and VSR Multiservice Integrated Service Adapter and Extended Services Appliance Guide* for more information about MS-ISA service adapters.

See the *MAG-c Control Plane Function Guide* for more information about NAT terminology and an overview of Residential NAT that describes the division of NAT responsibilities between the MAG-c and BNG-UP.

5.2 UP NAT policy template

A UP NAT policy template contains parameters that define NAT behavior for a group of subscribers within a NAT pool. This includes configuring the maximum number of port-blocks per subscriber, the size of the extended port-blocks, support for ALGs, setting limits for the number of NAT flows per subscriber, protocol timer definitions, flow-based logging, watermarks, and so on. The UP NAT policy configuration allows the NAT behavior to be customized for different groups of subscribers within the same NAT pool.

Although the UP NAT policy template is configured on the BNG-UP, its assignment to the NAT-enabled session is performed on the MAG-c during the authentication phase, using a reference in the CP NAT profile configuration.

The roles of the CP NAT profile and UP NAT policy can be summarized as follows:

- The CP NAT profile is configured on the MAG-c and identifies NAT subscribers during the authentication phase. Parameters defined in the CP NAT profile affect the selection of the NAT pool within a specific outside routing context. This includes the allocation of the outside IP addresses, port-blocks, NAT mode of operation (NAPT or 1:1), size of the initial port-block, the number of subscriber per outside IP address and port-forwarding parameters. These resources are managed by the MAG-c.

- The UP NAT policy template is configured on the BNG-UP and is used to define NAT behavior for a group of subscribers within a NAT pool. This behavior is closer to the NAT translation in the forwarding plane (for example, ALGs and protocol timers).

5.3 Guidelines for configuring extended port blocks

In addition to configuring MAG-C for extended Port Blocks (PBs) (see *MAG-c Control Plane Function Guide, Multiple Port Blocks Per Subscriber*), it is necessary to configure the following two options in the UP NAT policy in BNG-UP.

- size of the extended PBs that a subscriber can allocate in a NAT pool
- total number of PBs (initial and extended) that a subscriber can allocate

Use the following commands to configure the maximum number of PBs per NAT subscriber and the maximum number of ports per extended PB.

```
configure service nat up-nat-policy block-limit
configure service nat up-nat-policy port-block-extension ports
```

Example

```
[ex:/configure service nat]
A:admin@node-2# info
  up-nat-policy "policy1" {
    block-limit 10
    port-block-extension {
      ports 335
      ...
    }
  }
}
```



Note: It is also necessary to configure the high and low watermarks for the extended port blocks; see [Guidelines for configuring watermarks](#).

5.4 Guidelines for configuring NAT subscribers in the sub-profile

Many NAT configuration parameters are defined in the UP NAT policy template (**up-nat-policy**) or the CP NAT profile (see [UP NAT policy template](#)). There are also some parameters that may be used for NAT configuration that require further granularity of definition, such as the UPNP policy that enables the dynamic port forward allocation.

Use the commands in the following context to configure a UPNP policy for NAT.

```
configure subscriber-management sub-profile upnp-policy
```

5.5 Guidelines for configuring NAT groups

A NAT group represents a collection of ISAs that are used to process NAT traffic for subscribers. NAT traffic is distributed over multiple ISAs in a NAT group to achieve better performance and scale. BNG CUPS supports a single NAT group per BNG-UP, however, other NAT groups can be configured in the system outside CUPS.

A NAT group is a mandatory configuration. After the NAT group is defined, it must be referenced by a PFCP association. A NAT group is configured using commands in the **configure isa nat-group** context.

See [Provisioning residential NAT for BNG CUPS](#) for a configuration example.

5.6 Guidelines for configuring accounting and logging

Aggregated NAT logging based on port blocks is performed on the MAG-c, and flow-based logging can be enabled on the BNG-UP. Because a number of logs are produced in flow logging, flow logs are exported directly from the ISA, bypassing the MAG-c and the CPM on the BNG-UP. The BNG-UP supports flow logging in IPFIX format.

Use the commands in the following context to configure an IPFIX export policy.

```
configure service ipfix export-policy
```

Use the following command to associate the configured export policy with a UP NAT policy.

```
configure service nat up-nat-policy flow-log-policy ipfix
```

5.7 Guidelines for configuring watermarks

The following watermarks are supported on the BNG-UP:

- The session-level watermarks on the member ISA level monitor the NAT flow usage against the configured limit per member ISA. They are configured using the NAT group. Use the **high** and **low** command options in the following context to configure the watermarks.

```
configure isa nat-group session-limits watermarks
```

- The session-level watermarks on the subscriber level monitor the NAT flows usage against the configured limit per subscriber. They are configured using the UP NAT policy. Use the **high** and **low** command options in the following context to configure the watermarks.

```
configure service nat up-nat-policy session-limits watermarks
```

- The port usage watermarks on the subscriber level monitor port usage against the configured limit per subscriber. They are configured using the UP NAT policy. Use the **high** and **low** command options in the following context to configure the watermarks.

```
configure service nat up-nat-policy port-limits watermarks
```

- Extended Port Blocks (PBs) are monitored per outside IP address and alerts are sent if subscribers (on that outside IP address) may soon be denied additional service because they are running out of extended PBs. This is configured per NAT UP policy.



Note:

- You can map multiple CP NAT profiles with the same NAT UP policy to the same NAT pool. Alternatively, the same NAT UP policy can point to two different pools, using different CP NAT profiles. In other words, this is monitored for all subscribers that use all instances of the referenced NAT UP policy.
- For the system to generate these events, the NAT log event `tmnxNatPIMemberExtBlockUsageHigh` (ID 2045) must be enabled.

Use the **high** and **low** command options in the following context to configure the watermarks to monitor utilization of the PB space per outside IP in a NAT pool reserved for extended PBs.

```
configure service nat up-nat-policy port-block-extension watermarks
```

- On the MAG-c, a watermark threshold can be configured in either absolute value or percentages to monitor micronet usage within a NAT outside pool. See the *MAG-c Control Plane Function Guide* for more information.

5.8 Guidelines for configuring intra-chassis redundancy

ISA redundancy on the BNG-UP level supports the following modes of operation:

- N:M active/standby mode**

M number of standby ISAs protect *N* number of active ISAs.

- all active mode**

This mode supports failure of up to two ISAs simultaneously. During an ISA failure, the configuration from the failed ISA is distributed over the remaining operational ISAs.

Both modes are stateless which means that NAT binding must be re-established after the switchover.

Use the commands in the following context to configure ISA redundancy.

```
configure isa nat-group
```

The commands in the preceding context associate MDAs with the NAT group, set the mode of operation to active/standby, and configure the number of active ISAs in the NAT group. Any ISAs within the NAT group that are in excess of the configured number are automatically considered standby.

Use the following command to enable active and standby mode.

```
configure isa nat-group redundancy intra-chassis active-standby
```

Use the following commands to enable the all-active mode.

```
configure isa nat-group redundancy intra-chassis active-active
configure isa nat-group redundancy intra-chassis active-active failed-mda-limit
```

5.9 Provisioning residential NAT for BNG CUPS

Prerequisites

- Review the overview information for residential NAT for BNG CUPS; see [Network Address Translation](#).
- A UP NAT policy is required. You can configure a policy for the BNG-UP or use the default settings. See [Guidelines for configuring NAT groups](#).

To configure residential NAT on BNG CUPS, perform the following minimum configuration steps:

Procedure

Step 1. Configure the MAG-c as described in the *MAG-c Control Plane Function Guide*.

Step 2. Configure the BNG-UP.

a. Configure the NAT policy template.

```
configure {
  service {
    nat {
      up-nat-policy "pol-1" {
      }
    }
  }
}
```

b. Configure the NAT group, including the ISA redundancy mode.

```
configure {
  isa {
    nat-group 1 {
      mda 1/2
      mda 3/1
      mda 2/2
      redundancy {
        active-mds-limit 2
        intra-chassis {
          active-standby
        }
      }
    }
  }
}
```

c. Associate the NAT group created in step 2.b with the PFCP interface.

```
configure {
  subscriber-mgmt {
    pfc {
      association "profile-1" {
        nat {
          nat-group 1
        }
      }
    }
  }
}
```

6 BNG-UP resiliency

This chapter provides information about BNG-UP resiliency.

6.1 Resiliency based on Fate Sharing Group

The MAG-c groups the sessions in FSGs. All sessions in an FSG share their fate, that is, they become active or standby together. The MAG-c provides the following parameters to the BNG-UP per FSG:

- FSG ID
- status (active or standby)
- unique FSG MAC address also known as a virtual MAC
- list of associated sessions and one or more aggregate routes associated with these sessions. For more information, see [IP gateway, services, and routing](#).
- FSG template

When the MAG-c does not provide an FSG template, the template with the name default is used. If there is no default template, the setup of the FSG and any associated session fails.

Use the following command to configure FSG templates.

```
configure subscriber-mgmt up-resiliency fate-sharing-group-template
```

After this, the active BNG-UP and standby BNG-UP are used in the context of a single FSG. Each BNG-UP can have multiple FSGs and can have a different status for each FSG.

To attract traffic from the access network, an active BNG-UP replies to ARP requests or ND messages for any IP gateway associated with the FSG. A standby BNG-UP never replies to those ARP or ND messages. To expedite convergence when switching from standby to active, the new active BNG-UP sends Gratuitous ARP (GARP) messages using the IP gateway address for the FSG, or the system IP address if no IP gateway is known. Afterward, the BNG-UP keeps sending periodic GARP messages to ensure traffic is attracted at all times to the correct BNG-UP.

Use the following command to configure the granularity of GARP messages for QinQ SAPs.

```
configure subscriber-mgmt up-resiliency fate-sharing-group-template gratuitous-arp
```

You can configure the BNG-UP to send a single GARP message per SAP or per outer tag. Configure the **fsg-active**, **fsg-active-path-restoration** and **fsg-standby** options for the following command to correctly draw traffic to the active BNG-UP.

```
configure policy-options policy-statement entry from state
```

All routes received from PFCP, including per-session framed routes, have one of these values as an option. You can use this option to adjust values in routing export policies; for example, adjust a metric or a preference to the needs of the used routing protocol.

The following reduced configuration example shows a simplified policy that sets a metric of 100 for active routes, a metric of 150 for active routes while in headless, and a metric of 200 for standby routes.

Example: Policy statement configuration

```
[ex:/configure policy-options policy-statement "upf_resiliency_aware_export"]
A:admin@BNG-UPF# info
  entry 20 {
    from {
      origin pfc
      state fsg-active
    }
    action {
      action-type accept
      metric {
        set 100
      }
    }
  }
  entry 30 {
    from {
      origin pfc
      state fsg-active-path-restoration
    }
    action {
      action-type accept
      metric {
        set 150
      }
    }
  }
  entry 40 {
    from {
      origin pfc
      state fsg-standby
    }
    action {
      action-type accept
      metric {
        set 200
      }
    }
  }
}
```

An active BNG-UP always forwards traffic in both directions. It uses the FSG MAC as source MAC for downlink unicast traffic. A standby BNG-UP by default forwards downlink traffic using its local port MAC as source MAC and drops all received uplink traffic. You can modify the default behavior in the following ways:

- To shunt downlink traffic from the standby to the active BNG-UP and have the active BNG-UP forward that downlink traffic, do the following:
 - Use the following command to configure a redundant interface.

```
configure subscriber-mgmt up-resiliency fate-sharing-group-template redundant-interface
```

- Use the following command to configure the same shunt ID on the active and the standby BNG-UP for the applicable service.

```
configure service ies subscriber-mgmt multi-chassis-shunt-id
configure service vprn subscriber-mgmt multi-chassis-shunt-id
```


- Use the following command to enable forwarding of uplink traffic by the standby BNG-UP.

```
configure subscriber-mgmt up-resiliency fate-sharing-group-template uplink-forwarding-while-standby
```



Caution: Enabling the **uplink-forwarding-while-standby** command can lead to packet replication toward the core network. To prevent the possibility of packet replication toward the core network, provision the access network not to replicate unknown unicast packets to the BNG-UP.

When the standby BNG-UP forwards uplink traffic, it can significantly lower packet loss during transition scenarios. The following examples illustrate this benefit:

- If the current active BNG-UP fails and an access node detects this faster than the MAG-c (for example, using BFD), the access node can start sending packets to the standby BNG-UP before that BNG-UP has become active. When the **uplink-forwarding-while-standby** command is enabled, the uplink packets are not lost because the standby BNG-UP forwards them.
- During scheduled maintenance, the MAG-c can switch the roles of the active and standby BNG-UP while both are healthy. For some time, the access node continues to send packets to the previously active BNG-UP that has become the standby BNG-UP. When the **uplink-forwarding-while-standby** command is enabled, the uplink packets are not lost because the previously active BNG-UP still forwards them. When the **uplink-forwarding-while-standby** command is disabled, the previously active BNG-UP drops the packets until the access node learns the path to the new active BNG-UP (using GARP).

The resiliency based on FSG does not use the SRRP protocol, but the system internally consumes an SRRP instance for each unique combination of FSG, port, and group interface template. To avoid potential conflicts with pre-configured SRRP instance IDs, use the following command to define a range of SRRP instance IDs for the inter-BNG-UP resiliency functionality.

```
configure redundancy srrp auto-srrp-id-range
```

6.2 BNG-UP health reporting

The BNG-UP can send health reports to the MAG-c using PFCP Node Report messages. The MAG-c uses the health reports to determine the need for a BNG-UP status change (active or standby). Per FSG, the MAG-c selects the active and the standby BNG-UP. For example, the MAG-c can base its decision on link failures in the access network.

The BNG-UP supports health reports for the following contexts:

- **per network instance**

Use the commands in the following context to configure the health monitoring for the applicable service.

```
configure service ies subscriber-mgmt up-resiliency
configure service vprn subscriber-mgmt up-resiliency
```

The health reports per network instance can, for example, be used to indicate the status of the network where the subscriber is serviced.

- **per Layer 2 access ID**

Use the commands in the following context to configure health monitoring.

```
configure service vpls capture-sap pfc up-resiliency
```

The health reports per Layer 2 access ID can, for example, be used to indicate the status of the access links.

Each health report generates a health value between 0 (unhealthy) and 100 (healthy). The base health value is 100 and decreases with the number of failed members in the operation group x the configured health drop number for the operational group.

Whenever a member of the operational group changes its state (fails or recovers), the BNG-UP calculates the health value and sends an updated report to the MAG-c.

To configure the operational group and the health drop number, use the **monitor-oper-group** and the **monitor-oper-group health-drop** commands in the previously mentioned contexts.

For more information about operational groups, see *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN*, sections *Object grouping and state monitoring*.

With the following example configuration, the BNG-UP sends health reports for Layer 2 access ID (port) lag-access. The operational group has five members (port 1/1/20 to port 1/1/24) and the health value decreases with 51 per failed member, that is, with 20% of the base health value.

Example: Configuration of health reports for Layer 2 access ID (port) lag-access

```
[ex:/configure port 1/1/20]
A:admin@BNG-UPF# info
    oper-group "lag_access_health"
[ex:/configure port 1/1/21]
A:admin@BNG-UPF# info
    oper-group "lag_access_health"
[ex:/configure port 1/1/22]
A:admin@BNG-UPF# info
    oper-group "lag_access_health"
[ex:/configure port 1/1/23]
A:admin@BNG-UPF# info
    oper-group "lag_access_health"
[ex:/configure port 1/1/24]
A:admin@BNG-UPF# info
    oper-group "lag_access_health"
[ex:/configure service vpls "access" capture-sap lag-access:*. * pfc up-resiliency]
A:admin@BNG-UPF# info
    monitor-oper-group "lag_access_health" {
        health-drop 20
    }
```

The BNG-UP sends a health report for every status change in the operational group. Additionally, it sends all health reports periodically (every 60 seconds) and when a PFCP audit is requested.

With the following example configuration, the BNG-UP sends health reports for network instance HSI. The health drop number is not configured, that is, the default value of 100 is used. The health is based on a BFD session that is used to check if the BNG-UP is isolated from the rest of the network. When the BFD session is up, the health value equals 100, otherwise, the health value equals 0.

Example: Configuration of health reports for network instance HSI

```
[ex:/configure service oper-group "hsi-bfd"]
A:admin@BNG-UPF# info
    bfd-liveness {
```

```

router-instance "to_uplink_router"
  interface-name "endpoint"
  dest-ip 203.0.113.10
}
[ex:/configure service vprn "hsi" subscriber-mgmt up-resiliency]
A:admin@BNG-UPF# info
  monitor-oper-group "hsi-bfd" {
}

```

6.3 Interaction with headless mode

If the PFCP path between the BNG-UP and the MAG-c fails and the BNG-UP becomes headless, use the following command to determine the behavior for active FSGs.

```
configure subscriber-mgmt up-resiliency fate-sharing-group-template path-restoration-state
```

When using the **standby** option for this command, the BNG-UP forces all active FSGs to become standby. This avoids any possibility of active/active UP behavior. However, if all the BNG-UPs become headless, for example, because of a routing issue to the MAG-c, all FSGs on all BNG-UPs become standby, and no forwarding is possible.



Note: Nokia recommends leaving this command set to the default (**auto**). Only enable **standby** if the network cannot handle the described active/active scenarios and avoidance.

If you use the **auto** option, the BNG-UP uses an heuristic process to decide whether to keep the FSG as active or move it to standby. The BNG-UP autonomously changes FSGs to standby if any of the following conditions are met:

- No single network instance is monitored for health.
- At least one of the monitored network instances indicates health failure.
- No GARP messages are snooped from another BNG-UP.

This process detects the difference between an isolated BNG-UP becoming headless, or all BNG-UPs becoming headless. When the BNG-UP estimates that all BNG-UPs are headless, it keeps the FSGs active. Alternatively, the BNG-UP keeps FSGs as standby, because the MAG-c activates another BNG-UP.

The heuristic check of network health determines if the failure is a more generic network failure, which is more likely to be BNG-UP local (for example, a network link failure). If the PFCP fails but the local network is fine, the failure is probably central, and all BNG-UPs became headless. In addition, if the network link is down the system probably cannot forward the session traffic anyway.

The heuristic check of GARP snooping is used to determine if another BNG-UP became active while this UP is headless. If another BNG-UP sends a GARP message this means it was updated by the MAG-c to become active, which in turns means it cannot be headless. Because of this it can be estimated that the headless mode is contained to a single node and it is safe to become standby.

These heuristic checks are best-effort and may fail to detect active/active conditions. However, by correctly setting routing metrics to differentiate between the **fsg-active** and **fsg-active-path-restoration** options, you can avoid the worst of active/active scenarios. By giving the headless BNG-UP a worse metric or preference, only the non-headless active BNG-UP draws downlink traffic. The headless BNG-UP may still erroneously answer ARP requests and update forwarding databases in the access node for the FSG virtual MAC. However, typically this is quickly corrected by downlink traffic using the vMAC as the source MAC address, coming from the non-headless BNG-UP.

The following are other risks of active/active scenarios:

- When the aggregation network replicates unknown unicast packets to both BNG-UPs, it forwards these packets twice, leading to duplicate packets in the network. However, it is unlikely that the FSG virtual MAC is ever an unknown MAC. To further reduce the risk, disable unicast replication.
- Exact QoS enforcement is not guaranteed in this situation.

The standby FSGs remain as-is during headless conditions. Whichever option you use (**standby** or **auto**), the BNG-UP reverts the FSG state to active after the headless condition is no longer valid.

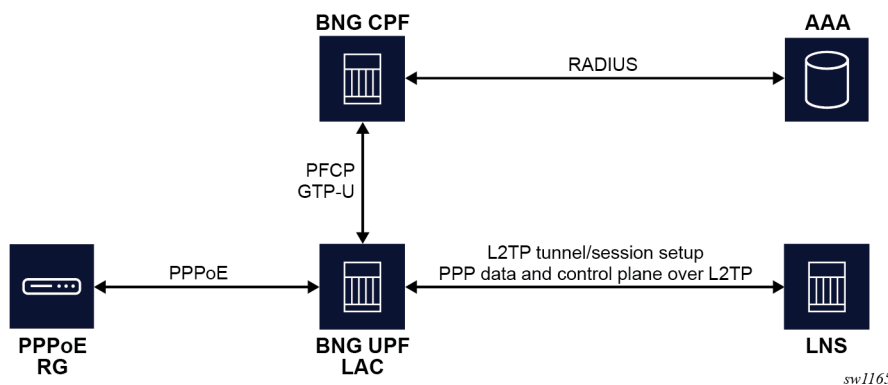
7 Layer 2 Tunneling Protocol

The BNG CUPS environment supports L2TP functionality.

7.1 BNG-UP-triggered L2TP access concentrator

The BNG-UP supports PPPoE LAC sessions with L2TP control plane signaling running on the BNG-UP. This relies on the SR L2TP functionality described in the *7450 ESS*, *7750 SR*, and *VSR Triple Play Service Delivery Architecture Guide*, section *Layer 2 Tunneling Protocol (L2TP)*.

Figure 2: L2TP LAC network components



To run a PPPoE LAC session with L2TP control plane signaling on the BNG-UP, the MAG-c provides the following parameters during the PFCP session data plane setup:

- list of potential L2TP Network Server (LNS) tunnel endpoints
- common or per-LNS parameters such as retry timers, preferences, and authorization IDs; see the *MAG-c Control Plane Function Guide* for more information about how to configure these parameters
- per-session PPP LCP and PPP authentication attributes to be sent by proxy to the chosen LNS, so the LNS can avoid renegotiating LCP and authentication with the PPPoE session
- unique name to identify the tunnel group over shared sessions

The BNG-UP performs normal L2TP tunnel management over the group of tunnels, including deny list management, as described in the *7450 ESS*, *7750 SR*, and *VSR Triple Play Service Delivery Architecture Guide*, section *Traffic Steering on L2TP LAC*. To accomplish this, the BNG-UP creates an internal L2TP group in the VPRN service or base router provided by the PFCP. This internal group is not configurable, but common parameters configured under **configure router l2tp** or **configure vprn service-name l2tp** are applied to the setup of the tunnels. If a parameter is configured locally and also in PFCP, the PFCP parameter takes precedence. To set up tunnels, the **l2tp** context must also be administratively enabled in the router or VPRN.

Based on the tunnel management, the BNG-UP selects a tunnel for the session, or creates the tunnel if it does not exist, and an L2TP session is created for that specific PPPoE session. After the tunnel and

session setup are complete, a response including the following parameters is sent to the triggering PFCP message for operational management on the MAG-c:

- L2TP tunnel and session IDs of both the LAC and the LNS
- LNS and LAC hostnames, also known as auth IDs
- L2TP tunnel and session IDs for both LNS and LAC
- Call Serial Number (CSN)

After the connection is established, the PPP data packets are forwarded to the LNS. The PFCP message instructs the BNG-UP to stop sending PPP control plane messages to the MAG-c, and to send them to the LNS instead. Only the PPPoE discovery packets are forwarded to the MAG-c. For example, a PADT packet is handled by the MAG-c, while an LCP terminate packet is handled by the LNS.

If L2TP session or tunnel setup fails, the normal L2TP backoff and deny list behavior applies. If the setup fails for all provided tunnels, an explicit PFCP error message is sent to the MAG-c. However, the BNG-UP does not time out the PFCP message if the setup takes too long, for example, because multiple LNSs are unreachable. The BNG triggers an explicit delete after the PFCP message times out, using its local PFCP N1 or T1 configuration. If a delete is received for an in-progress setup, the setup is aborted.



Note: An aborted setup does not count as a tunnel timeout and the tunnel is not placed in automatic deny lists for timeout hold-off purposes. You must configure the **max-retries-not-estab** command to ensure that the total timeout does not exceed the PFCP session timeout on the MAG-c. See the *MAG-c Control Plane Function Guide* for more information about how to correctly align the timeouts.

8 Lawful intercept

This chapter provides an overview of the Lawful Intercept (LI) functionality for BNG CUPS.

8.1 Overview of the LI implementation on the BNG-UP

To perform LI for BNG CUPS, configurations are required on both the MAG-c and BNG-UP:

- The MAG-c supports reporting of subscriber and LI events; see the *MAG-c Control Plane Function Guide* and the *MAG-c CLI Reference Guide* for more information about MAG-c configuration.
- The BNG-UP supports the provisioning of LI targets and mirroring of LI packets.

After the LI mediation gateway identifies an LI subscriber through the MAG-c-reported events, the provisioning of the LI subscriber can be performed directly on the BNG-UP, using the **configure li li-source** commands as described in the *7450 ESS, 7750 SR, 7950 XRS, and VSR OAM and Diagnostics Guide*. Provisioning can be performed using CLI or SNMPv3, however, SNMPv3 is the preferred platform.



Note: In CUPS architecture, the BNG-UP creates a new subscriber ID every time the subscriber or LI subscriber logs in. For this reason, it is highly recommended to use a mediation device to automate the LI configuration. It is not recommended to perform LI configuration through CLI on the BNG-UP.

When the MAG-c is configured, it notifies the LI mediation device about the UP subscriber IDs and IP addresses. The LI mediation device sends an SNMPv3 command directly to the BNG-UP IP address to set up an LI target. Li targets typically include the following parameters:

- mirror destination service; can be a layer 3 encapsulation or a SAP
- subscriber ID; for example, "_cups_549"



Note: The BNG-UP automatically appends "_cups_" to the auto-generated subscriber ID.

- ingress and egress direction
- session ID and intercept ID, which allow the LI mediation device to correlate subscriber events and mirrored packets (optional); see the *7450 ESS, 7750 SR, 7950 XRS, and VSR OAM and Diagnostics Guide*, section "Lawful Intercept" for information about additional parameters

When the subscriber logs out, the LI mediation device removes the subscriber from the LI source through SNMPv3. When the same subscriber logs in again, the system auto-generates a new BNG-UP subscriber ID.

For the procedure to configure SNMPv3 and BNG CUPS, see [Provisioning SNMPv3 and LI subscribers for the BNG-UP](#)

8.2 Provisioning SNMPv3 and LI subscribers for the BNG-UP

Prerequisites

Before you begin, review [Overview of the LI implementation on the BNG-UP](#).

To provision SNMPv3 and LI subscribers for the BNG-UP, perform the following steps:

Procedure

- Step 1.** Create the SNMPv3 group for LI.
- Step 2.** Provision an LI administrator for the BNG-UP with both LI access and SNMP access.
- Step 3.** Associate the SNMPv3 group created in step 1 with the LI administrator.
See the *7450 ESS, 7750 SR, 7950 XRS, and VSR System Management Guide*, for more information about LI users and SNMPv3 setup.
- Step 4.** Provision the LI subscriber directly on the BNG-UP, using the **configure li li-source** commands.
See the *7450 ESS, 7750 SR, 7950 XRS, and VSR OAM and Diagnostics Guide* for information about user plane LI management and procedures.
See the *MAG-c Control Plane Function Guide* and the *MAG-c CLI Reference Guide* for information about the related MAG-c configuration.

9 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.

9.1 Access Node Control Protocol (ANCP)

draft-ietf-ancp-protocol-02, *Protocol for Access Node Control Mechanism in Broadband Networks*

RFC 5851, *Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks*

9.2 Bidirectional Forwarding Detection (BFD)

draft-ietf-lsr-ospf-bfd-strict-mode-10, *OSPF BFD Strict-Mode*

RFC 5880, *Bidirectional Forwarding Detection (BFD)*

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

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

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

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

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

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

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

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

RFC 9247, *BGP - Link State (BGP-LS) Extensions for Seamless Bidirectional Forwarding Detection (S-BFD)*

9.3 Border Gateway Protocol (BGP)

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

draft-hares-idr-update-attrib-low-bits-fix-01, *Update Attribute Flag Low Bits Clarification*

draft-ietf-idr-add-paths-guidelines-08, *Best Practices for Advertisement of Multiple Paths in IBGP*

draft-ietf-idr-best-external-03, *Advertisement of the best external route in BGP*

draft-ietf-idr-bgp-flowspec-oid-03, *Revised Validation Procedure for BGP Flow Specifications*
draft-ietf-idr-bgp-gr-notification-01, *Notification Message support for BGP Graceful Restart*
draft-ietf-idr-bgp-ls-app-specific-attr-16, *Application-Specific Attributes Advertisement with BGP Link-State*
draft-ietf-idr-bgp-ls-flex-algo-06, *Flexible Algorithm Definition Advertisement with BGP Link-State*
draft-ietf-idr-bgp-optimal-route-reflection-10, *BGP Optimal Route Reflection (BGP-ORR)*
draft-ietf-idr-error-handling-03, *Revised Error Handling for BGP UPDATE Messages*
draft-ietf-idr-flowspec-interfaceset-03, *Applying BGP flowspec rules on a specific interface set*
draft-ietf-idr-flowspec-path-redirect-05, *Flowspec Indirection-id Redirect – localised ID*
draft-ietf-idr-flowspec-redirect-ip-02, *BGP Flow-Spec Redirect to IP Action*
draft-ietf-idr-link-bandwidth-03, *BGP Link Bandwidth Extended Community*
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 7674, *Clarification of the Flowspec Redirect Extended Community*
RFC 7854, *BGP Monitoring Protocol (BMP)*
RFC 7911, *Advertisement of Multiple Paths in BGP*
RFC 7999, *BLACKHOLE Community*
RFC 8092, *BGP Large Communities Attribute*
RFC 8097, *BGP Prefix Origin Validation State Extended Community*
RFC 8212, *Default External BGP (EBGP) Route Propagation Behavior without Policies*
RFC 8277, *Using BGP to Bind MPLS Labels to Address Prefixes*
RFC 8571, *BGP - Link State (BGP-LS) Advertisement of IGP Traffic Engineering Performance Metric Extensions*
RFC 8950, *Advertising IPv4 Network Layer Reachability Information (NLRI) with an IPv6 Next Hop*
RFC 8955, *Dissemination of Flow Specification Rules*
RFC 8956, *Dissemination of Flow Specification Rules for IPv6*
RFC 9086, *Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing BGP Egress Peer Engineering*
RFC 9494, *Long-Lived Graceful Restart for BGP*
RFC 9552, *Distribution of Link-State and Traffic Engineering Information Using BGP*

9.4 Bridging and management

IEEE 802.1AB, *Station and Media Access Control Connectivity Discovery*
IEEE 802.1ad, *Provider Bridges*
IEEE 802.1ag, *Connectivity Fault Management*
IEEE 802.1ah, *Provider Backbone Bridges*
IEEE 802.1ak, *Multiple Registration Protocol*
IEEE 802.1aq, *Shortest Path Bridging*
IEEE 802.1AX, *Link Aggregation*
IEEE 802.1D, *MAC Bridges*
IEEE 802.1p, *Traffic Class Expediting*
IEEE 802.1Q, *Virtual LANs*

IEEE 802.1s, *Multiple Spanning Trees*
IEEE 802.1w, *Rapid Reconfiguration of Spanning Tree*
IEEE 802.1X, *Port Based Network Access Control*

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

3GPP TS 23.003, *Numbering, addressing and identification*
3GPP TS 23.007, *Restoration procedures*
3GPP TS 23.402, *Architecture enhancements for non-3GPP accesses – S2a roaming based on GPRS*
3GPP TS 23.501, *System architecture for the 5G System (5GS)*
3GPP TS 23.502, *Procedures for the 5G System (5GS)*
3GPP TS 23.503, *Policy and charging control framework for the 5G System (5GS)*
3GPP TS 24.501, *Non-Access-Stratum (NAS) protocol for 5G System (5GS)*
3GPP TS 29.244, *Interface between the Control Plane and the User Plane nodes*
3GPP TS 29.281, *General Packet Radio System (GPRS) Tunnelling Protocol User Plane (GTPv1-U)*
3GPP TS 29.500, *Technical Realization of Service Based Architecture*
3GPP TS 29.501, *Principles and Guidelines for Services Definition*
3GPP TS 29.502, *Session Management Services*
3GPP TS 29.503, *Unified Data Management Services*
3GPP TS 29.512, *Session Management Policy Control Service*
3GPP TS 29.518, *Access and Mobility Management Services*
3GPP TS 32.255, *5G data connectivity domain charging*
3GPP TS 32.290, *Services, operations and procedures of charging using Service Based Interface (SBI)*
3GPP TS 32.291, *5G system, charging service*
BBF TR-459, *Control and User Plane Separation for a Disaggregated BNG*
BBF TR-459.2, *Multi-Service Disaggregated BNG with CUPS: Integrated Carrier Grade NAT function*
RFC 8300, *Network Service Header (NSH)*
RFC 8910, *Captive-Portal Identification in DHCP and Router Advertisements (RAs)*

9.6 Certificate management

RFC 4210, *Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)*
RFC 4211, *Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)*
RFC 5280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*

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

RFC 7030, *Enrollment over Secure Transport*

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

9.7 Circuit emulation

RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)*

RFC 5086, *Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)*

RFC 5287, *Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks*

9.8 Ethernet

IEEE 802.3ah, *Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks*

IEEE 802.3x, *Ethernet Flow Control*

ITU-T G.8031/Y.1342, *Ethernet Linear Protection Switching*

ITU-T G.8032/Y.1344, *Ethernet Ring Protection Switching*

ITU-T Y.1731, *OAM functions and mechanisms for Ethernet based networks*

9.9 Ethernet VPN (EVPN)

draft-ietf-bess-bgp-srv6-args-00, *SRv6 Argument Signaling for BGP Services*

draft-ietf-bess-evpn-ip-aliasing-00, *EVPN Support for L3 Fast Convergence and Aliasing/Backup Path – IP Prefix routes*

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

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

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

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

draft-sr-bess-evpn-vpws-gateway-03, *Ethernet VPN Virtual Private Wire Services Gateway Solution*

RFC 7432, *BGP MPLS-Based Ethernet VPN*

RFC 7623, *Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)*

RFC 8214, *Virtual Private Wire Service Support in Ethernet VPN*

RFC 8317, *Ethernet-Tree (E-Tree) Support in Ethernet VPN (EVPN) an Provider Backbone Bridging EVPN (PBB-EVPN)*

RFC 8365, *A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)*

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

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

RFC 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)*

RFC 9541, *Flush Mechanism for Customer MAC Addresses Based on Service Instance Identifier (I-SID) in Provider Backbone Bridging EVPN (PBB-EVPN)*

RFC 9625, *EVPN Optimized Inter-Subnet Multicast (OISM) Forwarding – ingress replication and mLDP*

9.10 gRPC Remote Procedure Calls (gRPC)

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

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

gnmi.proto version 0.8.0, *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*

9.11 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 6119, *IPv6 Traffic Engineering in IS-IS*
RFC 6213, *IS-IS BFD-Enabled TLV*
RFC 6232, *Purge Originator Identification TLV for IS-IS*
RFC 6233, *IS-IS Registry Extension for Purges*
RFC 6329, *IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging*
RFC 7775, *IS-IS Route Preference for Extended IP and IPv6 Reachability*
RFC 7794, *IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability – 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*

9.12 Internet Protocol (IP) Fast Reroute (FRR)

draft-ietf-rtgwg-lfa-manageability-08, *Operational management of Loop Free Alternates*
RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*
RFC 7431, *Multicast-Only Fast Reroute*
RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*
RFC 8518, *Selection of Loop-Free Alternates for Multi-Homed Prefixes*

9.13 Internet Protocol (IP) general

draft-grant-tacacs-02, *The TACACS+ Protocol*

RFC 768, *User Datagram Protocol*

RFC 793, *Transmission Control Protocol*

RFC 854, *Telnet Protocol Specifications*

RFC 1350, *The TFTP Protocol (revision 2)*

RFC 2347, *TFTP Option Extension*

RFC 2348, *TFTP Blocksize Option*

RFC 2349, *TFTP Timeout Interval and Transfer Size Options*

RFC 2428, *FTP Extensions for IPv6 and NATs*

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

RFC 2784, *Generic Routing Encapsulation (GRE)*

RFC 2818, *HTTP Over TLS*

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

RFC 3164, *The BSD syslog Protocol*

RFC 4250, *The Secure Shell (SSH) Protocol Assigned Numbers*

RFC 4251, *The Secure Shell (SSH) Protocol Architecture*

RFC 4252, *The Secure Shell (SSH) Authentication Protocol – publickey, password*

RFC 4253, *The Secure Shell (SSH) Transport Layer Protocol*

RFC 4254, *The Secure Shell (SSH) Connection Protocol*

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

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

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

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

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

RFC 5289, *TLS Elliptic Curve Cipher Suites with SHA-256/384 and AES Galois Counter Mode (GCM)*

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

RFC 5656, *Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer – ECDSA*

RFC 5925, *The TCP Authentication Option*

RFC 5926, *Cryptographic Algorithms for the TCP Authentication Option (TCP-AO)*

RFC 6398, *IP Router Alert Considerations and Usage – MLD*

RFC 6528, *Defending against Sequence Number Attacks*

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

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

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

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

RFC 7232, *Hypertext Transfer Protocol (HTTP/1.1): Conditional Requests*
RFC 7301, *Transport Layer Security (TLS) Application Layer Protocol Negotiation Extension*
RFC 7616, *HTTP Digest Access Authentication*
RFC 8446, *The Transport Layer Security (TLS) Protocol Version 1.3*

9.14 Internet Protocol (IP) multicast

cisco-ipmulticast/pim-autorp-spec01, *Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast* – version 1
draft-ietf-bier-pim-signaling-08, *PIM Signaling Through BIER Core*
draft-ietf-idmr-traceroute-ipm-07, *A "traceroute" facility for IP Multicast*
draft-ietf-l2vpn-vpls-pim-snooping-07, *Protocol Independent Multicast (PIM) over Virtual Private LAN Service (VPLS)*
RFC 1112, *Host Extensions for IP Multicasting*
RFC 2236, *Internet Group Management Protocol, Version 2*
RFC 2365, *Administratively Scoped IP Multicast*
RFC 2375, *IPv6 Multicast Address Assignments*
RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*
RFC 3306, *Unicast-Prefix-based IPv6 Multicast Addresses*
RFC 3376, *Internet Group Management Protocol, Version 3*
RFC 3446, *Anycast 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 3956, *Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address*
RFC 3973, *Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised) – auto-RP groups*
RFC 4541, *Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches*
RFC 4604, *Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast*
RFC 4607, *Source-Specific Multicast for IP*
RFC 4608, *Source-Specific Protocol Independent Multicast in 232/8*
RFC 4610, *Anycast-RP Using Protocol Independent Multicast (PIM)*
RFC 4611, *Multicast Source Discovery Protocol (MSDP) Deployment Scenarios*
RFC 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*

RFC 5186, *Internet Group Management Protocol Version 3 (IGMPv3) / Multicast Listener Discovery Version 2 (MLDv2) and Multicast Routing Protocol Interaction*

RFC 5384, *The Protocol Independent Multicast (PIM) Join Attribute Format*

RFC 5496, *The Reverse Path Forwarding (RPF) Vector TLV*

RFC 6037, *Cisco Systems' Solution for Multicast in MPLS/BGP IP VPNs*

RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*

RFC 6513, *Multicast in MPLS/BGP IP VPNs*

RFC 6514, *BGP Encodings and Procedures for Multicast in MPLS/IP VPNs*

RFC 6515, *IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs*

RFC 6516, *IPv6 Multicast VPN (MVPN) Support Using PIM Control Plane and Selective Provider Multicast Service Interface (S-PMSI) Join Messages*

RFC 6625, *Wildcards in Multicast VPN Auto-Discover Routes*

RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path*

RFC 7246, *Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context*

RFC 7385, *IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points*

RFC 7716, *Global Table Multicast with BGP Multicast VPN (BGP-MVPN) Procedures*

RFC 7761, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*

RFC 8279, *Multicast Using Bit Index Explicit Replication (BIER)*

RFC 8296, *Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks – MPLS encapsulation*

RFC 8401, *Bit Index Explicit Replication (BIER) Support via IS-IS*

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

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

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

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

RFC 9573, *MVPN/EVPN Tunnel Aggregation with Common Labels – DCB and static service labels*

9.15 Internet Protocol (IP) version 4

RFC 791, *Internet Protocol*

RFC 792, *Internet Control Message Protocol*

RFC 826, *An Ethernet Address Resolution Protocol*

RFC 951, *Bootstrap Protocol (BOOTP) – relay*

RFC 1034, *Domain Names - Concepts and Facilities*

RFC 1035, *Domain Names - Implementation and Specification*

RFC 1191, *Path MTU Discovery – router specification*

RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*
RFC 1534, *Interoperation between DHCP and BOOTP*
RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*
RFC 1812, *Requirements for IPv4 Routers*
RFC 1918, *Address Allocation for Private Internets*
RFC 2003, *IP Encapsulation within IP*
RFC 2131, *Dynamic Host Configuration Protocol*
RFC 2132, *DHCP Options and BOOTP Vendor Extensions*
RFC 2401, *Security Architecture for Internet Protocol*
RFC 3021, *Using 31-Bit Prefixes on IPv4 Point-to-Point Links*
RFC 3046, *DHCP Relay Agent Information Option (Option 82)*
RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*
RFC 4884, *Extended ICMP to Support Multi-Part Messages – ICMPv4 and ICMPv6 Time Exceeded*

9.16 Internet Protocol (IP) version 6

RFC 2464, *Transmission of IPv6 Packets over Ethernet Networks*
RFC 2529, *Transmission of IPv6 over IPv4 Domains without Explicit Tunnels*
RFC 3122, *Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification*
RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*
RFC 3587, *IPv6 Global Unicast Address Format*
RFC 3596, *DNS Extensions to Support IP version 6*
RFC 3633, *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*
RFC 3646, *DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*
RFC 3736, *Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6*
RFC 3971, *SEcure Neighbor Discovery (SEND)*
RFC 3972, *Cryptographically Generated Addresses (CGA)*
RFC 4007, *IPv6 Scoped Address Architecture*
RFC 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 4862, *IPv6 Stateless Address Autoconfiguration – router functions*
RFC 4890, *Recommendations for Filtering ICMPv6 Messages in Firewalls*
RFC 4941, *Privacy Extensions for Stateless Address Autoconfiguration in IPv6*

RFC 5007, *DHCPv6 Leasequery*
RFC 5095, *Deprecation of Type 0 Routing Headers in IPv6*
RFC 5722, *Handling of Overlapping IPv6 Fragments*
RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6 – IPv6*
RFC 5952, *A Recommendation for IPv6 Address Text Representation*
RFC 6092, *Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service – Internet Control and Management, Upper-Layer Transport Protocols, UDP Filters, IPsec and Internet Key Exchange (IKE), TCP Filters*
RFC 6106, *IPv6 Router Advertisement Options for DNS Configuration*
RFC 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*
RFC 6221, *Lightweight DHCPv6 Relay Agent*
RFC 6437, *IPv6 Flow Label Specification*
RFC 6603, *Prefix Exclude Option for DHCPv6-based Prefix Delegation*
RFC 8021, *Generation of IPv6 Atomic Fragments Considered Harmful*
RFC 8200, *Internet Protocol, Version 6 (IPv6) Specification*
RFC 8201, *Path MTU Discovery for IP version 6*

9.17 Internet Protocol Security (IPsec)

draft-ietf-ipsec-isakmp-mode-cfg-05, *The ISAKMP Configuration Method*
draft-ietf-ipsec-isakmp-xauth-06, *Extended Authentication within ISAKMP/Oakley (XAUTH)*
RFC 2401, *Security Architecture for the Internet Protocol*
RFC 2403, *The Use of HMAC-MD5-96 within ESP and AH*
RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH*
RFC 2405, *The ESP DES-CBC Cipher Algorithm With Explicit IV*
RFC 2406, *IP Encapsulating Security Payload (ESP)*
RFC 2407, *IPsec Domain of Interpretation for ISAKMP (IPsec DoI)*
RFC 2408, *Internet Security Association and Key Management Protocol (ISAKMP)*
RFC 2409, *The Internet Key Exchange (IKE)*
RFC 2410, *The NULL Encryption Algorithm and Its Use With IPsec*
RFC 2560, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*
RFC 3526, *More Modular Exponential (MODP) Diffie-Hellman group for Internet Key Exchange (IKE)*
RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*
RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*
RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*
RFC 3947, *Negotiation of NAT-Traversal in the IKE*
RFC 3948, *UDP Encapsulation of IPsec ESP Packets*

RFC 4106, *The Use of Galois/Counter Mode (GCM) in IPsec ESP*
RFC 4109, *Algorithms for Internet Key Exchange version 1 (IKEv1)*
RFC 4301, *Security Architecture for the Internet Protocol*
RFC 4303, *IP Encapsulating Security Payload*
RFC 4307, *Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)*
RFC 4308, *Cryptographic Suites for IPsec*
RFC 4434, *The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE)*
RFC 4543, *The Use of Galois Message Authentication Code (GMAC) in IPsec ESP and AH*
RFC 4754, *IKE and IKEv2 Authentication Using the Elliptic Curve Digital Signature Algorithm (ECDSA)*
RFC 4835, *Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)*
RFC 4868, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec*
RFC 4945, *The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2 and PKIX*
RFC 5019, *The Lightweight Online Certificate Status Protocol (OCSP) Profile for High-Volume Environments*
RFC 5282, *Using Authenticated Encryption Algorithms with the Encrypted Payload of the IKEv2 Protocol*
RFC 5903, *ECP Groups for IKE and IKEv2*
RFC 5996, *Internet Key Exchange Protocol Version 2 (IKEv2)*
RFC 5998, *An Extension for EAP-Only Authentication in IKEv2*
RFC 6379, *Suite B Cryptographic Suites for IPsec*
RFC 6380, *Suite B Profile for Internet Protocol Security (IPsec)*
RFC 6960, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*
RFC 7296, *Internet Key Exchange Protocol Version 2 (IKEv2)*
RFC 7321, *Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)*
RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*
RFC 7427, *Signature Authentication in the Internet Key Exchange Version 2 (IKEv2)*
RFC 8784, *Mixing Preshared Keys in the Internet Key Exchange Protocol Version 2 (IKEv2) for Post-quantum Security*

9.18 Label Distribution Protocol (LDP)

draft-pdutta-mpls-ldp-adj-capability-00, *LDP Adjacency Capabilities*
draft-pdutta-mpls-ldp-v2-00, *LDP Version 2*
draft-pdutta-mpls-mldp-up-redundancy-00, *Upstream LSR Redundancy for Multi-point LDP Tunnels*
draft-pdutta-mpls-multi-ldp-instance-00, *Multiple LDP Instances*
draft-pdutta-mpls-tldp-hello-reduce-04, *Targeted LDP Hello Reduction*

RFC 3037, *LDP Applicability*
RFC 3478, *Graceful Restart Mechanism for Label Distribution Protocol – helper mode*
RFC 5036, *LDP Specification*
RFC 5283, *LDP Extension for Inter-Area Label Switched Paths (LSPs)*
RFC 5443, *LDP IGP Synchronization*
RFC 5561, *LDP Capabilities*
RFC 5919, *Signaling LDP Label Advertisement Completion*
RFC 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*
RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*
RFC 6826, *Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*
RFC 7032, *LDP Downstream-on-Demand in Seamless MPLS*
RFC 7473, *Controlling State Advertisements of Non-negotiated LDP Applications*
RFC 7552, *Updates to LDP for IPv6*

9.19 Layer Two Tunneling Protocol (L2TP) Network Server (LNS)

draft-mammoliti-l2tp-accessline-avp-04, *Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions*
RFC 2661, *Layer Two Tunneling Protocol "L2TP"*
RFC 2809, *Implementation of L2TP Compulsory Tunneling via RADIUS*
RFC 3438, *Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers: Internet Assigned Numbers Authority (IANA) Considerations Update*
RFC 3931, *Layer Two Tunneling Protocol - Version 3 (L2TPv3)*
RFC 4719, *Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)*
RFC 4951, *Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"*

9.20 Multiprotocol Label Switching (MPLS)

draft-ietf-mpls-lsp-ping-ospfv3-codepoint-02, *OSPFv3 CodePoint for MPLS LSP Ping*
RFC 3031, *Multiprotocol Label Switching Architecture*
RFC 3032, *MPLS Label Stack Encoding*
RFC 3270, *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services – E-LSP*
RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*
RFC 4023, *Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)*
RFC 4182, *Removing a Restriction on the use of MPLS Explicit NULL*

RFC 4950, *ICMP Extensions for Multiprotocol Label Switching*
RFC 5332, *MPLS Multicast Encapsulations*
RFC 5884, *Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)*
RFC 6374, *Packet Loss and Delay Measurement for MPLS Networks – Delay Measurement, Channel Type 0x000C*
RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*
RFC 6425, *Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping*
RFC 6790, *The Use of Entropy Labels in MPLS Forwarding*
RFC 7308, *Extended Administrative Groups in MPLS Traffic Engineering (MPLS-TE)*
RFC 7510, *Encapsulating MPLS in UDP*
RFC 7746, *Label Switched Path (LSP) Self-Ping*
RFC 7876, *UDP Return Path for Packet Loss and Delay Measurement for MPLS Networks – Delay Measurement*
RFC 8029, *Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures*

9.21 Multiprotocol Label Switching - Transport Profile (MPLS-TP)

RFC 5586, *MPLS Generic Associated Channel*
RFC 5921, *A Framework for MPLS in Transport Networks*
RFC 5960, *MPLS Transport Profile Data Plane Architecture*
RFC 6370, *MPLS Transport Profile (MPLS-TP) Identifiers*
RFC 6378, *MPLS Transport Profile (MPLS-TP) Linear Protection*
RFC 6426, *MPLS On-Demand Connectivity and Route Tracing*
RFC 6427, *MPLS Fault Management Operations, Administration, and Maintenance (OAM)*
RFC 6428, *Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile*
RFC 6478, *Pseudowire Status for Static Pseudowires*
RFC 7213, *MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing*

9.22 Network Address Translation (NAT)

draft-ietf-behave-address-format-10, *IPv6 Addressing of IPv4/IPv6 Translators*
draft-ietf-behave-v6v4-xlate-23, *IP/ICMP Translation Algorithm*
draft-miles-behave-l2nat-00, *Layer2-Aware NAT*
draft-nishitani-cgn-02, *Common Functions of Large Scale NAT (LSN)*
RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*

RFC 5382, *NAT Behavioral Requirements for TCP*
RFC 5508, *NAT Behavioral Requirements for ICMP*
RFC 6146, *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*
RFC 6333, *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*
RFC 6334, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite*
RFC 6887, *Port Control Protocol (PCP)*
RFC 6888, *Common Requirements For Carrier-Grade NATs (CGNs)*
RFC 7753, *Port Control Protocol (PCP) Extension for Port-Set Allocation*
RFC 7915, *IP/ICMP Translation Algorithm*

9.23 Network Configuration Protocol (NETCONF)

RFC 5277, *NETCONF Event Notifications*
RFC 6020, *YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)*
RFC 6022, *YANG Module for NETCONF Monitoring*
RFC 6241, *Network Configuration Protocol (NETCONF)*
RFC 6242, *Using the NETCONF Protocol over Secure Shell (SSH)*
RFC 6243, *With-defaults Capability for NETCONF*
RFC 8071, *NETCONF Call Home and RESTCONF Call Home – NETCONF*
RFC 8342, *Network Management Datastore Architecture (NMDA) – Startup, Candidate, Running and Intended datastores*
RFC 8525, *YANG Library*
RFC 8526, *NETCONF Extensions to Support the Network Management Datastore Architecture – <get-data> operation*

9.24 Open Shortest Path First (OSPF)

RFC 1765, *OSPF Database Overflow*
RFC 2328, *OSPF Version 2*
RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*
RFC 3509, *Alternative Implementations of OSPF Area Border Routers*
RFC 3623, *Graceful OSPF Restart Graceful OSPF Restart – helper mode*
RFC 3630, *Traffic Engineering (TE) Extensions to OSPF Version 2*
RFC 4222, *Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance*
RFC 4552, *Authentication/Confidentiality for OSPFv3*
RFC 4576, *Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 5185, *OSPF Multi-Area Adjacency*

RFC 5187, *OSPFv3 Graceful Restart – helper mode*

RFC 5243, *OSPF Database Exchange Summary List Optimization*

RFC 5250, *The OSPF Opaque LSA Option*

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

RFC 5340, *OSPF for IPv6*

RFC 5642, *Dynamic Hostname Exchange Mechanism for OSPF*

RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*

RFC 5838, *Support of Address Families in OSPFv3*

RFC 6549, *OSPFv2 Multi-Instance Extensions*

RFC 6987, *OSPF Stub Router Advertisement*

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

RFC 7684, *OSPFv2 Prefix/Link Attribute Advertisement*

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

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

RFC 8920, *OSPF Application-Specific Link Attributes*

9.25 OpenFlow

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

9.26 Path Computation Element Protocol (PCEP)

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

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

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

draft-ietf-pce-pceps-tls13-04, *Updates for PCEPS: TLS Connection Establishment Restrictions*

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*

9.27 Point-to-Point Protocol (PPP)

RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*

RFC 1661, *The Point-to-Point Protocol (PPP)*

RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*

RFC 1990, *The PPP Multilink Protocol (MP)*

RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*

RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*

RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*

RFC 5072, *IP Version 6 over PPP*

9.28 Policy management and credit control

3GPP TS 29.212 Release 11, *Policy and Charging Control (PCC); Reference points – Gx support as it applies to wireline environment (BNG)*

RFC 4006, *Diameter Credit-Control Application*

RFC 6733, *Diameter Base Protocol*

9.29 Pseudowire (PW)

draft-ietf-l2vpn-vpws-iw-oam-04, *OAM Procedures for VPWS Interworking*

MFA Forum 12.0.0, *Multiservice Interworking - Ethernet over MPLS*

MFA Forum 13.0.0, *Fault Management for Multiservice Interworking v1.0*

MFA Forum 16.0.0, *Multiservice Interworking - IP over MPLS*

RFC 3916, *Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)*

RFC 3985, *Pseudo Wire Emulation Edge-to-Edge (PWE3)*

RFC 4385, *Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN*

RFC 4446, *IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)*

RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*

RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*

RFC 5085, *Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires*

RFC 5659, *An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge*

RFC 5885, *Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)*

RFC 6073, *Segmented Pseudowire*
RFC 6310, *Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping*
RFC 6391, *Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network*
RFC 6575, *Address Resolution Protocol (ARP) Mediation for IP Interworking of Layer 2 VPNs*
RFC 6718, *Pseudowire Redundancy*
RFC 6829, *Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6*
RFC 6870, *Pseudowire Preferential Forwarding Status bit*
RFC 7023, *MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking*
RFC 7267, *Dynamic Placement of Multi-Segment Pseudowires*
RFC 7392, *Explicit Path Routing for Dynamic Multi-Segment Pseudowires – ER-TLV and ER-HOP IPv4 Prefix*
RFC 8395, *Extensions to BGP-Signaled Pseudowires to Support Flow-Aware Transport Labels*

9.30 Quality of Service (QoS)

RFC 2430, *A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)*
RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*
RFC 2597, *Assured Forwarding PHB Group*
RFC 3140, *Per Hop Behavior Identification Codes*
RFC 3246, *An Expedited Forwarding PHB (Per-Hop Behavior)*

9.31 Remote Authentication Dial In User Service (RADIUS)

draft-oscca-cfrg-sm3-02, *The SM3 Cryptographic Hash Function*
RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*
RFC 2866, *RADIUS Accounting*
RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support*
RFC 2868, *RADIUS Attributes for Tunnel Protocol Support*
RFC 2869, *RADIUS Extensions*
RFC 3162, *RADIUS and IPv6*
RFC 4818, *RADIUS Delegated-IPv6-Prefix Attribute*
RFC 5176, *Dynamic Authorization Extensions to RADIUS*
RFC 6613, *RADIUS over TCP – with TLS*
RFC 6614, *Transport Layer Security (TLS) Encryption for RADIUS*
RFC 6929, *Remote Authentication Dial-In User Service (RADIUS) Protocol Extensions*
RFC 6911, *RADIUS attributes for IPv6 Access Networks*

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

draft-newton-mpls-te-dynamic-overbooking-00, A Diffserv-TE Implementation Model to dynamically change booking factors during failure events

RFC 2702, *Requirements for Traffic Engineering over MPLS*

RFC 2747, *RSVP Cryptographic Authentication*

RFC 2961, *RSVP Refresh Overhead Reduction Extensions*

RFC 3097, *RSVP Cryptographic Authentication -- Updated Message Type Value*

RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*

RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*

RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*

RFC 3906, *Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels*

RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

RFC 4124, *Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering*

RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*

RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)*

RFC 5712, *MPLS Traffic Engineering Soft Preemption*

RFC 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks*

9.33 Routing Information Protocol (RIP)

RFC 1058, *Routing Information Protocol*

RFC 2080, *RIPng for IPv6*

RFC 2082, *RIP-2 MD5 Authentication*

RFC 2453, *RIP Version 2*

9.34 Segment Routing (SR)

draft-ietf-bess-mvpn-evpn-sr-p2mp-07, Multicast and Ethernet VPN with Segment Routing P2MP and Ingress Replication – MVPN

draft-bashandy-rtgwg-segment-routing-uloop-15, Loop avoidance using Segment Routing

draft-filsfils-spring-net-pgm-extension-srv6-usid-15, Network Programming extension: SRv6 uSID instruction

draft-filsfils-spring-srv6-net-pgm-insertion-08, SRv6 NET-PGM extension: Insertion

draft-ietf-idr-bgpls-srv6-ext-14, BGP Link State Extensions for SRv6

draft-ietf-idr-segment-routing-te-policy-23, Advertising Segment Routing Policies in BGP

draft-ietf-idr-ts-flowspec-srv6-policy-03, Traffic Steering using BGP FlowSpec with SR Policy

draft-ietf-pim-p2mp-policy-ping-03, P2MP Policy Ping

draft-ietf-pim-sr-p2mp-policy-06, Segment Routing Point-to-Multipoint Policy – MPLS

draft-ietf-rtgwg-segment-routing-ti-lfa-11, Topology Independent Fast Reroute using Segment Routing

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

draft-ietf-spring-sr-replication-segment-16, SR Replication segment for Multi-point Service Delivery – MPLS

draft-ietf-spring-srv6-srh-compression-18, Compressed SRv6 Segment List Encoding

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

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

RFC 8665, OSPF Extensions for Segment Routing

RFC 8666, OSPFv3 Extensions for Segment Routing

RFC 8667, IS-IS Extensions for Segment Routing

RFC 8669, Segment Routing Prefix Segment Identifier Extensions for BGP

RFC 8754, IPv6 Segment Routing Header (SRH)

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

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

RFC 9085, Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing

RFC 9088, Signaling Entropy Label Capability and Entropy Readable Label Depth Using IS-IS – advertising ELC

RFC 9089, Signaling Entropy Label Capability and Entropy Readable Label Depth Using OSPF – advertising ELC

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

RFC 9256, Segment Routing Policy Architecture

RFC 9259, Operations, Administration, and Maintenance (OAM) in Segment Routing over IPv6 (SRv6)

RFC 9350, IGP Flexible Algorithm

RFC 9352, IS-IS Extensions to Support Segment Routing over the IPv6 Data Plane

9.35 Simple Network Management Protocol (SNMP)

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

draft-ietf-isis-wg-mib-06, *Management Information Base for Intermediate System to Intermediate System (IS-IS)*

draft-ietf-mboned-msdp-mib-01, *Multicast Source Discovery protocol MIB*

draft-ietf-mpls-ldp-mib-07, *Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP)*

draft-ietf-mpls-lsr-mib-06, *Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2*

draft-ietf-mpls-te-mib-04, *Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base*

draft-ietf-ospf-mib-update-08, *OSPF Version 2 Management Information Base*

draft-ietf-rrrp-unified-mib-06, *Definitions of Managed Objects for the VRRP over IPv4 and IPv6 – IPv6*

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

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

IANAifType-MIB revision 200505270000Z, *ianaifType*

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

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

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

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

LLDP-MIB revision 200505060000Z, *lldpMIB*

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

RFC 1212, *Concise MIB Definitions*

RFC 1215, *A Convention for Defining Traps for use with the SNMP*

RFC 1724, *RIP Version 2 MIB Extension*

RFC 1901, *Introduction to Community-based SNMPv2*

RFC 2021, *Remote Network Monitoring Management Information Base Version 2 using SMIv2*

RFC 2206, *RSVP Management Information Base using SMIv2*

RFC 2213, *Integrated Services Management Information Base using SMIv2*

RFC 2494, *Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type*

RFC 2578, *Structure of Management Information Version 2 (SMIv2)*

RFC 2579, *Textual Conventions for SMIv2*

RFC 2580, *Conformance Statements for SMIv2*

RFC 2787, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol*

RFC 2819, *Remote Network Monitoring Management Information Base*

RFC 2856, *Textual Conventions for Additional High Capacity Data Types*

RFC 2863, *The Interfaces Group MIB*

RFC 2864, *The Inverted Stack Table Extension to the Interfaces Group MIB*

RFC 2933, *Internet Group Management Protocol MIB*

RFC 3014, *Notification Log MIB*

RFC 3165, *Definitions of Managed Objects for the Delegation of Management Scripts*

RFC 3231, *Definitions of Managed Objects for Scheduling Management Operations*

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

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

RFC 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 3498, *Definitions of Managed Objects for Synchronous Optical Network (SONET) Linear Automatic Protection Switching (APS) Architectures*

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

RFC 3592, *Definitions of Managed Objects for the Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) Interface Type*

RFC 3593, *Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals*

RFC 3635, *Definitions of Managed Objects for the Ethernet-like Interface Types*

RFC 3637, *Definitions of Managed Objects for the Ethernet WAN Interface Sublayer*

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

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

RFC 3895, *Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types*

RFC 3896, *Definitions of Managed Objects for the DS3/E3 Interface Type*

RFC 4001, *Textual Conventions for Internet Network Addresses*

RFC 4022, *Management Information Base for the Transmission Control Protocol (TCP)*
RFC 4113, *Management Information Base for the User Datagram Protocol (UDP)*
RFC 4220, *Traffic Engineering Link Management Information Base*
RFC 4273, *Definitions of Managed Objects for BGP-4*
RFC 4292, *IP Forwarding Table MIB*
RFC 4293, *Management Information Base for the Internet Protocol (IP)*
RFC 4631, *Link Management Protocol (LMP) Management Information Base (MIB)*
RFC 4878, *Definitions and Managed Objects for Operations, Administration, and Maintenance (OAM) Functions on Ethernet-Like Interfaces*
RFC 7420, *Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module*
RFC 7630, *HMAC-SHA-2 Authentication Protocols in the User-based Security Model (USM) for SNMPv3*
SFLOW-MIB revision 200309240000Z, *sFlowMIB*

9.36 Timing

GR-1244-CORE Issue 3, *Clocks for the Synchronized Network: Common Generic Criteria*
GR-253-CORE Issue 3, *SONET Transport Systems: Common Generic Criteria*
IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*
ITU-T G.781, *Synchronization layer functions*
ITU-T G.811, *Timing characteristics of primary reference clocks*
ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC)*
ITU-T G.8261, *Timing and synchronization aspects in packet networks*
ITU-T G.8262, *Timing characteristics of synchronous Ethernet equipment slave clock (EEC)*
ITU-T G.8262.1, *Timing characteristics of an enhanced synchronous Ethernet equipment slave clock (eEEC)*
ITU-T G.8264, *Distribution of timing information through packet networks*
ITU-T G.8265.1, *Precision time protocol telecom profile for frequency synchronization*
ITU-T G.8272, *Timing characteristics of primary reference time clocks – PRTC-A, PRTC-B*
ITU-T G.8275.1, *Precision time protocol telecom profile for phase/time synchronization with full timing support from the network*
ITU-T G.8275.2, *Precision time protocol telecom profile for phase/time synchronization with partial timing support from the network*
RFC 3339, *Date and Time on the Internet: Timestamps*
RFC 5905, *Network Time Protocol Version 4: Protocol and Algorithms Specification*
RFC 8573, *Message Authentication Code for the Network Time Protocol*

9.37 Two-Way Active Measurement Protocol (TWAMP)

- RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP) – server, unauthenticated mode*
- RFC 5938, *Individual Session Control Feature for the Two-Way Active Measurement Protocol (TWAMP)*
- RFC 6038, *Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features*
- RFC 8545, *Well-Known Port Assignments for the One-Way Active Measurement Protocol (OWAMP) and the Two-Way Active Measurement Protocol (TWAMP) – TWAMP*
- RFC 8762, *Simple Two-Way Active Measurement Protocol – unauthenticated*
- RFC 8972, *Simple Two-Way Active Measurement Protocol Optional Extensions – unauthenticated*
- RFC 9503, *Simple Two-Way Active Measurement Protocol (STAMP) Extensions for Segment Routing Networks – excluding Sections 3, 4.1.2 and 4.1.3*

9.38 Virtual Private LAN Service (VPLS)

- RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*
- RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*
- RFC 5501, *Requirements for Multicast Support in Virtual Private LAN Services*
- RFC 6074, *Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)*
- RFC 7041, *Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging*
- RFC 7117, *Multicast in Virtual Private LAN Service (VPLS)*

9.39 Voice and video

- DVB BlueBook A86, *Transport of MPEG-2 TS Based DVB Services over IP Based Networks*
- ETSI TS 101 329-5 Annex E, *QoS Measurement for VoIP - Method for determining an Equipment Impairment Factor using Passive Monitoring*
- ITU-T G.1020 Appendix I, *Performance Parameter Definitions for Quality of Speech and other Voiceband Applications Utilizing IP Networks - Mean Absolute Packet Delay Variation & Markov Models*
- ITU-T G.107, *The E Model - A computational model for use in planning*
- ITU-T P.564, *Conformance testing for voice over IP transmission quality assessment models*
- RFC 3550, *RTP: A Transport Protocol for Real-Time Applications – Appendix A.8*
- RFC 4585, *Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)*
- RFC 4588, *RTP Retransmission Payload Format*

9.40 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*

9.41 Yet Another Next Generation (YANG) OpenConfig Models

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

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

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

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

openconfig-alarms.yang version 0.3.2, *OpenConfig System Alarms Model*

openconfig-bfd.yang version 0.2.2, *OpenConfig BFD Model*

openconfig-bgp.yang version 6.1.0, *OpenConfig BGP Model*

openconfig-bgp-common.yang version 6.0.0, *OpenConfig BGP Common Model*

openconfig-bgp-common-multiprotocol.yang version 6.0.0, *OpenConfig BGP Common Multiprotocol Model*

openconfig-bgp-common-structure.yang version 6.0.0, *OpenConfig BGP Common Structure Model*

openconfig-bgp-global.yang version 6.0.0, *OpenConfig BGP Global Model*

openconfig-bgp-neighbor.yang version 6.1.0, *OpenConfig BGP Neighbor Model*

openconfig-bgp-peer-group.yang version 6.1.0, *OpenConfig BGP Peer Group Model*

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

openconfig-if-aggregate.yang version 2.4.3, *OpenConfig Interfaces Aggregated Model*

openconfig-if-ethernet.yang version 2.12.2, *OpenConfig Interfaces Ethernet Model*

openconfig-if-ip.yang version 3.1.0, *OpenConfig Interfaces IP Model*

openconfig-if-ip-ext.yang version 2.3.1, *OpenConfig Interfaces IP Extensions Model*

openconfig-igmp.yang version 0.3.1, *OpenConfig IGMP Model*

openconfig-interfaces.yang version 3.0.0, *OpenConfig Interfaces Model*

openconfig-isis.yang version 1.1.0, *OpenConfig IS-IS Model*

openconfig-isis-policy.yang version 0.5.0, *OpenConfig IS-IS Policy Model*

openconfig-isis-routing.yang version 1.1.0, *OpenConfig IS-IS Routing Model*

openconfig-lacp.yang version 1.3.0, *OpenConfig LACP Model*

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


openconfig-local-routing.yang version 1.2.0, *OpenConfig Local Routing Model*

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

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

openconfig-mpls-rsvp.yang version 2.3.0, *OpenConfig MPLS RSVP Model*
openconfig-mpls-te.yang version 2.3.0, *OpenConfig MPLS TE Model*
openconfig-network-instance.yang version 1.1.0, *OpenConfig Network Instance Model*
openconfig-network-instance-l3.yang version 0.11.1, *OpenConfig L3 Network Instance Model – static routes*
openconfig-ospfv2.yang version 0.4.0, *OpenConfig OSPFv2 Model*
openconfig-ospfv2-area.yang version 0.4.0, *OpenConfig OSPFv2 Area Model*
openconfig-ospfv2-area-interface.yang version 0.4.0, *OpenConfig OSPFv2 Area Interface Model*
openconfig-ospfv2-common.yang version 0.4.0, *OpenConfig OSPFv2 Common Model*
openconfig-ospfv2-global.yang version 0.4.0, *OpenConfig OSPFv2 Global Model*
openconfig-packet-match.yang version 1.0.0, *OpenConfig Packet Match Model*
openconfig-pim.yang version 0.4.3, *OpenConfig PIM Model*
openconfig-platform.yang version 0.15.0, *OpenConfig Platform Model*
openconfig-platform-fan.yang version 0.1.1, *OpenConfig Platform Fan Model*
openconfig-platform-linecard.yang version 0.1.2, *OpenConfig Platform Linecard Model*
openconfig-platform-port.yang version 0.4.2, *OpenConfig Port Model*
openconfig-platform-transceiver.yang version 0.9.0, *OpenConfig Transceiver Model*
openconfig-procmon.yang version 0.4.0, *OpenConfig Process Monitoring Model*
openconfig-relay-agent.yang version 0.1.0, *OpenConfig Relay Agent Model*
openconfig-routing-policy.yang version 3.0.0, *OpenConfig Routing Policy Model*
openconfig-rsvp-sr-ext.yang version 0.1.0, *OpenConfig RSVP-TE and SR Extensions Model*
openconfig-system.yang version 0.10.1, *OpenConfig System Model*
openconfig-system-grpc.yang version 1.0.0, *OpenConfig System gRPC Model*
openconfig-system-logging.yang version 0.3.1, *OpenConfig System Logging Model*
openconfig-system-terminal.yang version 0.3.0, *OpenConfig System Terminal Model*
openconfig-telemetry.yang version 0.5.0, *OpenConfig Telemetry Model*
openconfig-terminal-device.yang version 1.9.0, *OpenConfig Terminal Optics Device Model*
openconfig-vlan.yang version 2.0.0, *OpenConfig VLAN Model*

Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)