



# Multi-Access Gateway – controller

Release 23.3.R1

## Control Plane Function Guide

---

3HE 19258 AAAA TQZZA  
Edition 01  
March 2023

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.

© 2023 Nokia.

# Table of contents

<b>List of tables.....</b>	<b>7</b>
<b>List of figures.....</b>	<b>8</b>
<b>1 Getting started.....</b>	<b>10</b>
1.1 About this guide.....	10
1.2 Conventions.....	10
1.2.1 Precautionary and information messages.....	10
1.2.2 Options or substeps in procedures and sequential workflows.....	11
<b>2 Session management.....</b>	<b>12</b>
2.1 PFCP session.....	12
2.2 General session functionality.....	12
2.2.1 PFCP protocol.....	12
2.2.1.1 PFCP association and path.....	12
2.2.1.2 BNG-UP selection.....	13
2.2.1.3 PFCP session state.....	13
2.2.1.4 PFCP provisioning components.....	13
2.2.2 PFCP connectivity failure.....	15
2.2.2.1 Headless mode.....	15
2.2.2.2 Session timer alignment.....	16
2.2.3 Subscribers.....	17
2.2.4 QoS.....	17
2.2.5 Service selection.....	18
2.2.6 Operational commands and debugging.....	18
2.2.7 CP session state and VM resilience.....	21
2.2.8 Prefix delegation as a framed route.....	22
2.2.9 Lawful intercept.....	22
2.3 Fixed access sessions.....	24
2.3.1 Layer 2 circuit.....	24
2.3.2 In-band control plane and BNG-UP selection.....	24
2.3.3 Session keys and anti-spoofing.....	26
2.3.4 Subscriber identification.....	26
2.3.5 Session limits.....	27

2.3.6	Session lockout.....	28
2.3.6.1	Enabling session lockout.....	28
2.3.7	IPoE.....	29
2.3.8	PPPoE.....	30
2.4	Address assignment protocols.....	39
2.4.1	DHCP.....	39
2.4.2	IPCP.....	40
2.4.3	ICMPv6 Router Advertisements (RA) and SLAAC.....	40
2.4.4	DHCPv6.....	42
2.5	Fixed Wireless Access sessions.....	43
2.5.1	Introduction to FWA.....	43
2.5.2	Residential Gateway models.....	44
2.5.3	Selected and real APN.....	45
2.5.4	Session identification, subscriber identification, and multi-APN support.....	46
2.5.5	FWA-UP selection.....	47
2.5.6	Address signaling methods and deferred allocation.....	47
2.5.7	QoS.....	48
2.5.8	PAP/CHAP authentication.....	49
2.5.9	Session lifetime and held addresses.....	49
2.5.10	Headless mode for FWA.....	50
2.5.11	4G and 5G NSA option 3 sessions.....	51
2.6	YANG state.....	55
<b>3</b>	<b>Address assignment.....</b>	<b>58</b>
3.1	Overview of address assignment.....	58
3.2	ODSA and local address assignment.....	58
3.2.1	ODSA.....	58
3.2.2	Variable prefix length and micro-net length.....	60
3.2.3	Local address assignment.....	61
3.3	AAA-based address assignment.....	62
3.4	Non-provisioned address assignment.....	63
3.5	Local static address assignment via authentication database.....	64
<b>4</b>	<b>Authentication.....</b>	<b>65</b>
4.1	Overview of the authentication process.....	65
4.2	BNG entry point (EP).....	65

4.3	Authentication database (ADB).....	66
4.4	Authentication flow.....	67
4.5	BNG EP and ADB lookup.....	68
4.6	Required minimal configuration for a session creation.....	71
4.7	RADIUS authentication profile.....	72
4.8	RADIUS CoA and DM.....	73
4.9	Example configuration.....	74
4.10	Web portal authentication.....	76
4.10.1	Configuring WPP.....	79
<b>5</b>	<b>Accounting and charging.....</b>	<b>81</b>
5.1	BNG charging profiles.....	81
5.2	Statistics collection from the BNG-UP.....	81
5.3	MAG-c-based charging.....	84
5.4	RADIUS accounting.....	84
5.4.1	Enabling RADIUS accounting.....	85
5.4.2	Session accounting.....	85
5.4.3	Message retransmission and buffering.....	87
<b>6</b>	<b>Residential NAT.....</b>	<b>89</b>
6.1	NAT terminology and references.....	89
6.2	Residential NAT44 on BNG CUPS.....	90
6.3	Functional split between MAG-c and BNG-UP.....	91
6.4	Management of NAT outside prefixes.....	91
6.5	CP NAT profile.....	92
6.6	Port forwards.....	92
6.7	Extended port blocks.....	93
6.7.1	Port space division.....	93
6.7.2	Managing port block space.....	94
6.8	NAT logging.....	96
6.8.1	RADIUS-based logging.....	97
6.8.1.1	Enabling RADIUS logging on MAG-c.....	101
6.8.1.2	Timestamp interpretation.....	102
6.8.1.3	High logging rates.....	103
6.8.1.4	Buffering during RADIUS failure.....	103
6.9	Watermarks.....	104

6.10	Minimum configuration steps.....	104
6.11	Operational commands.....	105
<b>7</b>	<b>Geo-redundancy.....</b>	<b>106</b>
7.1	Geo-redundancy overview.....	106
7.2	Operational and administrative roles.....	106
7.3	Traffic detection.....	107
7.4	State synchronization.....	108
7.5	Routing.....	108
7.6	Shunting.....	109
7.7	Manual switchover.....	110
7.8	Deploying and configuring geo-redundancy.....	110
<b>8</b>	<b>Python support.....</b>	<b>113</b>
8.1	Configuring a Python script.....	115
8.2	Protecting a Python script file.....	116
<b>9</b>	<b>BNG-UP resiliency.....</b>	<b>117</b>
9.1	Terminology for BNG-UP resiliency.....	117
9.2	Introduction to MAG-c-driven BNG-UP resiliency.....	117
9.3	Modeling a resilient BNG-UP deployment using UP groups.....	119
9.3.1	Fate sharing group creation.....	119
9.3.2	Fixed access.....	120
9.4	Fate sharing groups.....	125
9.4.1	Session-to-FSG mapping.....	125
9.4.2	Traffic steering parameters.....	125
9.4.3	BNG-UP health determination.....	127
9.4.4	Active/standby selection triggers.....	130
9.4.5	Active/standby selection.....	131
9.4.6	Active/standby change or switchover.....	133
9.4.7	UP Lockout.....	135
9.5	Warm and hot standby.....	136
9.6	Interaction with headless mode.....	136
9.7	Operational commands.....	137

---

## List of tables

Table 1: Address assignment protocols per session type.....	39
Table 2: Use cases and context for the apn-format command.....	46
Table 3: Minimal configuration for a session creation.....	72
Table 4: RADIUS based logging for residential NAT.....	96
Table 5: Integrated subscriber management and NAT RADIUS accounting.....	97
Table 6: Supported direction for RADIUS messages.....	114
Table 7: Supported direction for RADIUS CoA messages.....	114
Table 8: Supported direction for PPPoE messages.....	114
Table 9: Summary of BNG-UP states.....	129

# List of figures

Figure 1: HQoS example.....	17
Figure 2: IPoE session setup with RADIUS authentication.....	29
Figure 3: PPPoE session setup flow.....	31
Figure 4: Resiliency based on PADO delay.....	35
Figure 5: L2TP LAC network components.....	36
Figure 6: LAC-enabled PPPoE session setup flow.....	37
Figure 7: A separate model with PPP connectivity.....	45
Figure 8: Example of a bearer or QoS flow rate enforcement.....	49
Figure 9: Basic FWA network.....	51
Figure 10: 4G FWA session setup.....	52
Figure 11: WPP on BNG CUPS call flow example.....	77
Figure 12: Statistics collection using the pull model.....	82
Figure 13: Statistics collection using the push model.....	82
Figure 14: Residential NAT example.....	90
Figure 15: Alc-ISA-Event-Timestamp triggered Interim-Update message.....	103
Figure 16: Geo-redundant MAG-c deployment.....	106
Figure 17: High-level overview of communication for BNG-UP resiliency.....	118
Figure 18: Multiple backup BNG-UPs.....	118
Figure 19: Multiple Layer 2 access IDs per UP group.....	120
Figure 20: 1:1 hot standby resiliency example.....	122
Figure 21: Per S-tag 1:1 hot standby resiliency example.....	124



---

Figure 22: N:1 warm standby resiliency.....	
Figure 23: Example of the relationship between FSGs, MAC addresses, and subnets.....	127
Figure 24: GARP race conditions.....	135

# 1 Getting started

Find general information about this guide.

## 1.1 About this guide

This guide describes the Nokia Multi-Access Gateway – controller (MAG-c) for the BNG CUPS solution. The MAG-c is based on the Packet Forwarding Control Protocol (PFCP) interface as defined in 3GPP TS 29.244 for mobile 4G CUPS and 5G.

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

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

The CLI trees and command descriptions can be found in the *MAG-c CLI Reference Guide*.



**Note:** This guide covers content for the release specified on the title page of the guide, and may also contain content that will be released in later maintenance loads. See the applicable *MAG-c Release Notes* for information about features supported in each load of the release software.



**Note:** The information in this guide is intended to be used in conjunction with the 7750 SR software user guides. The 7750 SR software user guides describe SR OS service features that are supported by the MAG-c. For specific guide titles, see the *7450 ESS, 7750 SR, 7950 XRS, and VSR Documentation Suite Overview Card 20.10.R1*.

## 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 Session management

*Get a general overview of the session functionality and the address assignment protocols, and details on the IPoE and PPPoE fixed access and the fixed wireless access (FWA) session types.*

### 2.1 PFCP session

A session is the basic operational object of the BNG and represents the connectivity of a single device such as a residential gateway. Address assignment, authentication, accounting, and BNG-UP communication are all done in the scope of a single session. A PFCP association is needed to create a PFCP session.

### 2.2 General session functionality

*Get a high level overview of the PFCP protocol and general session related functionality including grouping sessions for a subscriber, QoS, service selection, session state, and lawful intercept.*

#### 2.2.1 PFCP protocol

*Get a high level overview of the core protocol of session management.*

The core of session management is the PFCP protocol as defined in 3GPP TS 29.244, with BNG-specific extensions defined in BBF TR-459.

##### 2.2.1.1 PFCP association and path

*The PFCP association and path define the connectivity between the MAG-c and BNG-UP.*

To send a session to a BNG-UP, a PFCP association needs to be established. While establishing this association, the BNG-UP and the MAG-c exchange capabilities, functional features, and parameters; for example, a BNG-UP sends functional features such as PPPoE support, IPoE support, and LCP Keep-alive Offload support. Capability exchange can influence the IE applicability in PFCP session messages.

- **PFCP association**

A PFCP association must be set up before sessions can be established between the BNG-UP and the MAG-c. Only one association per MAG-c and BNG-UP pair is allowed. The identifiers of the association are the MAG-c and the BNG-UP node IDs, which can be IP addresses or domain names. Provisioning commands specific to PFCP associations allow to enable the PFCP protocol.

- **PFCP path**

Multiple paths are possible per PFCP association. The identifier of a PFCP path is the pair of IP addresses used to communicate between the MAG-c and the BNG-UP. Paths are not negotiated but are learned while using PFCP signaling. Each IP address is called a PFCP entity. Each pair of MAG-c and BNG-UP IP addresses is called a PFCP path.

**Note:**

- The Nokia MAG-c uses only one IP address per association, although in general a MAG-c or BNG-UP (also called a PFCP node) could use multiple IP addresses for communication within the same PFCP association. Because the Nokia MAG-c and BNG-UP use only one PFCP path per association, the terms path and association are often used interchangeably.
- Both the MAG-c and BNG-UP verify that all known paths are alive using PFCP heartbeat messages. The heartbeat parameters are configured in the context of the PFCP profile. When a path fails, all related sessions are removed.

### 2.2.1.2 BNG-UP selection

*The MAG-c selects a BNG-UP for each session depending on the session type.*

The BNG-UP selection varies from very static (for example, because of hard wiring) to very dynamic (for example, for FWA).

See [In-band control plane and BNG-UP selection](#) for more information about the selection process for fixed access session.

See [FWA-UP selection](#) for more information about the selection process for FWA sessions.

### 2.2.1.3 PFCP session state

*PFCP sessions require the creation of a forwarding state on the BNG-UP device. Session operations allow to manage the forwarding state on the BNG-UP.*

The forwarding state includes rules (for example, encapsulation and decapsulation), information about routing context forwarding, QoS rules, and requested statistics collection.

The PFCP session establishment procedure creates the initial forwarding state. The path used for the PFCP session establishment procedure is tied to the session.

The following operations are supported for an established session:

- **PFCP session modification**  
The MAG-c modifies the state or performs a state query (for example, to fetch statistics).
- **PFCP session deletion**  
The MAG-c removes all state information.
- **PFCP session report**  
The BNG-UP sends information unsolicited (for example, to report statistics or a connectivity failure).

In stable conditions, the BNG-UP only modifies or deletes the state if instructed by the MAG-c. If the BNG-UP detects failure, for example, a link failure, it does not delete the state but sends a report and keeps the local state. The BNG-UP deletes the state only when the MAG-c sends a Delete Request.

### 2.2.1.4 PFCP provisioning components

*The PFCP protocol uses components that must be provisioned using the CLI.*

To enable the PFCP protocol, the following components must be provisioned and referred to in the PDN configuration:

- **pfc-association-peer-list**

The **pcfp-association-peer-list** command in the **config>mobile>profile>pcfp** context provisions the list of peer BNG-UP devices. It can be left empty if the BNG-UP initiates the PFCP association. By default, PFCP association requests coming from a peer that is not configured in the PFCP association list are accepted and the requesting (BNG-UP) peers are added dynamically to the PFCP association list for the MAG-c.



**Note:**

If the peer association list is enforced with the **enforced-pfcfp-association-list** command in the **config>mobile>pdn>sx-n4** context, all BNG-UP devices must already be provisioned in the PFCP peer association list. If enforcement is enabled, PFCP association requests coming from a peer that is not configured in the PFCP association list are not accepted and the requesting peer is not dynamically added to the list.

- **up-peer-list**

The **up-peer-list** command in the **config>mobile>profile>pcfp** context creates a list of UPs with the supported APNs and the UE IP address pools. In a MAG-c deployment, it must be configured for FWA sessions. It can be empty for the other types of session.

- **pcfp-profile**

The **pcfp-profile** command in the **config>mobile>profile>pcfp** context creates a PFCP profile in which the PFCP parameters can be configured.



**Note:**

The heartbeat and the retransmit options must be configured with the same values in both the BNG-UP and MAG-c, to prevent the BNG-UP and MAG-c from going out of sync if a link failure occurs.

The following example shows the provisioning in the **config>mobile>profile>pcfp** context and the **config>mobile>pdn** context.

```
A:MAG-c>config>mobile>profile>pcfp# info
-----
      pfcfp-association-peer-list "peers"
      exit
      up-peer-list "ups"
      exit
-----
A:MAG-c>config>mobile>pdn# info
-----
      instance-type control
      sx-n4 "default"
      pfcfp-association-list "peers"
      interface
      pfcfp "system"
      ibcp "system"
      exit
      signaling
      pfcfp
      profile "default"
      exit
      ibcp
      bng-entry-point "start"
      triggers pppoe-discover ipoe-dhcp ipoe-dhcpv6 ipoe-router-solicit
      exit
      exit
      exit
      up-peer-list "ups"
```

---

## Related topics

[Headless mode](#)

## 2.2.2 PFCP connectivity failure

*To protect against temporary PFCP connectivity failures, MAG-c supports a headless mode. Following the rules for the configuration of the sessions timers makes sure that the headless mode works as expected.*

### 2.2.2.1 Headless mode

*To prevent the removal of sessions with a temporary heartbeat failure, MAG-c supports a short-lived headless mode to restore connectivity.*

PFCP heartbeat messages check the connectivity of a PFCP path. When the heartbeat procedure fails, all state information for the corresponding path is removed and all sessions using that path are terminated. The association remains in place.

Use the following command to configure the heartbeat parameters:

```
configure mobile-gateway profile pfcpc pfcpc-profile heart-beat
```

To protect against temporary failures, the MAG-c and BNG-UP support a headless mode. Use the following command to enable the headless mode:

```
configure mobile-gateway profile pfcpc pfcpc-profile path-restoration-time
```

To enable BFD, use the optional **bfd-enable** keyword in the preceding command.

When BFD is enabled, the system starts a BFD session for each known PFCP path on the BNG-UP. If a BFD failure is detected, the system immediately brings down the associated path. BFD does not affect the path recovery detection, which requires the configuration of PFCP heartbeats.

When headless mode is enabled, the sessions are not removed when there is a heartbeat failure. Instead, the configured timer is started and heartbeats continue to be sent. Subsequently, one of the following events occurs:

- The timer expires and all sessions are removed. The association remains in place.
- The path is restored (a successful heartbeat is completed) but a BNG-UP restart is detected and all sessions are removed.
- The path is restored (a successful heartbeat is completed), the sessions are kept, and a PFCP audit procedure is started to ensure that the BNG-UP and MAG-c states are synchronized.



#### Note:

- To prevent the MAG-c or BNG-UP from deleting all sessions while the other node keeps all the sessions, Nokia recommends that the path restoration time is at least twice as large as the sum of the **heart-beat interval** plus the total heartbeat timeout (total heartbeat timeout = **heart-beat retry-count** N1 × **heart-beat timeout** T1). This ensures that the MAG-c and BNG-UP nodes each run an audit or delete all the sessions in their respective nodes.
- All parameter configurations must be identical between the MAG-c and BNG-UP.

To avoid hanging resources on a BNG-UP, the MAG-c only removes a session after it receives confirmation that the BNG-UP has removed the session. The MAG-c may receive confirmation in the following messages:

- PFCP Session Deletion Response message (most common case)
- PFCP message including a Cause IE that indicates an error (the BNG-UP lost the session)
- an indication that the BNG-UP restarted and lost all its sessions; for example, a new PFCP Association Setup Request

To remove a session manually, use the **force** keyword with the following command:

```
clear mobile-gateway bng session
```

This removes operational BNG (non-FWA) sessions on the MAG-c without synchronization with any external server or the client.

#### Related topics

[Operational commands and debugging](#)

### 2.2.2.2 Session timer alignment

Nokia recommends aligning the session timers (signaled to the BNG RG) with the path restoration time. If the session timers are not aligned with the path restoration time, a session may time out autonomously before the headless mode could restore the path.

The following configurations for the session timers guarantee that the headless mode kicks in as expected.

- For DHCP, the DHCP lease time must at least equal the renew time plus the path restoration time. In the default case, where the renew time is half of the lease time, the lease time must be at least twice the path restoration time.
- For all IPv6 enabled sessions, the router lifetime included in RA messages must be at least equal to the maximum advertisement interval plus the path restoration time. In the default case, where the router lifetime is three times the maximum advertisement interval, the maximum advertisement interval must be equal to at least twice the path restoration time.
- For SLAAC, the IPv6 preferred lifetime must be at least equal to the maximum router advertisement interval plus the path restoration time.
- For DHCPv6, the IPv6 preferred lifetime must be at least equal to the renew timer (T1 timer) plus the path restoration time. In the default case, where the renew timer is half of the preferred lifetime, the preferred lifetime must be equal to at least twice the path restoration time.

The parameters can be locally configured or received from an external AAA server.

To configure or get info about the following locally configured parameters, use the indicated CLI commands:

- path restoration time: **path-restoration-time** in the **config>mobile>profile>pfcp>pfcp-profile** context
- DHCP lease time: **option** with *option-number* 51 in the **config>mobile>profile>bng>dhcp-profile>options** context
- Renew time: **option** with *option-number* 58 in the **config>mobile>profile>bng>dhcp-profile>options** context



- Router lifetime in RA messages: **options router-lifetime** in the **config>mobile>profile>bng>ra-profile** context
- Maximum advertisement interval: **advertisement-interval max** in the **config>mobile>profile>bng>ra-profile** context
- IPv6 preferred lifetime: **preferred** in the **config>mobile>profile>adb>entry>address-assignment>lifetimes** context

## 2.2.3 Subscribers

*The MAG-c supports bundling a group of sessions for a single subscriber.*

Grouping of sessions is useful in cases where a subscription consists of multiple directly connected devices. For example, a subscription may consist of a routed residential gateway for Internet access, VoIP phones, and set-top boxes. The residential gateway bridges traffic for voice and video services to the VoIP phones and to the set-top boxes. The MAG-c automatically creates a subscriber based on keys it derives from the session types, and allocates an auto-generated subscriber ID to the sessions.

See [Subscriber identification](#) for fixed access sessions and [Session identification, subscriber identification, and multi-APN support](#) for FWA sessions for more information about how the subscriber ID is generated.

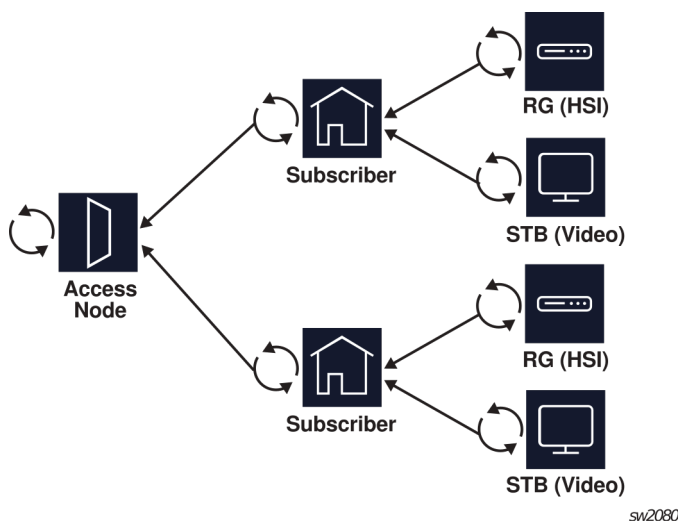
A subscriber ID alias can be provided via AAA interfaces, but this alias cannot change the scope of a subscriber. For example, if the key of a subscriber contains a Layer 2 circuit (I2-circuit), the AAA subscriber ID alias cannot group two sessions with two different I2-circuit values.

## 2.2.4 QoS

*The MAG-c enables the appropriate HQoS configuration by sending subscriber profiles and SLA profiles to the BNG-UP.*

A BNG connection uses HQoS structures, in which there are multiple levels of rate limiting and scheduling. For example, one structure has an aggregate rate per MSAN, a second structure has an aggregate rate per subscriber level, and a third structure has an aggregate rate per session.

*Figure 1: HQoS example*



HQoS models can be complex and very hardware specific, the MAG-c signals the BNG-UP profiles to enable the appropriate HQoS configuration. The Nokia MAG-c signals a subscriber profile and an SLA profile in the Activate Predefined Rules IE of the PFCP message. The profiles are provisioned during authentication.

The subscriber profile should be kept consistent for all sessions of a subscriber, but the MAG-c does not enforce consistency. Short-lived inconsistencies are allowed while changing a subscriber profile; for example, when sending a CoA message to all sessions of the subscribers. However, long-lived inconsistencies may lead to unexpected behavior, including reverting to an old subscriber profile.

## 2.2.5 Service selection

*The service selection requires an APN with a provisioned network realm.*

Because the MAG-c model is based on 3GPP, the service selection relies on APNs. An APN must be provided during the session authentication to avoid session setup failure.

To map the APN to a specific routing service on the BNG-UP, a network realm must be provisioned in the APN context. This realm is sent over the PFCP interface and maps to a Layer 3 service on the BNG-UP. APNs using different pools can map to the same realm.

The following example shows an APN configuration.

```
A:MAG_C>config>mobile>pdn>apn# info
-----
                network-realm "hsi"
                no shutdown
-----
```



### Note:

Most of the APN configurations do not apply to a MAG-c.

## 2.2.6 Operational commands and debugging

*Commands in the **show**, **clear**, and **tools** context allow to display MAG-c sessions and subscribers, to remove a session, and to debug a failing session setup. The call trace feature can be used for advanced debugging.*

To display the MAG-c sessions and subscribers, use the following commands in the **show** context:

- Use the **session** command in the **show>mobile>bng** context to display an overview of all basic data related to a session.
  - Add filters supported for the **session** command to display the data of specific sets of sessions.
  - Add the **count** keyword to display only the number of matching sessions.
  - Use sub-commands to get details on a specific aspect of the session (for example, **charging**).
- Use the **subscriber** command in the **show>mobile>bng** context to display an overview of the operational data related to a subscriber. The command has similar options as the **session** command.
- Use the **ibcp** keyword with the **ref-point-stats** command in the **show>mobile>pdn** context to display IBCP statistics for fixed access sessions.

To remove a session from the MAG-c, use the **session** command in the **clear>mobile>bng** context. When you issue this command, the MAG-c sends a PFCP Session Deletion Request to the BNG-UP.

To delete an operational BNG (non-FWA) session on the MAG-c without synchronization with any external server or the client session, use the **force** keyword with the **session** command. This bypasses the headless mode mechanism.



**Tip:** Always add filters to the **session** command in the **clear>mobile>bng** context to avoid accidental clearing of all sessions.

To debug a failing setup session, use the **error-history** command in the **tools>dump>mobile>bng** context to display a basic error log.

## Call trace

For more advanced debugging, the MAG-c supports the call trace feature which allows to trace control plane packets or events during the lifetime of a session. Such a trace allows to correlate multiple signaling interfaces and detect issues in a session setup flow. Call trace supports the most frequently used BNG and FWA protocols such as GTP, SBI, DHCP, DHCPv6, RADIUS, and Diameter. Important non-protocol related events are also logged, such as ADB lookups or running a Python script.

To enable call trace debugging of fixed access sessions, use the following command:

```
debug mobile-gateway call-insight bng
```

The parameters of the command, including MAC address, layer 2 access ID, layer 2 circuit ID, and remote ID, define the sessions to trace.

To enable call trace debugging of FWA sessions, use the following command:

```
debug mobile-gateway call-insight ue
```

The parameters of the command, including IMSI, MSISDN, and IMEI, define the sessions to trace.

Call trace supports the following output options to make traces available for further inspection. The options are mutually exclusive:

- **as a PCAP file on local MAG-c storage**

The MAG-c creates the PCAP files automatically on local CF storage. Use the following command to configure the CF storage location:

```
configure mobile-gateway call-insight location
```

The default for the location is CF1. The MAG-c stores the files in the following locations:

- for ongoing traces in the `calltrace/running/` folder
- for completed traces in the `calltrace/finished/` folder

The PCAP files can be downloaded from the system for further offline inspection. Use the following command to list all PCAP files currently on the system:

```
show mobile-gateway call-insight files
```

- **streaming to an external service**

The external service can monitor the packets directly; for example, using the Wireshark packet inspection tool.

- **in the MAG-c debug logging system**

The MAG-c parses and displays packets in a user-friendly format in the system debug logs. The messages are available in the log file and can be viewed on the router depending on the log file configuration. Viewing messages on the router simplifies troubleshooting by eliminating the need to extract PCAP files from the router to view the messages.



**Caution:** The parsing and formatting of packets uses a lot of processing time and may affect the performance of the system. Therefore, Nokia recommends not to send call trace output to the debug logging system when a large number of messages is expected.

For both the PCAP file and streaming output options, the MAG-c can add metadata to the traced packet. This metadata includes information such as the associated IMSI, MAC address, reference point, or UE ID. Nokia makes a plugin available for the Wireshark application to decode and display that metadata. By default, call trace uses the local PCAP output option. To change this and other advanced options:

- Create a profile using the following command:

```
configure mobile-gateway profile call-insight ue
```

- Apply the configured profile to a specific call trace session using the **profile** option of the debug commands mentioned above.

The following advanced configuration options are available:

- The **debug-output** command enables the debug logging output option.
- The **live-output** command enables the streaming output option.
- The **ref-point** and **sba-service** commands allow limiting the call trace to only specific reference points or SBA services. The MAG-c enables by default all reference points and SBA services.
- The **events** command allows tracing internal MAG-c events that are not directly related to an on-the-wire protocol message. Examples include ADB lookups or running a Python script.
- The **size-limit** and **time-limit** commands allow to specify conditions to automatically finish the tracing. By default these limits are set to 10 MB and 1 day.

To show the details for an ongoing call trace capture, use the following commands:

```
show mobile-gateway call-insight ue
show mobile-gateway call-insight bng
```

For example:

```
A:MAG-c>config# /show mobile-gateway call-insight bng mac-address 00:00:00:00:01:01 detail
=====
Call-insight BNG detail
=====
MAC address      : 00:00:00:00:01:01      Status       : running
Circuit Id       : --                    Profile      : default
Remote Id        : --                    Capture format: simulated-p*
UP Node Id       : --                    Time limit   : 86400s
L2 Access Id     : --                    Data limit   : 10MB
L2 Circuit       : --                    Session limit : --
Nr. of captured msgs : 125
|--control-plane msgs: 125
|--user-plane msgs  : 0
|--events          : 0
Size of captured msgs: 22244B
Started          : JAN 11 2023, 14:18:03 UTC
```

```

Ref-points      : dhcp,dns,ga,gn,gp,gx,gxc,pi,radius,rf,ro,s1,s11,s12,s2a
                  ,s2b,s4,s5,s6b,s8,sd,swm,swu,swu-cleartext,sww,sta,sx-
                  n4
SBA-services    : all
User-traffic    : none
Events         : none
Mask-name      : N/A
Live output     : N/A
Output file base : bng_000000000101_230111_1418
-----
Number of call-insight debug jobs: 1
Nr. of dropped user-plane packets: 0
-----
Note:Reference points field above is applicable only to control-plane messages.
=====

```

### Related topics

[In-band control plane and BNG-UP selection](#)

[PFCP connectivity failure](#)

## 2.2.7 CP session state and VM resilience

For scalability, the Nokia MAG-c distributes the session state functionality over multiple CMG virtual machines. This section provides a high-level overview of the MAG-c virtual machines.

The following types of virtual machines are used in a MAG-c to handle session state:

- **OAM-VM**

The OAM-VM handles all centralized functionality, such as IP address allocation and session ID assignments, and is the management interface for the configuration (CLI, SNMP, and so on). It creates the sessions, allocates basic resources such as session ID, and distributes the sessions over the SM-VMs. It balances the load of the sessions over the SM-VMs. To provide redundancy, two OAM-VMs should be deployed on different host systems.

- **SM-VM**

The SM-VM is also known as the CP-VM, or the MSCP. The SM-VM handles the session management after session creation. -VMs must be configured using N:K redundancy using the DB-VM for resilience. The MAG-c does not support 1:1 SM-VM redundancy. The following example shows the configuration of the resource pool.

```

*A:G-c>config>mobile>system# info
-----
    resource-pool 1 redundancy many-to-many gateway 1
        card 3
        card 4
        card 5
        card 6
    exit
    group 1 resource-pool 1
        no shutdown
    exit
    group 2 resource-pool 1
        no shutdown
    exit
    group 3 resource-pool 1
        no shutdown
    exit
-----

```

- **LB-VM**

The LB-VM forwards packets to a SM-VM or an OAM-VM based on load-balancing decisions. The LB-VM directly forwards session management messages to the correct -VM. To provide redundancy, two LB-VMs should be deployed on different host systems.

- **DB-VM**

The DB-VM stores SM-VM session states. In the case of a SM-VM failure, a redundant SM-VM recovers the state from the database to take over. An active redundant DB-VM is not required. Whenever a restart of a DB-VM is detected, all -VMs repopulate the database. Therefore, it is sufficient that a VM orchestrator spawns a new DB-VM upon failure. For redundancy reasons, the DB-VM should not be deployed on the same host as any SM-VM. If DB-VM and SM-VM are on the same host and the system fails, the sessions on the M-VM are irretrievably lost. The following example shows the configuration of database connectivity.

```
A:MAG-c>>config>mobile>pdn>cdbx# info
-----
                cloud-db-profile db
                interface "toDB" router "Base"
-----
A:MAG-c>>config>mobile>profile# info
-----
                cloud-db "db"
                no description
                server 203.0.113.5 port 5678
                no shutdown
                exit
                exit
-----
```

## 2.2.8 Prefix delegation as a framed route

*When configured properly, the MAG-c can signal a PD prefix as a framed route to the BNG-UP.*

When a PD prefix is allocated to a session, the MAG-c can be configured to signal the PD prefix as a framed route instead of as an explicit session address to the BNG-UP. When the PD prefix is signaled as a framed route, the BNG-UP cannot identify that the signaled prefix originated from a DHCPv6 PD lease, and treats it as any other IPv6 framed route.

To have the PD prefix signaled as a framed route, set the **pd-as-framed-route** command in the **config>mobile>profile>adb>entry>address-assignment** context to **true**.

To optimize host resource consumption on a Nokia BNG-UP, the framed route is signaled with a :: next-hop address. As with regular framed routes, it is a requirement that the **ip-anti-spoof** command in the **config>mobile>profile>adb>entry** context is set to **false**.

### Related topics

[AAA-based address assignment](#)

## 2.2.9 Lawful intercept

*Configure LI on BNG CUPS and learn about the content of the LI notifications.*

### Description of LI

Lawful intercept is a legally sanctioned official access to private communications. To provide intercepted private communications to law enforcement officials, a service provider or network operator collects communication of a private subscriber or organization using an LI security process.

### Subscriber events on the MAG-c

To perform LI on BNG CUPS, you need to configure both the MAG-c and the BNG-UP. To configure that the MAG-c includes RADIUS attributes in the accounting messages, use the **up-info**, **up-subscriber-id**, and **subscriber-id** commands in the **config>mobile>profile>charging>bng>radius>session>include-attribute** context.

The MAG-c and the BNG-UP have the following responsibilities for the LI functionality.

- When configured to perform LI, the MAG-c reports the subscriber and LI events.
- The BNG-UP provisions LI targets and supports mirroring of LI packets.

The MAG-c reports the subscriber and LI events through SNMPv3 and RADIUS to the LI mediation gateway. The LI mediation gateway uses the reported subscriber ID to enable LI on the BNG-UP.

The BNG-UP creates a new subscriber ID every time the subscriber logs on. For more information about LI on the BNG-UP, see *7750 SR and VSR BNG CUPS User Plane Function Guide*.

### Subscriber ID and IP address notifications for LI meditation devices

The MAG-c notifies the LI meditation gateway about the LI and subscriber events. The key parameter in the notifications includes the BNG-UP subscriber ID and the BNG-UP IP address. The LI mediation gateway uses the key parameter to provision the LI targets (using the subscriber ID) directly on the BNG-UP (using the IP address).

The MAG-c writes the following information in logs and includes it in RADIUS accounting messages:

- real subscriber name; for example, John Smith
- auto-generated BNG-UP subscriber ID; for example, 549
- BNG-UP IP address; for example, 3.3.3.3

The following is an example of a MAG-c log.

```
767 2020/08/07 13:24:41.990 UTC WARNING: MOBILE_CUPS_BNG #2003 Base CUPS-BNG
"CUPS BNG new subscriber created: Sub-Id '549', externally assigned alias (if any)
'John Smith', UP IP 3.3.3.3"
```

The following is an example of the VSAs in a MAG-c RADIUS accounting message.

```
Alc-sub-Id-str = John Smith
Alc-UP-IP-Address = 3.3.3.3
Alc-UP-Subscriber-Id = 549
```

The LI mediation gateway uses the information in the log and in the accounting message to detect possible LI targets. If the information points to an LI target, the LI mediation gateway sends an SNMPv3 command to the IP address of the BNG-UP to set up an LI target on the subscriber ID on the BNG-UP.



**Note:** The BNG-UP automatically adds `_cups_` to the auto-generated subscriber ID; for example, `_cups_549`.

For further information about the BNG-UP LI target provisioning and LI packet mirroring, see *7750 SR and VSR BNG CUPS User Plane Function Guide*.

For further information about LI access through SNMPv3, see *7450 ESS, 7750 SR, 7950 XRS, and VSR System Management Guide*.

## 2.3 Fixed access sessions

*Learn about the key identifiers for fixed access sessions, IBCP tunnels, BNG-UP selection, limits on the number of sessions, session lockout, and the PPPoE and IPoE setup flows.*

### 2.3.1 Layer 2 circuit

*The layer 2 circuit is the combination of the layer 2 access ID and the VLAN parameters.*

Fixed access sessions have direct Ethernet connectivity from the client device to the BNG. The MAG-c assumes that the Ethernet connectivity is configured. The MAG-c makes an abstraction of the underlying technology and topology. The MAG-c only needs the opaque Layer 2 access ID (I2-access-id) that uniquely identifies the access context of a session. The access context can be a port, a LAG, a pseudowire, an EVPN service, or anything that provides Ethernet connectivity. The BNG-UP defines the content of the I2-access-id. The MAG-c does not interpret it.

The I2-access-id is called the logical port in BBF TR-459. The MAG-c is aware of all Ethernet parameters such as MAC address, S-VLAN, and C-VLAN. Nokia uses the I2-circuit for the combination of the I2-access-id and VLAN parameters.

### 2.3.2 In-band control plane and BNG-UP selection

*The MAG-c creates IBCP tunnels for messages between the BNG-UP and MAG-c. Initial packets use the default IBCP tunnel. At session creation, the MAG-c sets up a per-session tunnel. The BNG-UP selection is based on the default tunnel used for the triggering packet.*

Fixed access sessions send in-band MAG-c messages over the connection to the BNG-UP. As defined in BBF TR-459, the MAG-c creates GTP-U tunnels (called IBCP tunnels) between the BNG-UP and the MAG-c to forward in-band MAG-c messages.

To enable IBCP tunneling, configure an IBCP listening interface using the following command:

```
configure mobile-gateway pdn sx-n4 interface ibcp
```

The following example shows the configuration of the interfaces for an Sx-N4 reference point in the Base routing instance.

```
A:BNG-CP>config>mobile>pdn>sx-n4>if# info
-----
pfcf "system"
ibcp "system"
-----
```



The initial MAG-c packets of a fixed access session are sent over a default IBCP tunnel. The MAG-c automatically sets up the default IBCP tunnel when the BNG-UP indicates support for PPPoE or IPoE sessions. Packets sent over the default tunnel signal the Layer 2 access ID (l2-access-id) and the local BNG-UP MAC address to the MAG-c in an NSH header (as defined in BBF TR-459). The MAG-c matches packets coming in over the default tunnel against a UP group and subsequently to a BNG entry point (EP):

- The UP group identifies interconnected UPs on which steering or resiliency is available. This step is optional and if no UP group is matched, the session is set up only in the context of the UP associated with the incoming IBCP tunnel.
- The BNG EP acts as a gateway mechanism and provides basic setup parameters. It is mandatory for a packet to match a BNG EP.

The MAG-c can instruct the BNG-UP to forward only specified packet triggers over the default IBCP tunnel and to ignore other triggers. For example, ignore IPoE triggers if only PPPoE is deployed or the other way around.

Configure the BNG EP and triggers using the following commands:

```
configure mobile-gateway pdn sx-n4 signaling ibcp bng-entry-point
configure mobile-gateway pdn sx-n4 signaling ibcp triggers
```

The following example shows the signaling configuration.

```
A:BNG-CP>config>mobile>pdn>sx-n4>signaling# info
-----
pfcf
  profile "default"
exit
ibcp
  bng-entry-point "start"
  triggers pppoe-discover ipoe-dhcp
exit
-----
```

The MAG-c sets up a per-session tunnel at session creation.

To display IBCP statistics for both the default tunnel and the per-session tunnel, use the following command with the **ibcp** keyword:

```
show mobile-gateway pdn ref-point-stats
```

To clear the IBCP statistics, use the following command with the **ibcp** keyword:

```
clear mobile-gateway pdn ref-point-stats
```

When multiple BNG-UP devices are managed by the same MAG-c, each BNG-UP has its own default tunnel. At session creation, the BNG-UP selection is based on the triggering packet:

- If the triggering packet matches a UP group, the session is tied to an FSG of that UP group. The MAG-c creates the session on the active (and optionally standby) BNG-UPs of the FSG.
- If the triggering packet does not match a UP group, the MAG-c installs the session on the BNG-UP that corresponds with the default tunnel on which the packet was received.

### Related topics

[BNG entry point \(EP\)](#)

[Modeling a resilient BNG-UP deployment using UP groups](#)

### 2.3.3 Session keys and anti-spoofing

*When multiple fixed access sessions use the same key identifier, data plane and user plane discriminators can be used.*

The MAG-c creates a session when receiving the initial triggering packet for a specific session type. Fixed access sessions are by default identified by the key <UP, I2-circuit, MAC address>, for IPoE as well as for PPPoE. In case a session is set up in the context of a resilient UP group, the UP and L2 circuit keys are internally mapped to a common key based on the UP group configuration. This guarantees that packets coming from different UPs within the same UP group all match the same session.

When multiple sessions share the same key, the remote ID or circuit ID can be used as a discriminator. The circuit ID and the remote ID are in the following sources:

- DHCP: option 82, sub-option 1 (circuit ID) and sub-option 2 (remote ID) as defined in RFC 3046
- DHCPv6: Option 18 (circuit ID) as defined in RFC 8415 and option 37 (remote ID) as defined in RFC 4649. The enterprise ID in the remote ID option is ignored for DHCPv6.
- PPPoE: BBF vendor-specific tags 1 (circuit ID) and 2 (remote ID) as defined in TR 101

To enable discrimination of multiple sessions with the same key, use the following command:

```
configure mobile-gateway profile bng entry-point entry multiple-sessions-per-mac
```

If multiple sessions per key are enabled, and no remote ID or circuit ID can be found, the MAG-c sets up the session anyway. No other sessions for the same <UP, I2-circuit, MAC address> key can be set up.



#### Note:

Nokia does not recommend enabling multiple sessions for the same key for IPoE. It is only supported for single-stack IPv4 sessions that are DHCP-triggered. IPv6 and any other triggers are not supported.

Session setup needs to finish within a hard-coded 60 second timeout. The session is set up when it is stable and the protocol-specific timers or state machines are running (for example, lease timers, or LCP state machine). If the session setup does not finish in time, the session is deleted.

Because data packets do not contain a remote ID or a circuit ID, data plane rules need an additional data plane differentiator. The session ID is an additional data plane differentiator for PPPoE. IPoE requires the enabling of IP anti-spoofing to get an additional data plane differentiator.

Data plane rules on the BNG-UP use the following keys to match data traffic to a specific session:

- IPoE: <I2-circuit, source MAC address> or <I2-circuit, source MAC address, source IP address>
- PPPoE: <I2-circuit, source MAC address, session ID> or <I2-circuit, source MAC address, session ID, source IP address>

To add the source IP address to the key, set the following command to **true**:

```
configure mobile-gateway profile authentication-database entry ip-anti-spoof
```

If not explicitly specified, the **ip-anti-spoof** command is enabled for all sessions, except when L2TP LAC is enabled. For more information about L2TP LAC, see [L2TP Access Concentrator](#).

## 2.3.4 Subscriber identification

*The subscriber identification is equal to the session key, but you can define alternative subscriber keys.*

For fixed access sessions, the BNG subscriber key is by default equal to the session key, that is, there is a single session per subscriber. When a different subscriber scope is needed, alternative subscriber keys can be defined.

To define alternative subscriber keys, use the **multi-session-key** command in the **config>mobile>profile>bng>ep>entry>subscriber-identification multi-session** context.

### Related topics

[Subscribers](#)

## 2.3.5 Session limits

*The MAG-c enforces limits on the number of sessions. These limits can be configured within a specific scope; for example, per MAC address.*

The MAG-c supports limits on the number of sessions within a specific scope; for example, per Layer 2 access ID. The session limits are applied when the session is created, before authentication. Changing a session limit only affects new sessions. Existing sessions are not removed to align with new session limits.

The session limits are configured in the following context:

```
configure mobile-gateway profile bng entry-point entry
```

The MAG-c enforces session limits in the following order:

### 1. per MAC address

When enabled, multiple fixed access sessions for the same MAC address are supported. By default, the support for multiple sessions per MAC is disabled. Use the **multiple-sessions-per-mac limit** command to limit the maximum number of sessions per MAC address. To change the limit, you must first administratively disable the BNG entry point using the following command:

```
configure mobile-gateway profile bng entry-point shutdown
```

### 2. per subscriber

When enabled, multiple fixed access sessions for the same subscriber are supported. By default, the support for multiple sessions per subscriber is disabled. Use the **subscriber-identification multi-session session-limit** command to limit the maximum number of sessions per subscriber. To change the limit, you must first administratively disable the BNG entry point using the following command:

```
configure mobile-gateway profile bng entry-point shutdown
```

and then enable the **subscriber-identification multi-session** command.

### 3. per Layer 2 access ID

Use the **session-limits per-l2-access-id** command to set a limit for the maximum number of sessions per Layer 2 access ID (l2-access-id). The system limits the maximum number of sessions per Layer 2 access ID (for example, port) and the associated BNG-UP. By default, two sessions with the same Layer 2 access ID on two different BNG-UPs do not count for the same session limit. However, if sessions are set up in the context of a UP group, the system maintains the limit per UP group and Layer 2 access ID combination. In this case, two sessions with the same Layer 2 access ID on different BNG-UPs in the same UP group count for the same session limit.

#### 4. per Layer 2 circuit

Use the **session-limits per-l2-circuit** command to set a limit for the maximum number of sessions per Layer 2 circuit (l2-circuit). The system limits the maximum number of sessions per Layer 2 circuit and the associated BNG-UP. By default, two sessions with the same Layer 2 circuit on two different BNG-UPs do not count for the same session limit. However, if sessions are set up in the context of a UP group, the system maintains the limit per UP group and Layer 2 circuit ID combination. In this case, two sessions with the same Layer 2 circuit ID on different BNG-UPs in the same UP group count for the same session limit.

#### 5. per UP

Use the **session-limits per-up** command to set a limit for the maximum number of sessions per BNG-UP. In case the sessions are set up in the context of a UP group, this limit applies to the UP group instead because sessions can dynamically move between BNG-UPs in one UP group.

It is possible to configure conflicting limits for the same context. For example, two sessions with the same Layer 2 access ID can match two different BNG EP entries. Both entries can have a different session limit. The enforced limit at session creation is the limit that is configured in the matching BNG EP entry.

Limits enforced for non-resilient contexts and resilient UP group contexts are never mixed. For example, if a BNG EP is configured with a per UP session limit of 100, the system allows up to 100 non-resilient sessions on that BNG-UP, and another 100 resilient sessions on any UP group linked to that BNG-UP.

#### Related topics

[Session keys and anti-spoofing](#)

[Subscriber identification](#)

### 2.3.6 Session lockout

*Session lockout is an optional feature to prevent DoS attacks and credential brute-force attacks.*

When the session lockout feature is enabled, the MAG-c blocks a client if the sum of the number of the following triggering events reaches a specified threshold within a specified time window

- session setup failure
- disconnection of an established session

If a client is in the locked-out state, the MAG-c drops all packets coming from the client.

The following events unblock a locked-out client:

- expiration of the block timer
- clearing of the locked-out state of the specified client using the **session-lockout** command in the **clear>mobile>bng** context

To enable session lockout, see [Enabling session lockout](#).

To display information for sessions that are subject to session lockout monitoring or to display the number of sessions that are in locked-out or monitoring state, use the **session-lockout** command in the **show>mobile>bng** context.

### 2.3.6.1 Enabling session lockout

To prevent attacks, enable session lockout in the BNG EP.

#### Procedure

**Step 1.** Define a session lockout profile.

Use the **session-lockout-profile** command in the **config>mobile>profile>bng** context. The profile includes:

- failed-attempts (the threshold)
- attempt-window (the time window)
- block (the block timer)

**Step 2.** Reference the session lockout profile in the BNG EP entry.

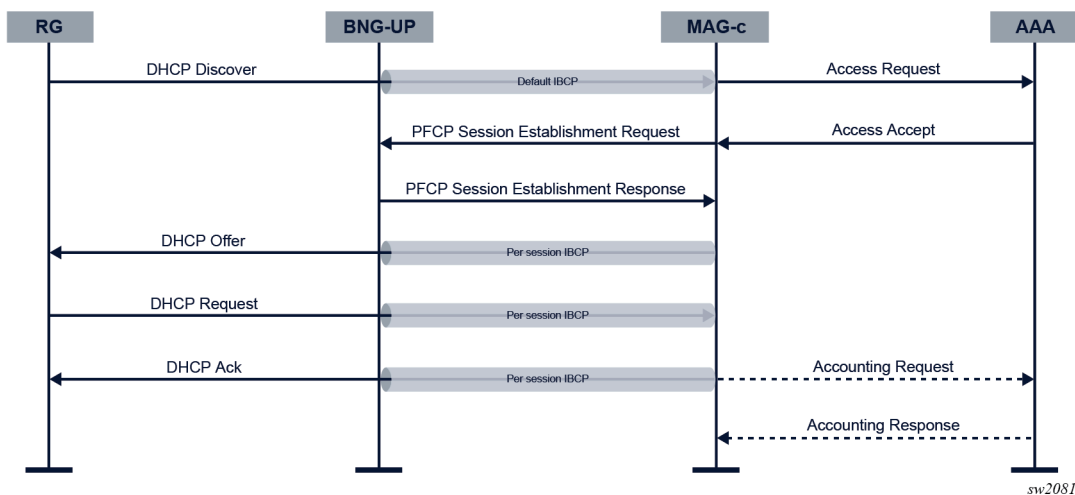
Use the **session-lockout-profile** command in the **config>mobile>profile>bng>ep>entry** context. Session lockout is enabled for the sessions that match a BNG EP entry with a referenced session lockout profile.

### 2.3.7 IPoE

IPoE does not involve a lower-layer connectivity protocol. Address assignment protocols directly trigger IPoE sessions. The MAG-c supports DHCP, DHCPv6, and ICMPv6 RS as triggering protocols.

The figure shows an example of an IPoE session setup flow with DHCP as triggering protocol.

Figure 2: IPoE session setup with RADIUS authentication



When the MAG-c receives the initial triggering packet over the default IBCP tunnel, the following actions are executed.

1. The MAG-c matches the triggering packet with a BNG EP, and creates an IPoE session using the configured keys. The MAG-c assigns the BNG EP and an IPoE profile. The IPoE profile defines the following behavior of the MAG-c.

- The MAG-c sets dot1p and DSCP values in Ethernet and IP headers of CP messages to the IPoE client.

- For packets that trigger session creation, the MAG-c verifies the DHCP client Ethernet address (chaddr field). If the DHCP client Ethernet address does not equal the Ethernet source MAC address, the packet is dropped and no session is created.
2. The MAG-c assigns a single authentication flow to the session and starts the authentication. Subsequent triggers for the same IPoE session do not trigger re-authentication.
  3. The MAG-c allocates addresses for the session.
  4. The MAG-c creates data plane rules and per-session IBCP tunnels on the BNG-UP.
  5. The MAG-c assigns addresses over the per-session IBCP tunnel.
  6. If accounting is provisioned, the MAG-c starts accounting for the session.

The following events cause deletion of an IPoE session.

- The session setup is not done within the hard-coded 60 second timeout. The timer starts on reception of the initial trigger and stops when an IP stack (for example, DHCP lease) is set up.
- The AAA triggers a failure; for example, RADIUS-initiated disconnect.
- An operator gives an explicit clear command for the session.
- No more DHCP or DHCPv6 lease is running, for example because of a lease timeout or an explicit release message. Because of the lack of client-generated messages, SLAAC assigned addresses are not tracked independently and require an active DHCP or DHCPv6 lease in the IPoE session.

#### **Related topics**

[Session keys and anti-spoofing](#)

[Authentication](#)

[Address assignment](#)

[Accounting and charging](#)


### **2.3.8 PPPoE**

PPPoE session setup is linked with the PPPoE discovery protocol as defined in RFC 2516. PPPoE session setup consists of the following phases:

1. PPPoE discovery phase to enable PPP over Ethernet
2. PPP LCP negotiation to negotiate the PPP link connection
3. Authentication
4. IP network connectivity using NCPs, such as IPCP and IPv6CP. In the case of IPv6, this is followed by IP assignment protocols, such as SLAAC or DHCPv6.

The figure shows an example PPPoE session setup flow for IPv4 network connectivity on a CUPS BNG, using PAP as the authentication protocol.

Figure 3: PPPoE session setup flow



To initiate a PPPoE session, the RG sends a PPPoE PADI discovery packet. The packet is sent over the default IBCP tunnel because no session context is known yet.

When the MAG-c receives the initial triggering packet over the default IBCP tunnel, the following actions are executed.

1. The MAG-c matches the triggering packet with a BNG EP. The BNG EP provides the following PPPoE-specific parameters:
  - **pppoe-profile**  
The PPPoE profile contains parameters for each phase in the PPPoE setup. It also specifies the dot1p value for all PPPoE control plane packets that the MAG-c sends.
  - **authentication-flow**  
The configuration of the authentication flow contains PADI authentication (optional) and PAP/CHAP authentication (mandatory).
  - **allocation-scope**  
By default, the MAG-c allocates a unique PPPoE session ID per combination of Layer 2 circuit ID and MAC address. When an aggregation device ignores MAC addresses and forwards based on PPPoE session ID only, the MAG-c allocates a unique session ID per Layer 2 circuit.
  - **random**  
By default, the MAG-c allocates the first free session ID within the allocation scope, starting with one. However, when randomization is configured, the MAG-c allocates a random unique session ID within the scope.

The following example shows the PPPoE configuration for a BNG EP.

```
A:MAG-c>config>mobile>profile>bng>ep>entry>pppoe# info detail
```

```

pppoe-profile "default"
authentication-flow
  no padi-adb
  pap-chap-adb "bng_auth"
exit
session-id
  allocation-scope l2-circuit-mac
  no random
exit
-----

```

- The MAG-c creates an initial session and allocates the PPPoE session ID. If a PPPoE session with the same key exists, the configuration of the PPPoE profile defines whether to delete or keep the existing PPPoE session. If the existing session is kept, the initial triggering PADI packet is ignored.

To define the behavior in case of conflicts, use the **padi-removes-existing-session** command in the **config>mobile>profile>bng>pppoe-profile** context.

- If the authentication flow configuration in the BNG EP requires PADI authentication, PADI authentication is done. Before sending the PADO, the MAG-c creates a PFCP session on the BNG-UP and a per-session IBCP tunnel. The remainder of the discovery phase uses the per-session IBCP tunnel. The following parameters in the PPPoE profile apply:

- **ac-name**  
By default, the AC name is the system name.
- **generate-ac-cookie**  
By default, the generation of an AC cookie is enabled.

The following example shows the discovery parameters in the configuration of the PPPoE profile.

```

A:MAG-c>config>mobile>profile>pppoe>discovery# info detail
-----
no ac-name
generate-ac-cookie
-----

```

For more information, see RFC 2516.

- The MAG-c performs LCP negotiation as defined in RFC 1661 and RFC 4638. The following parameters in the PPPoE profile apply:
  - **max-mtu** (default **1492**)  
The MAG-c derives a downstream PPPoE MTU based on the maximum MTU value and based on the MRU that is signaled by the PPPoE client. If the MRU is smaller than the maximum MTU, the MRU is used; otherwise the maximum MTU is used. The MAG-c sends the derived downstream PPPoE MTU to the BNG-UP. The BNG-UP applies it on the downstream data path.  
  
The BNG-UP may enforce a lower MTU than the derived downstream PPPoE MTU; for example, because of port limitations. The MAG-c does not learn this lower MTU and cannot send it as the MRU to the PPPoE client; therefore Nokia recommends to align the BNG-UP and MAG-c MTU configuration.
  - **mru** (default **1492**)  
To prevent the PPPoE client from sending packets that are dropped by the BNG-UP, Nokia recommends to align the MRU with the maximum packet size that a BNG-UP can receive.
  - **require-max-payload-tag**



The **require-max-payload-tag** parameter defines whether an MRU above 1492 can be negotiated. When enabled, the MAG-c uses the PPP-Max-Payload tag, optionally received from the PPPoE client during PPPoE discovery, for MRU negotiations. For more information, see RFC 4638.

When **require-max-payload-tag** is disabled, the MAG-c sends the MRU as configured and accepts any MRU value. When **require-max-payload-tag** is enabled, the MAG-c uses the value of the PPP-Max-Payload tag as a limit to MRU values.

If the PPP-Max-Payload tag is not present, or if it is lower than 1492, the limit for MRU values is set to 1492. If the configured MRU is bigger than the limit, the limit is sent as MRU. If the received MRU is bigger than the limit, the limit is used for MTU calculations.

- **keep-alive**

The **keep-alive** parameters consist of an **interval** and a maximum number of **tries**. If no response is received after the configured number of tries, the session is considered disconnected. MAG-c terminates the PPPoE session in that case.

- **renegotiation**

The **renegotiation** parameter defines the LCP renegotiation strategy. When a new LCP Configure Request is received while the LCP stack is already in the Opened state, the new request is ignored and dropped, or the PPPoE session is terminated.

The following example shows the LCP parameters in the configuration of the PPPoE profile.

```
A:MAG-c>config>mobile>profile>bng>pppoe-profile>lcp# info detail
-----
max-mtu 1492
mru 1492
renegotiation terminate-pppoe-session
require-max-payload-tag
keep-alive
  no ignore-magic-numbers
  interval 30
  tries 3
exit
-----
```

5. The MAG-c continues with authentication when the LCP stack is in the Opened state.

The MAG-c supports both PAP (RFC 1334) and CHAP (RFC 1994) authentication. PAP or CHAP authentication is required. The **authentication method** parameter in the PPPoE profile defines the priority between PAP and CHAP as follows:

- **pap**

Only PAP authentication is performed.

- **chap**

Only CHAP authentication is performed.

- **pref-pap**

PAP authentication is performed. When PAP authentication fails, CHAP authentication is performed.

- **pref-chap**

CHAP authentication is performed. When CHAP authentication fails, PAP authentication is performed.

In the case of CHAP, the MAG-c generates a challenge with a random length within the range defined in the **chap-challenge-length** parameter of the PPPoE profile.

The following example shows the authentication parameters in the configuration of the PPPoE profile.

```
A:MAG-c>config>mobile>profile>bng>pppoe>auth# info detail
-----
chap-challenge-length min 32 max 64
method pref-chap
-----
```

Both PAP and CHAP authentication are linked to the authentication flow. On successful authentication, any applicable IP address is allocated and the PFCP session on the BNG-UP gets data plane forwarding parameters. On successful completion of the first IP address assignment protocol, accounting is started if a charging profile is provisioned during authentication.

6. After successful authentication, the MAG-c starts the NCP stacks for the allocated addresses. PPPoE sessions support basic NCP renegotiation with the following limitations.
  - The MAG-c never triggers renegotiation.
  - Configuration Requests received while the NCP stack is in Opened state are discarded. A renegotiation consists of an NCP Termination Request, followed by an NCP Configuration Request.
  - NCP renegotiation does not trigger IP reallocation.
  - When the only remaining NCP stack changes from Opened to the Closed state, LCP termination is triggered.

When a PPPoE session must be terminated, the MAG-c fetches the final counters from the BNG-UP and tears down the session by sending an LCP Terminate Request to the PPPoE client. The following events cause termination of a PPPoE session:

- reception of a PADT or LCP Terminate Request from the PPPoE client
- reception of a conflicting PADI if the configuration of the PPPoE profile defines the deletion of the existing PPPoE session in case of conflicts (**padi-removes-existing-session**)
- reception of a new LCP Configuration Request, if the configuration of the PPPoE profile defines the termination the PPPoE session (**renegotiation** is set to **terminate-pppoe-session**)
- detection of an LCP keep-alive failure
- administrative removal; for example, RADIUS-initiated disconnect or clear commands
- charging quota exhaustion
- all NCP stacks changed from the Opened to the Closed state
- PFCP association loss between the BNG-UP and the MAG-c

### LCP keep-alive offload

The MAG-c offloads the LCP keep-alive function to the BNG-UP if the BNG-UP signals the LCP keep-alive offload capability during the PFCP association.

The MAG-c handles LCP keep-alive until the data plane session is created on the BNG-UP. During the data plane session creation, the MAG-c signals the LCP keep-alive offload to the BNG-UP. At that time, the MAG-c stops running timers for LCP keep-alive monitoring and relies on the LCP keep-alive failure reporting from the BNG-UP.

Detection of an LCP keep-alive failure terminates the PPPoE session.

## Resiliency based on PADO delay

PADO delay is a method to provision basic but deterministic (that is, determined from first use) BNG-UP dual homing. A PADO delay value is provisioned during PADI authentication.

In BNG-UP dual homing, a PPPoE client is connected to two BNG-UP devices. The MAG-c sends the PADO packet via one BNG-UP device without delay and via the other BNG-UP device with delay. The PPPoE client chooses the first received PADO and ignores the second received PADO. In stable conditions, the PADO delay determines a primary BNG-UP choice with the option to fall back to a secondary BNG-UP. The session setup continues on the primary BNG-UP device. The feature is supported regardless of whether there is one MAG-c device for both BNG-UP devices or a different MAG-c device per BNG-UP device.

A prerequisite for deterministic BNG-UP dual homing based on PADO delay is PADI authentication. If PADI authentication is not configured, the delay value is not known.




**Note:** PADI authentication is configured in the authentication-flow parameter of a PPPoE BNG entry.

To configure a delay for PADO packets, you can use one of the following options.

- Use the **pado-delay** command in the **config>mobile>profile>adb>entry>pppoe** context.
- Provide the delay via the Alc-PPPoE-PADO-Delay VSA during RADIUS authentication.

**Figure 4: Resiliency based on PADO delay** shows an example of deterministic BNG-UP dual homing. A PPPoE connection is dual homed to two BNG-UP devices, called East and West. East and West share a common MAG-c. The PPPoE client broadcasts the initial PADI to both BNG-UP devices. The MAG-c handles both PADIs and creates two sessions for it because they have different session keys. The session on East has no delay while the session on West has a 500 ms delay. The PPPoE client chooses the first received PADO sent via East. The session is established on East while the session on West times out. If the setup on East fails, the PPPoE client only gets the (delayed) PADO sent via West and the session is established on West.

*Figure 4: Resiliency based on PADO delay*



## IPCP subnet negotiation

IPCP subnet negotiation allows a BNG to assign an IPv4 subnet to a PPPoE session without the need for out-of-band provisioning (for example, using TR-69 signaling).

A prerequisite for IPCP subnet negotiation is to disable the IP address-based anti-spoofing functionality. To explicitly disable IP address-based anti-spoofing, use the **ip-anti-spoof** command in the **config>mobile>profile>adb>entry** context.

An IPCP subnet can be provisioned via an external AAA server or using local address assignment. To provision the subnet using local address assignment, use the **subnet-allocation** command in the **config>mobile>pdn>laa>network-realm>pool>ipv4** context.

If the client signals support for IPCP subnet allocation, the MAG-c signals the subnet using the non-standard IPCP sub-option 0x90. The MAG-c selects the main session IP address from the subnet and signals it as the regular IPv4 address in IPCP option 0x03. The main session IP address is:

- the address in the Framed-IP-Address attribute in case of an AAA-assigned address
- the first address of the subnet in case of a local assignment

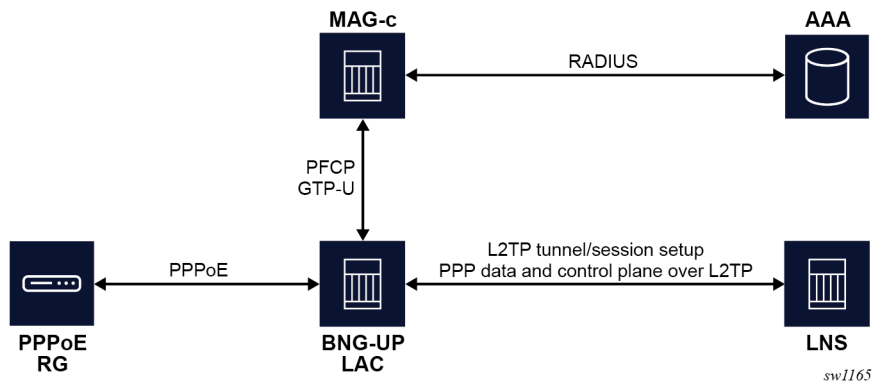
On the BNG-UP, the session is installed with the main session IP address as the session address and the remainder of the subnet as a framed route toward the main session IP address.

In RADIUS accounting, the main session IP address is added in the Framed-IP-Address attribute and the subnet mask is added in the Framed-IP-Netmask attribute. Those attributes are included in the accounting messages if the **address-information** command in the **config>mobile>profile>charging>bng>radius>session>include-attribute** context is enabled.

## L2TP Access Concentrator

When a BNG-UP acts as a LAC for a PPPoE session, it does not act as an IP gateway, but sends the PPP traffic to an LNS using the L2TP protocol.

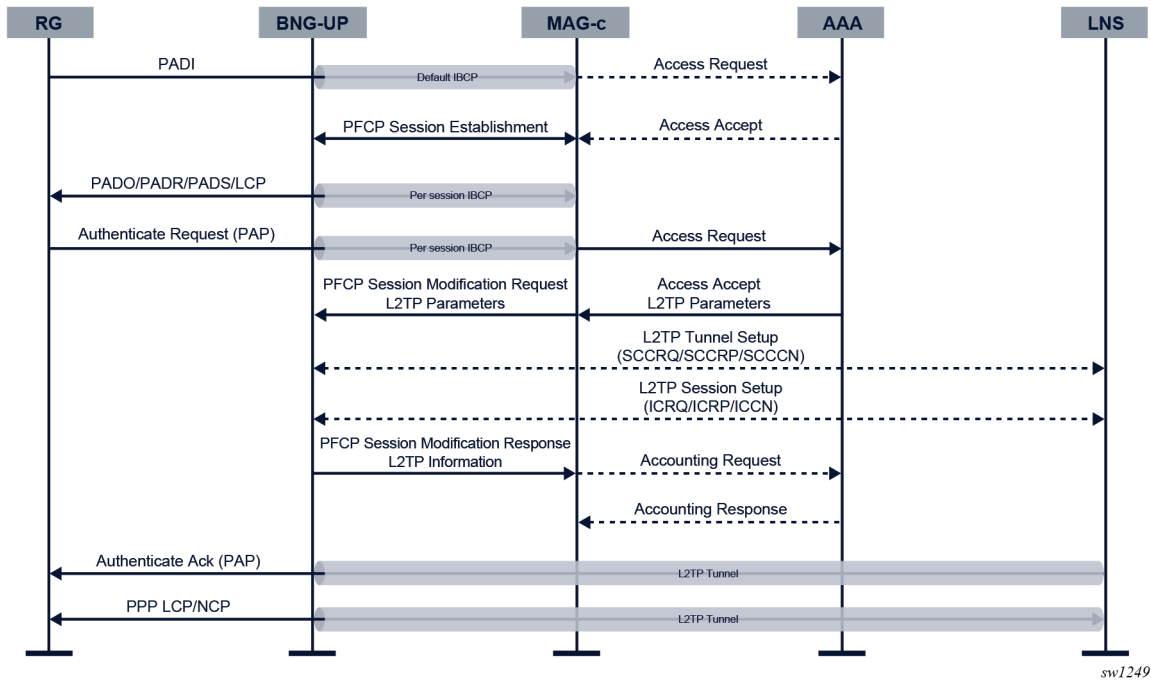
Figure 5: L2TP LAC network components



The functional split of the Nokia BNG CUPS LAC solution between the BNG-UP and the MAG-c is the following.

- The MAG-c does the initial PPPoE setup and authentication.
- The MAG-c gets the L2TP tunnels via local provisioning or via authentication protocols (for example, via RADIUS).
- The MAG-c passes the provisioned tunnels to the BNG-UP via the PFCP protocol.
- The L2TP protocol, including the control plane, runs on the BNG-UP. The BNG-UP signals the L2TP tunnels and the L2TP sessions within those tunnels.
- The BNG-UP does the L2TP tunnel management, including tunnel load-balancing, deny lists maintenance, and preference-based tunnel selection.

Figure 6: LAC-enabled PPPoE session setup flow



Setup of an L2TP LAC-enabled session is the same as the setup of a regular PPPoE session until the PPP authentication step (see [Figure 3: PPPoE session setup flow](#)). Passing a list of tunnel attributes or a name of a locally configured L2TP group in the authentication step enables the PPPoE session for LAC functionality. A mix of locally configured L2TP groups with tunnels provided via authentication is not supported.

To configure a local L2TP group, use the **l2tp-group** command in the **config>mobile>profile>bng** context. The L2TP group contains a list of potential L2TP tunnels and their associated parameters.



**Note:** If PADI authentication is enabled and L2TP parameters are passed, the PPP setup continues on the MAG-c until the PPP authentication is done. In this case, L2TP is triggered after the authentication.

When the authentication indicates L2TP, the PFCP session modification procedure to install PPP data plane rules differs from the one for a regular PPPoE session setup as follows.

- The data plane rules, which contain a list of candidate L2TP tunnels, instruct the BNG-UP to perform L2TP-based forwarding instead of IP-based forwarding. To prevent that the LCP and the authentication negotiation needs to restart between the client and the LNS, the MAG-c includes a list of PPP LCP and Auth proxy parameters or packets to be sent to the LNS.
- The modified control plane rules instruct the BNG-UP to forward PPP control plane traffic such as LCP over the L2TP tunnel. The BNG-UP keeps on sending PPPoE control plane traffic such as PADT packets to the MAG-c. LCP keep-alive offload is not enabled because the BNG-UP forwards LCP keep-alive messages to the LNS.

The BNG-UP performs the following tasks to set up the session and to complete the session setup from a BNG point of view.

1. The BNG-UP chooses an L2TP tunnel from the provisioned list based on parameters such as preference and current tunnel load. If the tunnel is not yet signaled, the BNG-UP sets it up using the Start-Control-Connection (SCC\*) L2TP messages.
2. The BNG-UP sets up an L2TP session using the Incoming-Call (IC\*) L2TP Messages. It includes the LCP and Auth proxy parameters in the session setup messages.
3. The BNG-UP forwards all PPP traffic (including control plane traffic) to the LNS and sends a PFCP Session Modification Response to the MAG-c. This response contains the selected L2TP tunnel, including some meta-data such as the tunnel ID, the session ID, and the CSN, which can be displayed on the MAG-c for operational purposes.

Because the LNS acts as the IP gateway, the MAG-c does not allocate addresses, or start any IP address assignment protocols. The MAG-c can perform accounting for a PPPoE LAC, but the accounting session does not include information that is only known by the LNS, such as IP addressing information.

Except for any LCP or NCP based triggers, the MAG-c can terminate a PPPoE LAC session in the same way as a regular PPPoE session. To terminate the session, the MAG-c deletes the PFCP session, which triggers the BNG-UP to delete the L2TP session using the L2TP Call-Disconnect-Notify message. The BNG-UP disconnects the L2TP tunnel when the last session on the tunnel is deleted and the tunnel idle-timeout has expired.

The L2TP protocol can terminate a session when the LNS disconnects the session, or when the underlying tunnel is removed (for example, because of a timeout). In this case, the BNG-UP sends to the MAG-c a PFCP Session Report Request with the User Plane Inactivity Report flag set in the Report Type IE. In response, the MAG-c sends a PFCP Session Deletion Request to the BNG-UP to clean up the session.

### Interaction between PFCP and L2TP timeout

The BNG-UP runs a function to detect unreachable L2TP LNS servers:

When an L2TP LNS server is unreachable (the L2TP message times out after the last retry), the BNG-UP puts the LNS server in a temporary deny list and does not reconsider the server for this and any following sessions. When all tunnels for a session are unreachable, the BNG-UP indicates setup failure in the PFCP Session Modification message.

The MAG-c runs a function to detect PFCP message timeouts:

The MAG-c times out the PFCP Session Modification message based on the configuration of the **message-retransmit** command in the **config>mobile>profile>pfcp>pfcp-profile** context. When a PFCP timeout is detected, the MAG-c triggers a PFCP Session Deletion message to clean up the BNG-UP state. Because the BNG-UP does not know the reason of deletion, it cancels the session and, or the tunnel setup without moving the tunnel to a temporary deny list.

Because those two detection functions run in contention with each other, it is important that the L2TP message timeout is lower than the PFCP message timeout. To configure the L2TP message timeout, use the **max-retries-non-established** and **max-retries-established** commands in the **config>mobile>profile>bng>l2tp-group** context, or their equivalents in the authentication attributes. Those commands configure the number of retries before a non-established or established tunnel is considered unreachable. The first retry is sent after one second, the second retry after two seconds, the third retry after four seconds, and each subsequent retry after eight seconds. The total timeout must be lower than the PFCP message timeout. For example, if a maximum of two retries is configured, the total timeout is seven seconds (1 + 2 + 4).

### Related topics

[In-band control plane and BNG-UP selection](#)  
[BNG entry point \(EP\)](#)

[Session keys and anti-spoofing](#)  
[Authentication](#)  
[Address assignment](#)

## 2.4 Address assignment protocols

Get a generic overview of the address assignment protocols supported in a MAG-c solution.

This section provides a generic overview of the address assignment protocols supported in a MAG-c solution.

Some address assignment protocols are specific to a session type (for example, IPCP for PPPoE) while other protocols are common for multiple session types (for example, SLAAC for IPoE and PPPoE).

The following table lists the supported address assignment protocols per session type.

Table 1: Address assignment protocols per session type

	IPoE session	PPPoE session
DHCP (see <a href="#">DHCP</a> )	✓	
IPCP (see <a href="#">IPCP</a> )		✓
DHCPv6 NA (see <a href="#">DHCPv6</a> )	✓	✓
DHCPv6 PD (see <a href="#">DHCPv6</a> )	✓	✓
SLAAC (see <a href="#">ICMPv6 Router Advertisements (RA) and SLAAC</a> )	✓	✓

### 2.4.1 DHCP

The DHCP protocol, as defined in RFCs 2131, 2132, 3046, 4679, and 6842, is supported for IPv4 address assignment. The address allocation provides the address and the associated default gateway address. The subnet mask, as signaled in DHCP, is based on the micro-net subnet as provided by ODSA.

For messages sent by the MAG-c, the source IP, the DHCP server IP option, and the siaddr option are by default equal to the default gateway. To override the default per APN, use the **dhcpv4-server-ip** command in the **config>mobile>pdn>apn** context.

The MAG-c maintains a DHCP lease for every successfully negotiated DHCP transaction and extends the lease on renew or rebind. If a lease expires, the MAG-c considers the IPv4 address for the session down and takes appropriate actions for the corresponding session (for example, bring the session down).

The following sources define the DHCP options sent in messages to the client:

- explicit option values provided by authentication sources
- bulk options signaled during authentication; for example, via the RADIUS Alc-ToClient-Dhcp-Options VSA
- bulk options derived from a locally configured DHCP profile using the **dhcp-profile** command in the **config>mobile>profile>bng** context

- DNS options from local address assignment (used in ODSA)

DHCP Offer and Ack messages to the client are constructed using the explicit option values. The option values of the two bulk sources (authentication and DHCP profile) are appended after the explicit option values.

The following rules apply to the options.

- Only one source can provide DNS or NBNS. If a source with higher priority provides DNS or NBNS, DNS or NBNS are filtered out of lower-priority bulk options if present. The sources have the following priority:
  1. explicit options
  2. authentication bulk options
  3. DHCP profile options
  4. local address assignment options
- Lease time, renew, and rebind timers are only provided by explicit per-session authentication sources and are filtered out of bulk options if present.
- Certain options cannot be configured in the message and are filtered out of authentication bulk options. Overriding these options leads to incorrect DHCP behavior. Examples of these options are subnet mask, router, and DHCP message type.

#### Related topics

[Address assignment](#)

## 2.4.2 IPCP

IPCP is used for PPPoE sessions and is defined in RFC 1332 and RFC 1877. IPCP is used to signal an allocated IPv4 address and any DNS or NBNS address.

After IPCP is successfully negotiated, the MAG-c never initiates an IPCP Terminate Request. The client can bring down the IPCP stack and renegotiate if the underlying session is still alive. However, this does not trigger a reallocation of the IP address.

## 2.4.3 ICMPv6 Router Advertisements (RA) and SLAAC

The MAG-c periodically generates ICMPv6 RA messages when an IPv6 address is allocated for the session. A client can trigger the generation of an ICMPv6 RA message by sending an ICMPv6 Router Solicitation (RS) message, but this is not mandatory.



**Note:** RFC 4861 and RFC 4443 define the ICMPv6 RA messages.

The client uses the source address of the ICMPv6 RA message as its default gateway address. By default, this source address is a link-local address derived from the MAC address of the MAG-c. The MAG-c installs the link-local address on the BNG-UP via PFCP so the BNG-UP can answer any ND request for it.

In some cases, this may lead to address conflicts; for example, when two BNG-UPs are connected to the same layer 2 aggregation. To solve this, you can override the link-local address using the **link-local-address** command in the **config>mobile>profile>adb>entry>interface** context. The MAG-c installs the



override on the BNG-UP via PFCP. While this allows very granular overrides, a Nokia BNG-UP can only have one unique link-local address per realm.

ICMPv6 RA messages use an RA profile. The RA profile is configured during authentication.

Use the **ra-profile** command in the **config>mobile>profile>bng** context to configure the RA profile locally.

The RA profile defines the following parameters:

- **advertisement-interval min** and **advertisement-interval max**

These parameters define the interval between periodical unsolicited ICMPv6 RA messages. The MAG-c sends periodical unsolicited ICMPv6 RA messages with a random interval between the configured **min** and **max**. The random interval is regenerated after every unsolicited RA message.

By default, the maximum advertisement interval is 600s and the minimum advertisement interval is 33% of the maximum interval.

- **force-unicast-mac**

This parameter defines which MAC address to use.

If **force-unicast-mac** is enabled, the MAG-c sends ICMPv6 RA messages to the unicast MAC address of the session, otherwise the MAG-c sends the ICMPv6 RA messages to the all-nodes multicast MAC address (33:33:00:00:00:01).

To avoid sending ICMPv6 RA messages to the wrong client, the **force-unicast-mac** parameter is by default enabled.



**Note:** The destination IP address is always the all-nodes multicast IP address (FF02::1).

- **router-lifetime**

This parameter defines the validity period of the default router after receipt of the ICMPv6 RA message. By default, the **router-lifetime** is equal to (**advertisement-interval max** × 3).

- **reachable-time** and **retransmit-timer**

The **reachable-time** parameter defines the period that a neighbor can be reached after receiving a reachability confirmation.

The **retransmit-timer** parameter defines the interval between retransmitted NS messages. By default, both parameters are set to zero, that is, the MAG-c does not specify a value, and the client can choose a value based on local configurations.

- **hop-limit**

This parameter defines the value of the Hop Limit field in the IPv6 header of outgoing ICMPv6 RA messages.

By default, the **hop-limit** value equals 255 hops.

- **mtu**

This parameter defines whether the MTU option is included in the ICMPv6 RA messages and, if included, what value the MTU option contains. By default, the MTU option equals **not-included**.

- **other-configuration**

This parameter defines whether the other-configuration flag in the ICMPv6 RA message is enabled. If the other-configuration flag is enabled, a client can receive options via DHCPv6 without acquiring an address via DHCPv6; for example, in combination with SLAAC based address assignment. By default, the **other-configuration** parameter is disabled. To indicate whether address assignment via DHCPv6 is

available, the related M flag is automatically set if an DHCPv6 IA-NA or a IA-PD prefix was allocated to the session.

- **on-link**

This parameter defines whether the on-link flag is set in the SLAAC prefixes that are present in the ICMPv6 RA messages.

By default, this flag is set.

When an SLAAC address is allocated to the client, each ICMPv6 RA message includes the SLAAC prefix with the A flag enabled. With the A flag enabled, the client can autonomously allocate an IPv6 address from the signaled SLAAC prefix (as defined in RFC 4862).

The SLAAC prefix contains the preferred and valid lifetime that is learned during authentication. The default values of the preferred and valid lifetime are equal to 7 and 30 days respectively.

The ICMPv6 RA messages do not contain any other prefixes. A prefix that is derived from either DHCPv6 IA-PD or IA-NA, is not present.

The ICMPv6 RA messages include all IPv6 DNS servers that are discovered during session authentication (as defined in RFC 8106).

## 2.4.4 DHCPv6

The MAG-c supports the DHCPv6 protocol, as defined in RFC 8415, with additional support for a Lightweight DHCPv6 Relay Agent (LDRA) between the DHCPv6 client and the BNG-UP/MAG-c as defined in RFC 6221. The MAG-c does not support a full DHCPv6 relay between the client and the BNG-UP and MAG-c.

Within the DHCPv6 lease, the following is signaled to the client:

- an allocated IA-NA address, a IA-PD prefix, or both
- preferred and valid lifetimes
- IPv6 DNS servers
- DUID of the server

Preferred and valid lifetimes can be locally configured or received from an external AAA server in the Alc-v6-Preferred-Lifetime and Alc-v6-Valid-Lifetime VSAs. To locally configure the lifetimes, use the **valid** and **preferred** commands in the **config>mobile>profile>adb>entry>address-assignment>lifetimes** context. Valid and preferred lifetimes are common for all IPv6 addresses of a session.

The server DUID is by default based on the MAG-c system name. To override the default server DUID per APN, use the **dhcpv6-server-duid** command in the **config>mobile>pdn>apn** context.

The MAG-c maintains a DHCPv6 lease for every successfully negotiated DHCPv6 transaction and extends the lease on renew or rebind. The lease time is based on the IPv6 valid lifetime. If a lease expires, the MAG-c considers the IA-NA address or IA-PD prefix for the session down and takes appropriate actions for the corresponding session (for example, bring the session down, or bring NCPv6 down).

As with DHCP, DHCPv6 options can come from multiple sources. The following sources define the DHCPv6 options sent in messages to the client:

- explicit option values provided by authentication sources
- bulk options signaled during authentication; for example, via the RADIUS Alc-ToClient-Dhcp6-Options attribute

- bulk options derived from a locally configured DHCP profile using the **dhcpv6-profile** command in the **config>mobile>profile>bng** context
- DNS options from Local Address Assignment (ODSA)

DHCPv6 Advertise and Reply messages to the client are constructed using the explicit option values. The option values of the two bulk sources (authentication and DHCPv6 profile) are appended after the explicit options.

The following rules apply to the options.

- Only one source can provide DNS. If a source with higher priority provides DNS options, they are filtered out of lower priority bulk options if present. The sources have the following priority:
  - explicit options
  - authentication bulk options
  - DHCPv6 profile options
  - local address assignment options
- Identity Association (IA) options, server DUID, server unicast, relay message, status code, interface ID, and other similar options are filtered out of the bulk options because these must be in full control of the MAG-c.
- In case of LDRA, all options are included in the Relayed message.

If an IA-PD prefix or IA-NA address is allocated, the MAG-c sends ICMPv6 RA messages so the client can learn its default gateway address. For consistency, the MAG-c sends DHCPv6 messages with a link-local address that is the same as the source address in ICMPv6 RA messages.

#### Related topics

[DHCP](#)

[ICMPv6 Router Advertisements \(RA\) and SLAAC](#)

## 2.5 Fixed Wireless Access sessions

*FWA sessions obtain network connectivity through a mobile RAN such as (e-)UTRAN (4G) or NR (5G). FWA sessions use the same BNG-UP, charging, authentication, and address management methods (including ODSA) as fixed BNG sessions.*

### 2.5.1 Introduction to FWA

Because FWA sessions obtain network connectivity through a mobile RAN such as (e-)UTRAN (4G) or NR (5G), there are two important differences with fixed access sessions.

1. The initial setup is performed out-of-band and is signaled directly to the MAG-c which acts as a PGW-C, SGW-C + PGW-C, or SMF, whichever is applicable.
2. Sessions have tunneled IP connectivity from the client device to the FWA-UP, which acts as a 4G PGW-U, 4G SGW-U + PGW-U, or 5G UPF, whichever is applicable.

Because FWA sessions are BNG subscriptions, they have the same service requirements as fixed sessions. QoS, service selection, pool management, and charging are also applicable to FWA sessions.

FWA sessions often use mobile-specific terminology. The following lists the most common terms and how they map to an FWA deployment.

<b>User Equipment (UE)</b>	The end device used by the subscriber. The UE and RG can be integrated in one device or can be two separate devices.
<b>International Mobile Subscriber Identity (IMSI)</b>	A globally unique number identifying the subscriber. For FWA, the IMSI is mapped to a subscriber ID. Its value consists of the following three sub-fields. <ul style="list-style-type: none"> <li>• Mobile Country Code (MCC): uniquely identifies a country</li> <li>• Mobile Network Code (MNC): uniquely identifies a mobile network within the country</li> <li>• Mobile Subscriber Identification Number (MSIN): a second globally unique number identifying the subscriber</li> </ul>
<b>Mobile Station International Subscriber Directory Number (MSISDN)</b>	For mobile subscriptions, the public and well-known phone number. For FWA sessions, it can be used as an identifier for compatibility with AAA systems that are mostly MSISDN based.
<b>International Mobile Equipment Identity (IMEI)</b>	A unique identifier for the hardware device used for the connection. Part of the IMEI is the Type Allocation Code (TAC) which is allocated to a specific device model.
<b>Packet Data Network (PDN)</b>	A network the subscriber connects to; for example, the public Internet.
<b>Access Point Name (APN)</b>	Identifies a specific PDN but can also be used to identify a specific gateway within a PDN, or even a specific service type within a PDN. A single UE can be connected to multiple APNs. Each such session is called a PDN session. This model maps directly to FWA and is also common with fixed BNG sessions. An APN consists of the following two sub-fields. <ul style="list-style-type: none"> <li>• Network Identifier (NI): defines the network (PDN).</li> <li>• Operator Identifier (OI): defines the operator's mobile network in which the PDN gateway resides. This consists of the operator's MNC and MCC. This field is optional.</li> </ul>
<b>Radio Access Network (RAN)</b>	The antennas providing wireless connectivity.
<b>Radio Access Bearer (RAB)</b>	A radio channel between a UE and the RAN. Multiple RABs per PDN session can exist.

#### Related topics

[Residential Gateway models](#)

[Service selection](#)

[Fixed access sessions](#)

## 2.5.2 Residential Gateway models

For FWA, there are two common models for a RG.

- An integrated model where the RG provides classic BNG RG functions and acts as a mobile UE. A single IP address can be used to manage this entity. The RG or UE authentication is based on mobile SIM authentication.
- A separate model where there is a logical RG function and a logical UE function. The MAG-c abstracts from how these functions are connected. Often both functions require their own address for management purposes. It is possible to do additional RG device authentication based on PAP/CHAP.



**Note:** The separation is logical and not mapped to on-premises hardware components.

- An integrated RG and UE can have a separate outdoor unit for radio signaling and therefore, can consist of two hardware components.
- A separate model can be a single hardware device that implements the RG and UE function separately with an internal link.

The following figure shows a model using a PPP link between the RG and the UE.

Figure 7: A separate model with PPP connectivity



#### Related topics

[PAP/CHAP authentication](#)

### 2.5.3 Selected and real APN

When a session is initially set up, only one APN is available, as signaled during session management (for example, in the GTP Create Session Request message). This APN is known as the real APN. It is used to pick up a BNG authentication flow that is configured using the **authentication-flow** command in the **config>mobile>pdn>apn>fixed-wireless-access** context.



After initial authentication, a new APN can be returned. This APN is known as the selected APN or the virtual APN. It determines the service selection for the FWA session. If no new APN is returned, the selected service is based on the real APN.

The APN name can be used in one of the following formats for both authentication and accounting purposes.

- **real**  
The real, unmodified APN is used as is, including the OI if it was signaled. For example, the received APN equals `internet.mnc001.mcc001.gprs` and is used as is.
- **real-ni-only**  
The real APN is used, but without the OI part if it was signaled. For example, the received APN equals `internet.mnc001.mcc001.gprs` and is used as `internet`.
- **selected**  
This is the default option. The selected APN is used as is. If no selected APN is available, the system falls back to the **real-ni-only** option.

To configure the format, use the **apn-format** command. The context of the command depends on the use case. The following table shows the different use cases and the related context for the **apn-format** command.

Table 2: Use cases and context for the **apn-format** command

Use case	Context
<p>Use the APN as a match criterion in the ADB.</p> <p>The APN format only applies if the <b>attribute</b> parameter in the <b>match</b> command equals apn.</p>	<b>config&gt;mobile&gt;profile&gt;adb&gt;match</b>
<p>Reflect the APN in the Called-Station-Id attribute for RADIUS authentication.</p> <p>The APN format also applies to the username, if the <b>user-name-format fixed-wireless-access format</b> command in the same context is configured to include the APN.</p> <p> <b>Note:</b> If the mobile UE sends a PPP username as a PCO during session setup signaling, the username is reflected as is. The <b>apn-format</b> command in this context does not modify the username even if it contains an APN.</p>	<b>config&gt;mobile&gt;profile&gt;bng&gt;radius-authentication-profile</b>
<p>Reflect the APN in the Called-Station-Id attribute for RADIUS accounting.</p> <p> <b>Note:</b> The user-name attribute from the authentication is reflected as is in RADIUS accounting. The <b>apn-format</b> command in this context does not modify the username in accounting even if it contains an APN.</p>	<b>config&gt;mobile&gt;profile&gt;charging&gt;bng&gt;radius&gt;session</b>

#### Related topics

[Service selection](#)

[PAP/CHAP authentication](#)

## 2.5.4 Session identification, subscriber identification, and multi-APN support

For FWA sessions, the MAG-c subscriber key is the IMSI and the FWA session key is the pair (IMSI, APN). So multiple FWA sessions can be set up for the same UE/RG if those sessions use a different APN. This is called multi-APN support. The following lists some use cases of how multi-APN can be used.

- Have a different APN for Internet, video, and voice services.
- Have a different APN for public Internet and private work-from-home VPN connections.
- Have a different APN for a UE management IP address in case of the separated UE/RG model.

Each session can be mapped to a different L3 service and a different FWA-UP. For example, to optimize the BNG-UP resource usage, a low-throughput management-only session can be installed on a different FWA-UP than a high-throughput Internet session.

### Related topics

[Residential Gateway models](#)

## 2.5.5 FWA-UP selection

Because FWA session setup messages are sent out-of-band, an FWA session is not tied to a FWA-UP as is the case with fixed access. To select a FWA-UP, the following configuration is needed.

- Define a UP peer list using the **up-peer-list** command in the **config>mobile>profile>pfcp** context.
- Configure all FWA-UPs as peer in the UP peer list using the **peer** command in the **config>mobile>profile>pfcp>up-peer-list** context. The IP address to use as parameter in the **peer** command is the source IP address of the PFCP association of the FWA-UP.
- Set the FWA-UP selection to **true** for each peer using the **upf-selection** command in the **config>mobile>profile>pfcp>up-peer-list>peer** context.
- Configure the APNs that are supported on the peer using the **apn** command in the **config>mobile>profile>pfcp>up-peer-list>peer** context.
- Make a reference to the configured UP peer list using the **up-peer-list** command in the **config>mobile>pdn** context.



**Note:** Other configuration commands under the **config>mobile>profile>pfcp>up-peer-list>peer** context including **network-realm**, **apn nas-ip**, **apn pco-option**, and **apn tai-lai-list** are not applicable to FWA-UPs and must not be configured.

When a session is set up, it automatically selects a peer from the UP peer list that is enabled for that APN. Load-balancing over the FWA-UPs is done in a round-robin fashion.

## 2.5.6 Address signaling methods and deferred allocation

The 3GPP specifications define two methods to signal an allocated IPv4 address.

- **directly in NAS messaging during session setup**  
From MAG-c point of view, the address is included in the session setup messages; for example, in the GTP Create Session Response message.
- **via DHCP**  
The IPv4 address is signaled after the session setup completes.

Which method is used depends on the following parameters, in order of precedence. Each parameter can be absent, or can have the value NAS or DHCP.

- the Alc-FWA-IPv4-Signaling-Method RADIUS VSA
- the configuration of the **ipv4-signaling-method** command in the **config>mobile>profile>adb>entry>fixed-wireless-access** context
- the 00AH (IP address allocation via NAS signaling) and the 000BH (IPv4 address allocation via DHCPv4) PCOs as signaled by the UE. If only one is present, that method is used.

- the configuration of the **default-ipv4-signaling-method** command in the **config>mobile>pdn>apn>fixed-wireless-access** context of the selected APN where the session is created
- NAS, by default

If IPv6 is enabled, a UE/RG interface-identifier is derived from the link-local address of the MAG-c. This interface-identifier is signaled via NAS to the UE/RG, which derives a link-local address from it when needed; for example, to send ICMPv6 RS messages.

IPv6 global addresses always use deferred allocation and are signaled in-band after session setup completion. In case of SLAAC, the UE/RG can optionally use the signaled interface-identifier to construct a global address.



**Note:** Any allocated SLAAC prefix is sent in a GTP Create Session Response message for the MME/HSS. It is not further propagated to the UE/RG via NAS signaling. If no SLAAC prefix is allocated but another IPv6 stack is allocated, the prefix signaled in GTP is set to 0:0:0:0.

For FWA, deferred allocation can be used to assign an address to an RG function in the separated UE/RG model. In this case, the UE can forward the relevant DHCP, DHCPv6, or ICMPv6 messages to the RG without maintaining the stack itself.

#### Related topics

[DHCP](#)

[Address assignment](#)

[Residential Gateway models](#)

## 2.5.7 QoS

FWA sessions support the generic BNG QoS functionality and are additionally subject to specific mobile QoS constructs. This topic gives a high-level overview of how mobile QoS works. For information about the specific subset that is supported for FWA, see [4G QoS attributes](#).

In mobile networks, the concept of a bearer (4G), or a QoS flow (5G), defines the QoS. At least one default bearer or QoS flow exists to which all traffic is mapped by default. More dedicated bearers and QoS flows can be created to which specific traffic is mapped using PCC rules.

Each bearer or QoS flow is classified as either GBR or non-GBR, and QoS treatment on the BNG-UP is applied accordingly. The specified rate values are sent to the BNG-UP using PFCP QER IEs.

- GBR bearer or QoS flow: a GBR and an MBR value are applied.
- Non-GBR bearer or QoS flow: a common MBR is applied to all bearers or QoS flows of this session. For 4G, this value is known as the APN-AMBR, for 5G this value is known as the session AMBR.

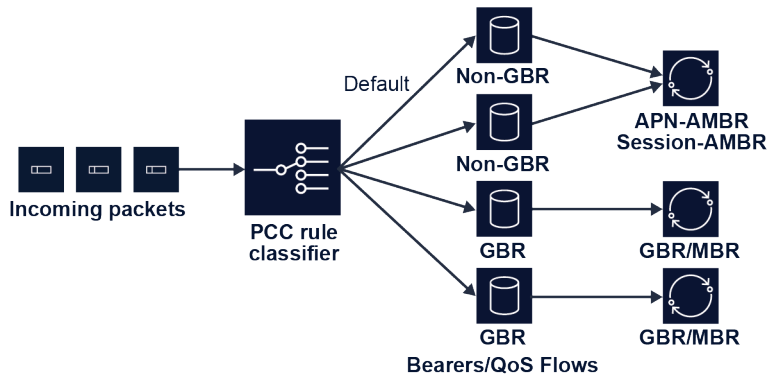
The following figure shows a high-level example of QoS handling. It shows a setup with four bearers or QoS flows, of which two are GBR and two are non-GBR (including the default).



**Note:** In case there is only a default bearer or QoS flow, a PCC rule classifier can be absent.



Figure 8: Example of a bearer or QoS flow rate enforcement



sw1251

Each bearer or QoS flow has QCI (4G), 5QI (5G), and ARP (4G/5G) values. The RAN uses these values, which are transparent to the MAG-c and the FWA-UP, with the exception of the GBR or non-GBR determination of the bearer or QoS flow.

- For 4G, standardized QCI values and the corresponding GBR or non-GBR classification are specified in 3GPP TS 23.203 section 6.1.7.2.
- For 5G, standardized 5QI values and the corresponding GBR or non-GBR classification are specified in 3GPP TS 23.501 section 5.7.4.
- Both QCI and 5QI can be operator-specific, and in this case GBR or non-GBR classification is provisioned together with the QCI or 5QI provisioning. The operator-specific range is from 128 to 254 as defined in 3GPP TS 24.301 section 9.9.4.3 (4G) and 3GPP TS 24.501 section 9.11.4.12 (5G).

The access technology defines how the system learns the different values.

## 2.5.8 PAP/CHAP authentication

A mobile UE can send PAP/CHAP authentication credentials as a PCO option during the session setup signaling. When such a PCO option is received, the MAG-c uses these credentials in the authentication instead of the locally configured credentials.

- The PAP/CHAP username can be used for the relevant ADB match criteria. For RADIUS authentication, the username has precedence over the configuration of the **user-name-format fixed-wireless-access format** command in the **config>mobile>profile>bng>radius-authentication-profile** context.
- The PAP/CHAP credentials are reflected in RADIUS instead of the **password** configured in the **config>mobile>profile>bng>radius-authentication-profile** context.

PAP/CHAP authentication is typically used for separate RG/UE models, where a PPP link is present between the RG and UE. PAP/CHAP authentication is performed over the link. The UE pretends successful authentication of PAP/CHAP and moves the session to the NCP state. The MAG-c gets a CHAP challenge and a CHAP response in PCOs. In cases where the PAP/CHAP based authentication on the MAG-c fails, the UE tears down the PPP session using the LCP terminate messages.

### Related topics

[Residential Gateway models](#)

## 2.5.9 Session lifetime and held addresses

The session management procedures define the lifetime of an FWA session; for example, a GTP Delete Session Request message. An FWA session can exist without any stack being signaled to the client if deferred allocation is used. Stack timeouts do not drive session deletion. If a session management procedure deletes a session with active IP stacks, those stacks are by default released.

For an integrated RG/UE, this is not a problem because the RG/UE knows of the session failure and knows that IP stacks are lost.

For a separated RG/UE model, the UE can have forwarded those stacks to a RG. Depending on the RG/UE connection model, it is possible that the RG is not aware of the failure and keeps using assigned addresses until their signaled lifetimes expire.

To solve this, FWA sessions support holding addresses when a session gets deleted. To configure this, use the **address-hold-time** command in the **config>mobile>profile>adb>entry>fixed-wireless-access** context. The hold time can be set to a fixed amount or be based on the longest remaining address lifetime. In both cases, the hold time is limited to 10 days.

When the hold time is configured, the MAG-c keeps a state for these addresses for the specified hold time. If the ODSA local address assignment allocated the addresses, the ODSA address hold time is automatically increased to the configured hold time. Other than that, there is no link between held addresses and the original allocation source of the addresses (for example, ODSA or AAA). For example, if a held address allocated by ODSA is removed, it is not explicitly released for ODSA. It needs to time out independently.

Held addresses are linked to the FWA session key pair (IMSI, APN). If a new session for the key of an held address is created, the held address is picked up and re-used. The interaction with specific address assignment methods is as follows:

- **local address assignment (ODSA)**

The new session requests the held address from ODSA. The request cancels any running ODSA hold time and ODSA links the address to the new session. In case ODSA cannot assign the address, the session setup fails; for example, ODSA was not the original allocation source of the held address.

- **AAA**

The held addresses are ignored and removed in favor of the AAA-signaled address.



**Note:** When a session setup fails, the held addresses are permanently removed to avoid infinite retries.

Because held addresses do not interact with the original allocation method, Nokia recommends to only enable holding addresses when the address assignment source of a particular session is stable over session creation and deletion. If the address assignment source is not stable, session setups can fail or addresses can be hold for a long time. For example, consider a locally assigned ODSA address being put in the hold state for ten days. If the session is re-established but indicates an AAA address, the held address is not used and not released toward ODSA. The ODSA address keeps its ten day hold time and is not usable during that period.

## 2.5.10 Headless mode for FWA

Headless mode as described in section [Headless mode](#) is partially supported for FWA sessions. When the connection between the FWA-UP and the MAG-c fails, the currently installed sessions on the FWA-UP remain unaffected and continue forwarding data.

However, when new FWA procedures are initiated, such as a mobility, idling, or reactivation, the headless mode procedure fails and immediately terminates the session on the MAG-c. Terminated FWA sessions do not wait for the FWA-UP confirmation as described in section [Headless mode](#), but release the addresses back to the ODSA while in headless mode. To prevent a FWA-UP from announcing these released addresses, Nokia recommends setting the path restoration time lower than the hold time configured for any ODSA pools used by FWA sessions. Use the following command to set the path restoration time:


```
configure mobile-gateway profile pfcg pfcg-profile path-restoration-time
```

## 2.5.11 4G and 5G NSA option 3 sessions

[Figure 9: Basic FWA network](#) shows the elements involved in a simple FWA network with a MAG-c operating as a combined PGW-C and SGW-C, and a FWA-UP operating as a combined PGW-U and SGW-U. The figure highlights the communication between the elements with the interface name if applicable and for the BNG relevant interfaces, the protocol being used. The high-level role of the MME, SGW, and PGW elements is as follows:

- **MME**  
The MME orchestrates the basic connectivity of the UE and how this connectivity evolves during the lifetime of a session (for example, mobility, idling, paging).
- **SGW**  
The SGW provides advanced data forwarding functionality such as buffering packets for idling UEs or providing a mobility anchor for inter-eNodeB mobility. Often this functionality is combined with a PGW in a single network element. In a CUPS environment the SGW is split in an SGW-C and an SGW-U element.
- **PGW**  
The PGW provides PDN connectivity to an IP network. It can perform additional authorization to PDN-specific AAA servers; for example, authorization using the RADIUS protocol. It enforces charging functionality. In a CUPS environment the PGW is split in a PGW-C and a PGW-U element.

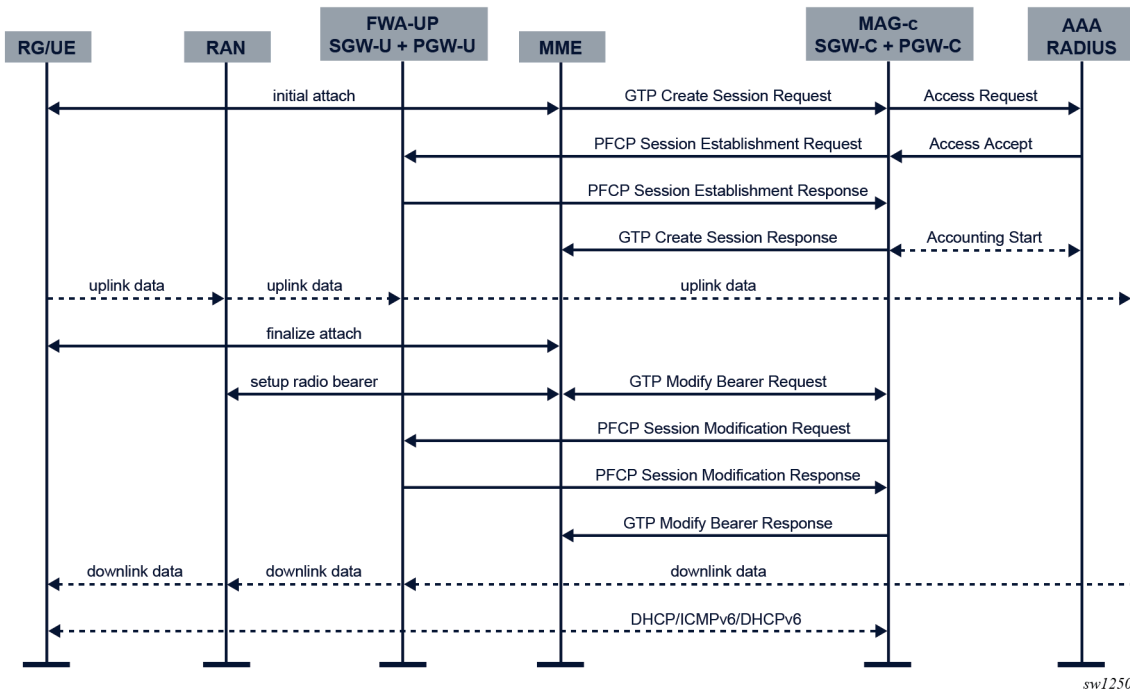
Figure 9: Basic FWA network



To enable 4G FWA sessions for a combined gateway, the GTP-C S11 and S5 interfaces need to be enabled on the MAG-c. The S5 interface is between the SGW and PGW and is internal for a combined gateway.

- To enable the S5 interface, use the **gtp-c** command in the **config>mobile>pdn>s5>interface** context. This interface is not necessarily used to terminate packets, but its IP must match the “PGW S5/S8 Address for Control Plane or PMIP” address that is signaled in the GTP Create Session Request.
- To enable the S11 interface, use the **gtp-c** command in the **config>mobile>pdn>s11>interface** context. This interface transfers the GTP-C packets.
- To identify the L3 service on the FWA-UP where the GTP-U packets are received, configure the **s1u-interface-realm** parameter of the **gtp-c** command in the **config>mobile>pdn>s11>interface** context.
- Configure the EPC node name using the **epc-node** command in the **config>mobile>pdn** context. The MCC and MNC value under this node must match the MCC and MNC of the IMSI for new FWA sessions. Otherwise the session is seen as roaming, which is not supported for FWA.

Figure 10: 4G FWA session setup



The following describes the message flow and actions to set up a 4G FWA session.

- An initial attach sequence is initiated between the UE and the MME, which leads to a GTP Create Session Request message toward the MAG-c. This message contains three sets of parameters:
  - protocol-specific data to initiate the stateful control session between the MME and the MAG-c
  - the required identification and subscription parameters to initiate session management for the PDN session



**Note:** These parameters always include IMSI, IMEI, APN, PDN type, and can be extended with additional values such as MSISDN, static IP addresses, and APN-AMBR.

- PCOs sent from the client; for example, a request for DNS servers, a request for deferred allocation, or PDN-specific authorization using encapsulated PAP/CHAP messages
2. Based on the signaled APN, the MAG-c looks up the APN under the **config>mobile>pdn>apn** context. To enable FWA functionality, the APN must be configured with an FWA node in the **no shutdown** state using the **fixed-wireless-access** command in the **config>mobile>pdn>apn** context. If there is no FWA node configured, it is not an FWA session and the regular mobile gateway session management applies. The authentication flow of the FWA node configured using the **authentication-flow** command in the **config>mobile>pdn>apn>fixed-wireless-access** context is used to authenticate the session. FWA-specific ADB match criteria or RADIUS include attributes can be used.
  3. The MAG-c selects a FWA-UP and sets up a PFCP session. The PFCP session carries the full session data, including all BNG-specific parameters. At this point, the RAN has not yet allocated a downlink GTP-U F-TEID, so the downlink data packet is buffered. For uplink traffic, the FWA-UP allocates an F-TEID and sends it to the MAG-c.
  4. The MAG-c sends a GTP Create Session Response message to the MME. This message includes the uplink GTP-U F-TEID as allocated by the FWA-UP. It includes the IPv4 address if one was allocated and signaled via NAS. The MME completes the attach toward the UE and sets up the radio bearers toward the RAN. At this point, uplink data traffic can be sent.
  5. As part of the radio bearer setup, the RAN allocates a downlink GTP-U F-TEID. The MME signals the downlink to the MAG-c using a GTP Modify Bearer Request Message. The MAG-c forwards it to the FWA-UP using a PFCP Session Modification Request message which changes the downlink rules from buffering to forwarding with the provisioned F-TEID. The FWA-UP acknowledges the PFCP message to the MAG-c. The MAG-c acknowledges the GTP message to the MME. At this point, downlink traffic is possible.
  6. The data layer connectivity is ready, but it is possible that not all IP addresses allocated to the RG are signaled to the UE. The addresses are signaled over the data connection using the DHCP, ICMPv6, and DHCPv6 protocols.

## Mobility and idling

FWA supports the X2-based handover, the tracking area update, the UE triggered service request, the network triggered service request, and the S1 release procedures. Variants of these procedures with an SGW change are not supported.

Many of these procedures require interoperability with standard RAN deployments which assume these procedures are available by default. Handover procedures are useful to provide RAN resiliency, where a UE/RG switches to a new antenna using handover procedures if its current antenna fails.

Most of these procedures require updating the FWA-UP to handle the new traffic rules. This update is done using the PFCP Session Modification Request messages.

- For an X2-based handover, the F-TEID field in the Forwarding Parameter IE is updated.
- For an S1 release procedure, the forwarding rule is removed and replaced by a buffering rule that instructs the FWA-UP to buffer packets.
- For a UE triggered service request procedure, the buffer rule is removed and replaced by a forwarding rule with an F-TEID. Any buffered packet is sent to the new F-TEID.

## PDN type selection

FWA sessions need an assigned PDN type, which can be ipv4, ipv6, or ipv4v6. During setup, an initial wanted PDN type is signaled in GTP. This PDN type is checked against the local configuration done with the **pdn-type** command in the **config>mobile>pdn>apn** context. The PDN type acts as an allow-list. If

the signaled PDN type does not match a configured PDN type, the selection depends on the following configuration in the same **config>mobile>pdn>apn** context.

- If the **pdn-type-conversion** command is not enabled, or the PDN types are not compatible (for example, ipv4 is signaled and ipv6 is allowed), the session setup fails.
- If the **pdn-type-conversion** command is enabled and the PDN types are compatible (for example, ipv4v6 is signaled and ipv4 is configured), the session is allowed and a downgrade to a compatible PDN type is performed (for example, ipv4). If a downgrade can happen to both ipv4 and ipv6 (for example, ipv4v6 is signaled and both ipv4 and ipv6 are configured), ipv4 is chosen by default, unless the **pdn-type-preferred-ipv6** command is enabled.

To optimize IP address usage, only ODSA addresses that are applicable for the selected PDN type are allocated. For example, if a local IPv4 pool and a local IPv6 PD pool are provisioned, and the PDN type is IPv6, no IPv4 address is allocated.

Similarly, if local or AAA based address assignment allocates addresses for only one stack, the PDN type is downgraded to the type matching that stack if applicable.

## 4G QoS attributes

4G FWA sessions only support default bearers and these default bearers must be of the non-GBR type. This means that only an APN-AMBR is supported for mobile-specific rate-limiting.

ARP, QCI, and APN-AMBR for the default bearer are determined using the following parameters, in order of precedence:

- from RADIUS, the 3GPP-GPRS-Negotiated-QoS-Profile attribute
- from the ADB, the QoS profile configured using the **qos-profile** command in the **config>mobile>profile>adb>entry>fixed-wireless-access** context



**Note:** The QoS profile must reference a profile configured in the **config>mobile>profile>qos-profile** context.

- from the MME/HSS, the parameters in the GTP Create Session Request message

The values signaled in the GTP Create Session Request message can optionally be sent to a RADIUS authentication server using the 3GPP-GPRS-Negotiated-QoS-Profile attribute. To enable this, use the **gprs-negotiated-qos-profile** command in the **config>mobile>profile>bng>radius-authentication-profile>include-attribute** context.

The final QoS parameters are used for the generic mobile QoS functionality. They are included in the GTP Create Session Response message so the MME can signal their values to the eNodeB and the UE.

The APN-AMBR can be dynamically changed during the session lifetime using the following procedures:

- a CoA including the 3GPP-GPRS-Negotiated-QoS-Profile attribute
- an HSS-initiated QoS procedure, where the new APN-AMBR is sent using a GTP Modify Bearer Command message

Both procedures trigger a forced change of the APN-AMBR on the FWA-UP using a PFCP Session Modification Request that updates the applicable QER IE. When successful, the MAG-c initiates the PGW-initiated QoS update procedure to signal the new APN-AMBR to the MME/HSS using a GTP Update Bearer Request message.

## Location tracking

The location for an FWA session can be tracked and learned with procedures such as X2-based handover, UE-triggered service request, location reporting, and tracking area update. The MAG-c requests the MME to report location if the MME signals the "Change Reporting Support Indication" flag and if the BNG charging profile is enabled to track the location.

When RADIUS location tracking is enabled, an Interim Update message is triggered whenever the user location changes. The user location is sent in the 3GPP-User-Location-Info attribute.

To enable location tracking for RADIUS, use the following commands.

- To enable location tracking, use the **user-location-change** command in the **config>mobile>profile>charging>bng>radius>session>update-triggers** context.
- To send the user location in the 3GPP-User-Location-Info attribute, use the **user-location-info** command in the **config>mobile>profile>charging>bng>radius>session>include-attribute** context.

### Related topics

[Address signaling methods and deferred allocation](#)

[Address assignment protocols](#)

[QoS](#)

## 2.6 YANG state

MAG-c supports YANG state data for various control plane states, including session and subscriber states. To retrieve session or subscriber states, first configure one or multiple queries using the commands in the following context:

```
configure mobile-gateway system bng queries
```

A query specifies a set of filters for sessions or subscribers, such as a filter on the MAC address, the Layer 2 access ID, and so on.

The states for the sessions or subscribers that match the filters of a query can be retrieved via NETCONF or CLI. See the 7450 ESS, 7750 SR, 7950 XRS, and VSR System Management Guide for more information about the configuration of NETCONF.

To access the state via CLI, you must switch to the MD-CLI engine. Use the following commands to switch to MD-CLI and to access the state:

```
configure system management-interface cli cli-engine md-cli
//
state mobile-gateway bng queries
```

The first command configures the MD-CLI engine, the second command makes the actual switch to MD-CLI, and the third command accesses the state (in this example BNG query states, but there are other states available in the state context).



### Note:

- The **state** command is the only supported MD-CLI command.
- Model driven configuration via MD-CLI or NETCONF is not supported.

- See the distributed YANG files for more information about the supported YANG states.

### Example: Get the number of sessions for the UP IP 5.5.5.5

```

config>mobile>system>bng>queries# info
-----
        session "sessionNumOnUPF1"
            up-ip 5.5.5.5
            output-options
            count
            exit
        exit

[state mobile-gateway bng queries]
A:admin@BNG-SMF# info
  sessions {
    session "sessionNumOnUPF1" {
      number-of-sessions 1000
    }
  }

```

### Example: Get detailed session state information for the MAC address 02:de:14:00:01:94 on the UP IP 5.5.5.5

```

config>mobile>system>bng>queries# info
-----
        session "session1"
            mac-address 02:de:14:00:01:94
            up-ip 5.5.5.5
        exit

[state mobile-gateway bng queries]
A:admin@BNG-SMF# info
  sessions {
    session "session1" {
      result 1 {
        user-access-type ipoe
        mac-address 02:de:14:00:01:94
        selected-apn "mybngvrf"
        network-realm "mybngvrf"
        ue-id 65856
        pdn-session-id 65856
        up-time 15078
        subscriber-name "auto_sub_5"
        acct-session-id "X000I014024F9491600000005"
        acct-multi-session-id "Y0000000524F9491600000004"
        sub-profile "base"
        sla-profile "base"
        sap-template "defaultsap"
        group-interface-template "defaultgrp"
        call-insight false
        ipv4 {
          prefix-length 24
          default-router 20.20.0.254
          ip-address {
            address 20.20.0.3
            pool "p1"
            origin local-pool
          }
          dhcp-lease {
            server-address 20.20.0.254
            lease-time 604800
          }
        }
      }
    }
  }

```



```
        expiration-time 2023-01-30T16:32:56.0+00:00
        up-time 15078
        last-renew-time 2023-01-23T16:32:56.0+00:00
        renew-time 302400
        rebind-time 529200
        remaining-lease-time 589722
    }
}
up {
    active {
        state created
        l2-access-id "1/1/2"
        s-vlan 334
        c-vlan 0
        pfc {
            remote-ip-address 5.5.5.5
            local-seid 0x00000000000010140
            remote-seid 0x30000000000000001
        }
        ibcp {
            remote-ip-address 5.5.5.5
            remote-teid 0x80020000
        }
    }
}
system-resources {
    group 1
    vm 1
}
ibcp {
    mac-address ea:ac:33:01:02:44
    ip-address 1.1.1.1
    teid 0x40010141
}
ipoe {
    ipv4 {
        ip-address 20.20.0.3
    }
}
}
```

## 3 Address assignment

*Overview of the address assignment, and details on the supported local (ODSA) and AAA-based address assignment.*

### 3.1 Overview of address assignment

For sessions that require direct connectivity to a Layer 3 network, the MAG-c supports the following address assignment options:

- local address assignment via ODSA
- local static address assignment via authentication database
- AAA-based address assignment
- non-provisioned address assignment

Additionally, the MAG-c allocates corresponding prefixes (micro-nets) for the BNG-UP, to allow a BNG-UP device to send aggregate routes without announcing the per-session routes. For IPv4, the MAG-c assigns a dedicated gateway address per prefix. It is possible to select different address allocation methods for different address types of the same session. For example, IPv4 can use AAA-based address assignment while IPv6 PD can use a local pool. However, all allocation methods should be known after session authentication.

After session authentication, an address is allocated based on the local address assignment or the AAA-based address assignment. Addresses are always set up on the BNG-UP as soon as they are allocated, independent of whether they are already signaled in associated assignment protocols such as DHCP or IPCP.

### 3.2 ODSA and local address assignment

*ODSA can be used to assign a local address, to assign an aggregate prefix per BNG-UP, and to derive the default gateway.*

#### 3.2.1 ODSA

*On demand subnet allocation (ODSA) is a dedicated CUPS address assignment system.*

ODSA is a dedicated CUPS address assignment system that can automatically split a common subnet into smaller subnets (micro-nets). The micro-nets are automatically installed on the associated BNG-UP. The BNG-UP announces the micro-nets in routing. ODSA can either assign an address itself (local address assignment), or work in combination with external address assignment systems (for example, AAA-based).

ODSA pools are configured on a per network-realm basis. A network realm represents a single IP routing context and maps to an IP service on the BNG-UP (for example, to a VPRN). ODSA guarantees that there is no overlap between addresses within one network realm.

The main function of ODSA is to assign subnets to an allocation context. The default allocation context is a single BNG-UP. In resilient environments, the allocation context is a single fate-sharing group (FSG). Each ODSA pool consists of one or more prefixes and is either configured in dedicated mode or with a target micro-net length.

- **dedicated mode**

In dedicated mode, a prefix is assigned directly to an allocation context. It is not divided into smaller micro-nets.

To enable the dedicated mode, use the following command:

```
configure mobile-gateway pdn local-address-assignment network-realm pool dedicated
```

- **target micro-net length**

With a target micro-net length, all prefixes are divided into smaller, equally sized, micro-nets. Those smaller micro-nets are assigned to an allocation context.



**Note:**

In case of DHCPv6 prefix delegation, you can allocate a variable prefix length per session and a variable micro-net size.

The following example shows the configuration of an ODSA pool with a target micro-net length.

```
A:BNG-CPF# /configure mobile-gateway pdn 1 local-address-assignment
A:BNG-CPF>config>mobile>pdn>laa# info
-----
      network-realm "hsi"
        pool "hsi"
          ipv4
            micro-net-length 28
            prefix 192.0.2.0/24
            exit
          exit
          ipv6
            pd
              micro-net length 48
              prefix 2001:db8:b00::/40
              exit
            exit
            na
              micro-net-length 120
              prefix 2001:db8:a00::/116
              exit
            exit
          exit
        exit
      exit
    -----
```

A subnet (either micro-net or dedicated prefix) can be assigned to only one context. When the first address of a subnet is assigned to a session, the subnet is assigned to the context of the session (for example, to a BNG-UP). To guarantee that the full subnet can always be announced in routing without introducing routing conflicts, the following applies.

- The subnet is only unlinked from the context after the last address of the subnet is released.
- While a subnet is linked to a context, no address of the subnet can be assigned to another context, even if ODSA does not do the address assignment for the other context.

To generate a log event when the number of available free micro-nets is minimal, set a threshold using the following command:

```
configure mobile-gateway pdn local-address-assignment network-realm pool minimum-free
```

For IPv4 subnets, ODSA also assigns a default gateway address. To define whether the first or the last address in the subnet is selected for the default gateway address, set the *choice* variable to respectively **first-address** or **last-address** in the following command:

```
configure mobile-gateway pdn local-address-assignment network-realm pool ipv4 default-gateway
```

To associate default DNS servers with the ODSA pool, use the following commands:

```
configure mobile-gateway pdn local-address-assignment network-realm pool ipv4 dns
configure mobile-gateway pdn local-address-assignment network-realm pool ipv6 dns
```

Default DNS servers can be reflected in protocols such as IPCP, DHCP, ICMPv6, and DHCPv6, but sessions can get more specific individual DNS servers.

### Related topics

[Variable prefix length and micro-net length](#)

## 3.2.2 Variable prefix length and micro-net length

*For DHCPv6 prefix delegation, you can allocate a variable prefix length per session and a variable micro-net size.*

You can allocate a variable prefix length per session instead of a fixed prefix length for the whole pool in case of DHCPv6 prefix delegation. The following enables a variable prefix length per session.

- Configure a **minimum** and **maximum** prefix length for ODSA using the **variable** command in the **config>mobile>pdn>laa>network-realm>pool>ipv6>pd>delegated-prefix** context.
- Provision different prefix lengths per session during session setup.

ODSA allocates a micro-net per prefix length. For example, consider the following sequence of session setups, all within the same ODSA pool and BNG-UP.

1. Session 1 with prefix length 56 leads to the allocation of micro-net A.
2. Session 2 with prefix length 64 leads to the allocation of a new micro-net B.
3. Session 3 with prefix length 64 re-uses micro-net B.
4. Session 4 with prefix length 56 re-uses micro-net A.
5. Session 5 with prefix length 60 leads to the allocation of a new micro-net C.

When only one micro-net length is configured, this leads to unequally loaded micro-nets. For example, for a micro-net length of 52 with variable prefix lengths 64 and 56, each micro-net can hold up to 4096 /64 prefixes, but only 16 /56 prefixes. When the variable prefix length range is big, this is not wanted. For example, an allowed prefix length range of 48 to 64 requires at least a /44 micro-net, which can cover up to 1048576 /64 prefixes.

To solve this, you can provision a variable micro-net length. Configure a **minimum** and a **maximum** micro-net length using the **micro-net variable** command in the **config>mobile>pdn>laa>network-realm>pool>ipv6>pd** context. With a variable micro-net length, the system automatically chooses the

micro-net length that best fits the corresponding prefix length. The following examples illustrate that the configuration must be done carefully.

- For a variable micro-net length between 42 and 50, and a variable prefix-length between 56 and 64, the system allocates a /50 micro-net for a /64 prefix, a /42 micro-net for a /56 prefix, and a /46 micro-net for a /60 prefix. Each micro-net holds up to 16384 prefixes, independent of its length.
- For a variable micro-net length between 48 and 50, and a variable prefix-length between 54 and 64, the system allocates a /50 micro-net for a /64 prefix which holds 16384 prefixes, and a /48 micro-net for a /54 prefix which holds 64 addresses.

When using variable micro-net lengths, the number of free micro-nets is not deterministic. A micro-net with a small length uses a bigger chunk of the pool than a micro-net with a large length. Therefore, the threshold mechanism, configured using the **minimum-free** command in the **config>mobile>pdn>laa>network-realm>pool** context, is adapted as follows.

- When the **minimum-free** configuration is an absolute value, the threshold takes into account the biggest micro-net size. A log event is generated as soon as there is not enough space in the pool to allocate the configured amount of micro-nets with the **minimum** length.
- When the **minimum-free** configuration is a percentage, the threshold takes into account the smallest micro-net size. A log event is generated as soon as there is not enough space in the pool to allocate the configured percentage of micro-nets with the **maximum** length.

### 3.2.3 Local address assignment

ODSA can act as a stand-alone subnet allocation mechanism for BNG-UP devices, but it can also assign addresses to individual sessions, without the need for additional configuration.

To allocate an address for a session, ODSA performs the following checks.

- Are there subnets already linked to the allocation context of the session?
- Do any of the linked subnets have available addresses?

If the answer to both questions is yes, an address from any of the linked subnets is allocated to the session.

If the answer to one of the questions is no, a new address is allocated from any subnet that is not yet linked to an allocation context. The subnet is automatically linked to the allocation context of the session. If no subnets are available, the address allocation fails.

To exclude one or more address ranges in a prefix from address allocation, use the **exclude-addresses** command (IPv4, NA) or the **exclude-prefixes** command (SLAAC, PD) in the following contexts:

- **config>mobile>pdn>laa>network-realm>pool>ipv4>prefix**
- **config>mobile>pdn>laa>network-realm>pool>ipv6>**
  - **slaac>prefix**
  - **na>prefix**
  - **pd>prefix**

Excluded address ranges can be assigned with other allocation methods (for example, via AAA). In case of IPv6, excluded address ranges are not used for default gateway selection.

To stop assigning addresses from a prefix, configure the drain mode for the prefix using the **drain** command in the following contexts:

- **config>mobile>pdn>laa>network-realm>pool>ipv4>prefix**
- **config>mobile>pdn>laa>network-realm>pool>ipv6>**
  - **slaac>prefix**
  - **na>prefix**
  - **pd>prefix**

Existing address allocations from the prefix remain allocated until the corresponding sessions are terminated.

Local address assignment can be combined with AAA-based address assignment for different address types. For example, IPv4 and IPv6 PD can use local address assignment while IPv6 NA can use AAA-based assignment.

#### Related topics

[AAA-based address assignment](#)

### 3.3 AAA-based address assignment

AAA services can provide an address during authentication. The MAG-c marks the AAA-based address as in use in the ODSA pools and allocates the micro-net to the corresponding context (for example, the BNG-UP). In the case of IPv4, the default gateway is assigned using ODSA.

An AAA-based address can fall within an exclude-addresses range.

Setup of the new session fails in situations such as the following:

- the address is already allocated to another session
- the corresponding micro-net is allocated to a context that does not match the context of the session

The prefix pool on which ODSA operates can be used in the following ways.

- If the AAA service provisions both an address pool and an explicit IP address for the same address type (for example, IPv4 or IPv6 PD), ODSA uses the explicit IP address for assignment and the pool for marking the address and allocating BNG-UP prefixes and IPv4 gateway addresses.
- In the absence of an AAA pool, a pool can be provisioned using the **unmanaged** command in the **config>mobile>profile>adb>entry>address-assignment** context.
- In the absence of any pool during authentication, a pool can be provisioned per APN using the **unmanaged** command in the **config>mobile>pdn>apn>address-assignment-defaults** context.

When no dedicated pools are available, ODSA assigns micro-nets to a context. It is important that the AAA service is aware of the micro-net sizes and that addresses are allocated per context within the scope of a micro-net.

For example, the prefix 192.168.0.0/16 is available, to which addresses are allocated per BNG-UP in the AAA. For example, all sessions of BNG-UP1 fall within 192.168.1.0/24 and all sessions of BNG-UP2 fall within 192.168.2.0/24. In this case, it is not necessary to provision these per-BNG-UP prefixes on the MAG-c. The MAG-c has provisioned a non-dedicated pool with prefix 192.168.0.0/16 and micro-net length 24 and automatically derives the /24 prefixes based on the AAA-based addresses.

The following requirements apply when using ODSA pools for a mix of AAA-based addresses and locally assigned addresses.

- The AAA-based addresses must fall within the configured exclude-addresses ranges to avoid conflicts with local assigned addresses.
- If a pool is not dedicated to a specific context (for example, the BNG-UP), the exclude-addresses ranges should align with a micro-net size. This is required to avoid the case where a locally-assigned address allocates the corresponding micro-net to a different context.

Because of the complexity of the requirements, Nokia recommends to have a non-dedicated pool for AAA-based address assignment and a separate non-dedicated pool for local address assignment.

The MAG-c supports AAA provisioned framed routes for sessions with **ip-anti-spoof** disabled (set to **false**); for example, using the Framed-Route and Framed-IPv6-Route RADIUS attributes. The MAG-c installs these routes on the BNG-UP using the PFCP protocol. The MAG-c does not check these routes for overlap with other framed routes or session allocated addresses.

### 3.4 Non-provisioned address assignment

*Configuration of the support for non-provisioned addresses in the ADB allows an external entity to assign the session address.*

When an external entity managed by a third party (for example, a RADIUS server) assigns the session address, the operator is not aware of the subnet/prefix where the address is assigned from. Therefore, the operator cannot provision this subnet/prefix on the BNG system.

To support the above use case, the operator can configure the support of non-provisioned addresses in the ADB for one or multiple address types. The address assigned by the external entity becomes the unmatching address for the session if the following applies.

- After ADB lookup, non-provisioned addresses are supported for the new session.
- The external source assigns an address of a configured address type.
- The address is not within any subnet/prefix of the configured ODSA pool.

To support unmatching addresses of specified address types, use the **unmatching-prefix allow** command in the **config>mobile>profile>adb>entry>address-assignment** context.

The MAG-c cannot send the subnet/prefix information to the BNG-UP because the subnet/prefix for the unmatching address is not provisioned on the BNG system. Therefore, the BNG-UP has /32 or /128 routes for each unmatching address, or the exact prefix route for an IPv6 SLAAC delegated prefix.

The following applies to an unmatching address via DHCPv4.

- The MAG-c automatically generates a default router address and returns it in the DHCP reply. This auto-generated default router address is not passed to the BNG-UP, so the BNG-UP uses the ARP proxy to answer the client's ARP request for the default router address.

To generate the default router address, the host address part of the assigned address is set to one, or to two if the host address part already equals one. The following examples illustrate the generation of the default router address.

- The assigned address is 172.16.3.139 and the netmask is /28, so the host bits are the last 4 bits. 139 equals the binary number 0b10001011. The value of the host bits does not equal 1. When setting the value of the last 4 bits to 1, it becomes 0b10000001 or 129. The default router address is 172.16.3.129.
- The assigned address is 172.16.3.129 and the netmask is /28, so the host bits are the last 4 bits. 129 equals the binary number 0b10000001. The value of the host bits already equals 1. When

setting the value of the last 4 bits to 2, it becomes 0b10000010 or 130. The default router address is 172.16.3.130.

- If the external source does not return a subnet mask, the MAG-c automatically generates one, and returns it in the DHCP reply.

For an unmatching IPv4 address of a PPPoE session, to configure the IPv4 loopback address used as the BNG address in the IPCP negotiation in the corresponding APN, use the **realm-loopback-address** in the **config>mobile>pdn>apn** context. The MAG-c sends the loopback address to the BNG-UP in the PFCP Session Establishment Request message.

#### Related topics

[BNG EP and ADB lookup](#)

## 3.5 Local static address assignment via authentication database

To return the same configured address to a specific client, the MAG-c supports the following static address options:

- IPv4 address
- IPv6 NA address
- IPv6 PD prefix
- IPv6 SLAAC prefix

Configure the address in the following context:

```
configure mobile-gateway profile authentication-database entry address-assignment unmanaged
```

Optionally, use the same context to configure the ODSA pool name.

- If the authentication returns an unmanaged ODSA pool name with the address, the address is with the ODSA pool.
- If no unmanaged ODSA pool name is returned with the address, the static address assignment is treated as a non-provisioned address assignment.



## 4 Authentication

*Configure authentication for a new MAG-c session, including the RADIUS authentication profile. Learn about the BNG EP and ADB lookup process.*

### 4.1 Overview of the authentication process

The authentication process for a new session on MAG-c performs a lookup in the following order:

1. BNG EP for fixed sessions or APN for FWA sessions
2. authentication flow

The BNG EP or APN lookup returns the following:

- basic configurations for the CP protocol negotiation (for example, the IPoE profile, the PPPoE profile)
- basic session configuration (for example, subscriber identification)
- the authentication flow used to authenticate the session

The authentication flow contains an ordered list of Authentication Databases (ADBs). The MAG-c performs a lookup in each ADB in the list, in the specified order. The lookup returns the following configurations required to create the session:

- session attributes (for example, the SLA profile and the subscriber profile)
- address assignment configuration (for example, the local address pool name)
- optional external AAA authentication (for example, RADIUS)

When both the BNG EP or APN lookup and the authentication flow lookup complete successfully, the MAG-c creates a full forwarding state on the BNG-UP for the session using the session management procedures.

#### Related topics

[Session management](#)

[QoS](#)

### 4.2 BNG entry point (EP)

The BNG entry point (EP) provides information needed in the authentication flow. This section describes how to create and configure a BNG EP.

Use the following command to create a BNG EP.

```
configure mobile-gateway pdn gw-id sx-n4 signaling ibcp bng-entry-point name
```

#### Example

```
configure mobile-gateway pdn 1 sx-n4 signaling ibcp bng-entry-point e1
```



**Note:** Each Sx-N4 reference point can reference only one BNG EP; that is, BNG EP e1 and e2 must not be referenced in the same **config>mobile>pdn>sx-n4** context.

To define the control packet types that trigger the BNG EP lookup, use the following command:

```
configure mobile-gateway pdn gw-id sx-n4 signaling ibcp triggers [pppoe-discover] [ipoe-dhcp]
[ipoe-dhcpv6] [ipoe-router-solicit]
```

### Example

```
configure mobile-gateway pdn 1 sx-n4 signaling ibcp triggers ipoe-dhcp
```

To configure the content of the BNG EP in the BNG profile, use the **entry-point** command in the **config>mobile>profile>bng** context.

The following example shows a BNG profile EP configuration.

```
config>mobile>profile>bng>
-----
  entry-point "e1"
    match 1 attribute up-ip
    exit
    entry "10"
      ipoe
        ipoe-profile "mydefault"
        authentication-flow
          adb "adb1" "adb2"
        exit
      exit
    match
      up-ip 172.16.10.50
    exit
  pppoe
    pppoe-profile "pppoeProf1"
    authentication-flow
      pap-chap-adb "adb3" "adb4"
    exit
  exit
  no shutdown
exit
no shutdown
-----
```

## 4.3 Authentication database (ADB)

Each ADB entry contains three groups of configuration parameters:

- match criteria
- action parameters
- session creation parameters (for example, **sla-profile**)

After the MAG-c chooses the best matched entry in the ADB, the MAG-c executes the configured action. The action can be any of the following types:

- **reject**

The session authentication fails and no subsequent ADB lookups are performed, even if they are configured as part of the authentication flow.

- **accept**  
The MAG-c includes the session creation configuration parameters of the chosen ADB entry for the session creation.
- **radius**  
The MAG-c performs the RADIUS authentication using the RADIUS authentication profile. Configure the RADIUS authentication profile using the **radius-authentication-profile** command in the **config>mobile>profile>bng** context. If the RADIUS authentication succeeds, the MAG-c includes the returned RADIUS authentication attributes and the session creation configuration parameters for the session creation. If the RADIUS authentication fails, the session authentication fails.
- **local\_auth**  
The MAG-c performs a PAP/CHAP authentication using the configured username and password in the ADB entry for the PPPoE session. Configure the username using the **username** command in the **config>mobile>profile>adb>entry>match** context. Configure the password using the **action local-auth** command in the **config>mobile>profile>adb>entry** context.

The MAG-c uses the session creation configuration parameters of all ADBs. The authentication flow contains an ordered list of ADBs. If ADB<sub>x</sub> comes before ADB<sub>y</sub> in the ordered list of ADBs, the values of the parameters in ADB<sub>y</sub> have priority over the values of the parameters in ADB<sub>x</sub>. For example, an authentication flow contains two ADBs, ADB1 and ADB2. If the matched entry in ADB1 returns **sla-profile** foo, and the matched entry in ADB2 returns **sla-profile** bar, a new session is created with **sla-profile** bar.

If a session creation configuration is not explicitly configured (for example, it equals the default value), the ADB lookup returns no value for this configuration. For example, an authentication flow contains two ADBs, ADB1 and ADB2. If ADB1 returns **sla-profile** bar, and the matched entry in ADB2 does not contain an explicit configuration for **sla-profile** (it equals the default value), a new session is created with **sla-profile** bar.

Some session creation configuration parameters support a special **discard** keyword. The **discard** keyword means that the previously returned ADB value for the attribute must be discarded. For example, an authentication flow contains two ADBs, ADB1 and ADB2. If ADB1 returns a value for **bng-charging-profile**, and ADB2 returns **discard** for **bng-charging-profile**, the MAG-c creates the session without **bng-charging-profile**.

#### Related topics

[BNG EP and ADB lookup](#)

## 4.4 Authentication flow

An authentication flow contains the following configuration items:

- trigger packet type; for example, DHCPv4 discovery or PPPoE PADI packet



**Note:** For FWA session, the trigger packet type is not explicitly configured. It is always the first session establishment message, such as a GTP Create Session Request message.

- ordered list of one or more ADBs for the specified trigger packet type

When the BNG-UP sends a trigger packet, the MAG-c performs a lookup in each ADB in the list, in the specified order. Each ADB can return session-related configurations. These session-related configurations can be locally configured or returned from an external AAA server.

An IPoE or FWA session has only one authentication flow. A PPPoE session requires at least one of the following independent authentication flows:

- PADI
- PAP/CHAP

If an ADB lookup fails, the session setup fails. The ADB lookup may fail, for example, if an entry is matched with an action reject or if there is an AAA authentication failure.

If all lookups complete successfully, the MAG-c continues session setup using the combined configurations from all ADB lookups. For example, the BNG EP lookup returns two authentication flows for a new PPPoE session. The authentication flows return the following configuration:

- PADI authentication flow with 1 ADB: ADB1 returns PADO delay value
- PAP/CHAP authentication flow with 2 ADBs: ADB2 configures RADIUS authentication, ADB3 returns a local address pool

In this example, the MAG-c uses the combined configuration result from the three ADB lookups to set up the PPPoE session.

Each session requires an APN for service selection, as described in [Service selection](#). The APN can also provide override for specific configurations. If different types of sources return the same type of configuration (for example, an address pool name), the MAG-c uses the value of the source with the highest ranking. The sources are ranked as follows, with the highest ranked first:

1. AAA
2. Local ADB
3. APN

If different sources of the same type (for example, different ADBs) return the same type of configuration, the MAG-c uses the last returned value. For example, if both ADB-1 and ADB-N return an SLA profile name, and ADB-1 returns SLA profile name X and ADB-N returns SLA profile name Y, the system uses SLA profile name Y because it is the last returned value.

#### Related topics

[Authentication database \(ADB\)](#)

## 4.5 BNG EP and ADB lookup

Both the BNG EP entries and the ADB entries contain session configuration and one or more ordered match criteria. The match criteria are used in the lookup. The session configuration is used in the creation of the session.

### Match criteria properties

Match criteria have the following properties:

- **match-id**  
The match ID defines the priority. The lower the ID, the higher the priority.
- **attribute**

The attribute defines the name of the attribute that is used for the lookup. It is the name of a session attribute. The attribute can be a control protocol field (for example, DHCP option 82 **circuit-id**, **vendor-class**) or metadata of the session (for example, **I2-access-id**).

- **value**  
The value defines the criteria value to which the session value must match for the specified attribute. If the attribute is optional, the value can be empty, meaning any session value matches with the criteria value.
- **optional**  
Match criteria can be optional or mandatory. The attribute of optional criteria does not need to be present in the session data to match the entry. If the attribute of optional criteria is present in the session data, the session value must equal the criteria value to match the entry. An attribute that is not present in the entry can have any value in the session (including not available).
- **string-mask**  
A string mask is applied to the value of the session attribute before comparing it with the value of the criteria. It can be used for supported attributes (for example, **I2-access-id**).

The string mask can be length-based or string-based and can be a suffix or a prefix, as follows:

- **prefix**
  - **length-based**  
The MAG-c removes the specified number of characters from the beginning of the session value.
  - **string-based**  
The MAG-c removes the specified string from the beginning of the session value. An asterisk (\*) can be used as a wild-card in the string mask.
- **suffix**
  - **length-based**  
The MAG-c removes the specified number of characters from the end of the session value.
  - **string-based**  
The MAG-c removes the specified string from the end of the session value. An asterisk (\*) can be used as a wild-card in the string mask.

The following examples show the string that is used to compare the session value of **I2-access-id** with the criteria value for a specific string mask configuration. Assume that the session value of **I2-access-id** equals 1/2/3.

- For **stringmask** equal to **prefix length 2**, the MAG-c removes the first two characters of the session value. The resulting value 2/3 is used to match with the end of the criteria value; for example, the resulting value 2/3 matches with the criteria value 4/2/3.
- For **stringmask** equal to **suffix string "/"**, the MAG-c removes the last slash (/) and everything after it at the end of the session value. The resulting value 1/2 is used to match with the beginning of the criteria value; for example, the resulting value 1/2 matches with the criteria value 1/2/4.

### Default entries

If a BNG EP entry or an ADB entry does not have any match criteria, this BNG EP entry or ADB entry is the default entry. The MAG-c chooses the default entry when there is no other matched entry. Only one default entry is allowed for the BNG EPs and for the ADBs.

## Entry matching

Entries of a BNG EP or of an ADB cannot have the same set of match criteria within the same BNG EP or ADB. In this case, the entry becomes operationally down. The system does allow entries with the same match criteria in different BNG EPs or ADBs.

During a BNG EP or ADB lookup, the MAG-c compares the attributes of the session with the match criteria of all entries in the BNG EP or in the ADB and creates a list of all matched entries. A matched entry is one for which all mandatory match criteria are fulfilled.

At the end of the lookup, the MAG-c chooses the best matched entry from the list of all matched entries for session creation. The MAG-c chooses an entry from the list as following:

- If the list of all matched entries contains only one entry, that entry is the best match.
- If the list of all matched entries contains more than one entry, the MAG-c reduces the list to the entries with the highest number of match criteria. If this list contains only one entry, that entry is the best match.
- If the reduced list of entries with the highest number of match criteria contains more than one entry, the MAG-c selects the entry with matches for the highest priority attributes.

## Example mandatory and optional match criteria

As an example, the match criteria for an ADB entry contain the attribute **l2-access-id** (marked **optional**) and the attribute **up-ip** (mandatory). To call the ADB entry a matched entry, one of the following must be true.

- Both **up-ip** and **l2-access-id** are present in the session and both match the values in the ADB entry.
- Only **up-ip** is present in the session and it matches the value in the ADB entry.

If both **l2-access-id** and **up-ip** are present in the session, but only **l2-access-id** matches the value in the ADB entry, the ADB entry is not a matched entry.

## Example entry matching and selection

The following output defines the configuration of four ADB entries.

```

-----
#first match criteria is UP's IP address
    match 1 attribute up-ip
        optional
    exit
#2nd match criteria is Layer 2 access ID
    match 2 attribute l2-access-id
        optional
    exit
#3rd match criteria is SVLAN
    match 3 attribute s-vlan
        optional
    exit
    entry "10"
        match
            l2-access-id "1/1/2"
            up-ip 172.16.10.50
            vlan
                s-vlan start 100 end 200
        exit
    exit
    subscriber-mgmt
        sla-profile "entry10"
        sub-profile "entry10"
    exit

```

```

        no shutdown
    exit
    entry "20"
        charging
            bng-charging-profile "mybngcharging"
        exit
        match
            l2-access-id "1/1/2"
            up-ip 172.16.10.50
        exit
        subscriber-mgmt
            sla-profile "entry20"
            sub-profile "entry20"
        exit
        no shutdown
    exit
    entry "30"
        match
            l2-access-id "1/1/2"
            vlan
                s-vlan start 100 end 200
            exit
        exit
        subscriber-mgmt
            sla-profile "entry30"
            sub-profile "entry30"
        exit
        no shutdown
    exit
    entry "40"
        match
            vlan
                s-vlan start 100 end 200
            exit
        exit
        subscriber-mgmt
            sla-profile "entry40"
            sub-profile "entry40"
        exit
        no shutdown
    exit
    no shutdown
-----

```

A new session has the following attributes and values:

- **up-ip** with value 172.16.10.50
- **l2-access-id** with value 1/1/2
- **s-vlan** with value 100

The session matches with all ADB entries. The MAG-c chooses the entry 10 because it has the highest number of matching criteria; that is, three matching criteria.

Assume the entry 10 is shut down. Both the entries 20 and 30 have the highest number of matching criteria; that is, two matching criteria. The MAG-c chooses the entry 20 because it has the matching criteria with the highest priority; that is, **up-ip**.

Assume all entries except 40 are shutdown. The MAG-c chooses the only matching entry; that is, the entry 40.

## 4.6 Required minimal configuration for a session creation

To create a session, the MAG-c requires a minimal number of session creation configuration parameters. The table lists the parameters that are required for session creation, as well as the source that contains those parameters.

Table 3: Minimal configuration for a session creation

Configuration	Source
ipoe-profile (IPoE session only)	BNG EP
pppoe-profile (PPPoE session only)	BNG EP
authentication-flow (fixed access)	BNG EP
authentication-flow (FWA)	APN
APN	ADB, RADIUS
address-assignment	ADB, RADIUS
sla-profile <sup>1, 2</sup>	ADB, RADIUS
sub-profile <sup>1, 2</sup>	ADB, RADIUS
group-interface-template <sup>1, 2</sup>	ADB, RADIUS
sap-template <sup>1, 2</sup>	ADB, RADIUS

## 4.7 RADIUS authentication profile

RADIUS authentication is performed when the action parameter in the best matched ADB entry equals **radius**. The RADIUS authentication profile defines the behavior of the RADIUS authentication. To define the profile, use the **radius-authentication-profile** command in the **config>mobile>profile>bng** context.

RADIUS authentication is triggered by ADB lookup, which means it is possible to have multiple rounds of RADIUS authentications during the authentication flow lookup. If during multiple rounds, the same attributes are returned in the Access-Accept message, the last attribute received is used.

A RADIUS authentication profile contains the following configuration commands:

- **radius-group**

The **radius-group** command contains RADIUS server configuration such as address, port, and shared secret. To define the RADIUS server configuration, use the **radius-group** command in the **config>mobile>profile** context. Afterwards, reference this **radius-group** in the **config>mobile>profile>bng>radius-authentication-profile** context.

<sup>1</sup> Required when using a Nokia BNG-UP.

<sup>2</sup> If the BNG-UP contains a template or a profile with the name `default`, the default template or profile is used when the authentication does not return a template or profile. If the BNG-UP does not contain a specific template or profile with the name `default`, the configuration of the parameters is required.



- **user-name-format**

The **user-name-format** command defines the username format for the RADIUS server. To define the username format, use the **user-name-format** command in the **config>mobile>profile>bng>radius-authentication-profile** context.

- **password**

The **password** command defines the password of the RADIUS user. To define the password, use the **password** command in the **config>mobile>profile>bng>radius-authentication-profile** context.

- **include-attribute**

The **include-attribute** command defines the RADIUS attributes to be included in an Access-Request message. To define the attributes to be included, use the **include-attribute** command in the **config>mobile>profile>bng>radius-authentication-profile** context.

The username and password configuration is required for IPoE authentication, PPPoE PADI authentication, and FWA authentication if no PAP/CHAP credentials are provided during FWA session setup.

## 4.8 RADIUS CoA and DM

A RADIUS CoA or a DM message asks for changes in the session or subscriber object.

To enable support for RADIUS CoA and DM messages, use the **interface** command in the **config>mobile>pdn>bng>radius-coa** context.

The interface defines one or more listening interfaces for incoming CoA and DM messages, and the shared secrets.

When the MAG-c receives a CoA or DM message, it makes the requested change to the target object. The *CMG BNG CUPS RADIUS Attributes* list defines the message attributes that can be used to identify one or multiple sessions as target object. Filter on the value yes for the CoA key column to find those attributes in the list. If a subscriber is specified in the request, the MAG-c applies the requested changes to all sessions of the targeted subscriber.

The CoA message contains one or more attributes that define the requested changes; for example, the *Alc-SLA-Prof-Str VSA* defines a new sla-profile for the target object. For more information about the supported attributes, see the *CMG BNG CUPS RADIUS Attributes*.

If the MAG-c applies all requested changes successfully to all targeted objects, the MAG-c sends a CoA-ACK message. If the MAG-c can only apply the requested changes partially or only on a subset of the target objects, the MAG-c sends a CoA-NAK message with an ERROR-CAUSE code 506, and rolls back the changes as follows.

- If the change request is for multiple attributes on a single session and only part of the attribute changes are successful, the MAG-c sends a CoA-NAK message with an ERROR-CAUSE code 404, and rolls back the already applied changes.
- If the change request is for multiple attributes on multiple sessions and the changes are only successful for a part of all the target sessions, the MAG-c sends a CoA-NAK message with an ERROR-CAUSE code 506, and rolls back the applied changes for the sessions that were only partially changed. For example, a CoA message requests to change three attributes on five sessions. The MAG-c successfully applies all attribute changes on session 1, session 2, and session 3, but only applies one attribute change successfully on session 4 and session 5. The MAG-c sends a CoA-NAK message and rolls back the attribute change on session 4 and session 5.

A DM message only contains target objects. The MAG-c removes the sessions of the target objects and sends an ACK message. If the target objects do not exist, the MAG-c sends a CoA-NAK message with ERROR-CAUSE code 503.

If a CoA or DM message contains an unsupported attribute, the MAG-c rejects the request with a CoA-NAK message by default. To ignore unsupported attributes, use the **ignore-unknown-attributes** command in the **config>mobile>pdn>bng>radius-coa** context.

## 4.9 Example configuration

The example configuration in this section is for the following setup:

- IPoE session
- RADIUS authentication
- address pool pool-up-1 for sessions from BNG-UP 1.1.1.1
- address pool pool-up-2 for sessions from BNG-UP 2.2.2.2
- sla-profile basic, sub-profile basic, and authentication with radius-auth-profile-1 for sessions with s-vlan 100
- sla-profile premium, sub-profile premium, and authentication with radius-auth-profile-2 for sessions with s-vlan 200

To achieve the above setup, an authentication flow with three ADBs is used. In this example, the following is returned.

- ADB adb1 only returns the address pool.
- ADB adb2 only returns the sla-profile and the sub-profile, and performs RADIUS authentication.
- ADB adb3 returns the other configuration parameters.

The following example shows the ADB configuration.

```
authentication-database "adb1"
  match 1 attribute up-ip
  exit
  entry "up-1"
    address-assignment
      local-dynamic
        ipv4-pool "pool-up-1"
    exit
  exit
  match
    up-ip 1.1.1.1
  exit
  no shutdown
exit
entry "up-2"
  address-assignment
    local-dynamic
      ipv4-pool "pool-up-2"
  exit
exit
match
  up-ip 2.2.2.2
exit
no shutdown
```

```

    exit
    no shutdown
  exit
  authentication-database "adb2"
  match 1 attribute s-vlan
  exit
  entry "basic"
  action radius radius-authentication-profile "radius-auth-profile-1"
  match
    vlan
      s-vlan start 100 end 100
    exit
  exit
  subscriber-mgmt
  sla-profile "basic"
  sub-profile "basic"
  exit
  no shutdown
exit
entry "premium"
action radius radius-authentication-profile "radius-auth-profile-2"
match
  vlan
    s-vlan start 200 end 200
  exit
exit
subscriber-mgmt
sla-profile "premium"
sub-profile "premium"
exit
no shutdown
exit
no shutdown
exit
authentication-database "adb3"
entry "default"
apn "mybngvrf"
interface
  group-interface-template "defaultgroup"
  sap-template "defaultsap"
  exit
  no shutdown
exit
no shutdown
exit

```

The following example shows the configuration of the BNG EP.

```

*A:BNG-CPF>config>mobile>profile>bng# info
-----
    entry-point "e1"
      entry "10"
        ipoe
          ipoe-profile "mydefault"
          authentication-flow
            adb "adb1" "adb2" "adb3"
          exit
        exit
      no shutdown
    exit
  no shutdown
exit

```

The following example shows the reference to the BNG EP in the **config>mobile>pdn>sx-n4>signaling>ibcp** context.

```
*A:BNG-CPF>config>mobile>pdn>sx-n4# info
-----
      pfcf-association-list "pfcfassoc1"
      interface
        pfcf "system"
        ibcp "system"
      exit
      signaling
        pfcf
          profile "pfcfpro-1"
        exit
        ibcp
          bng-entry-point "e1"
          triggers ipoe-dhcp
        exit
      exit
    exit
```

## 4.10 Web portal authentication

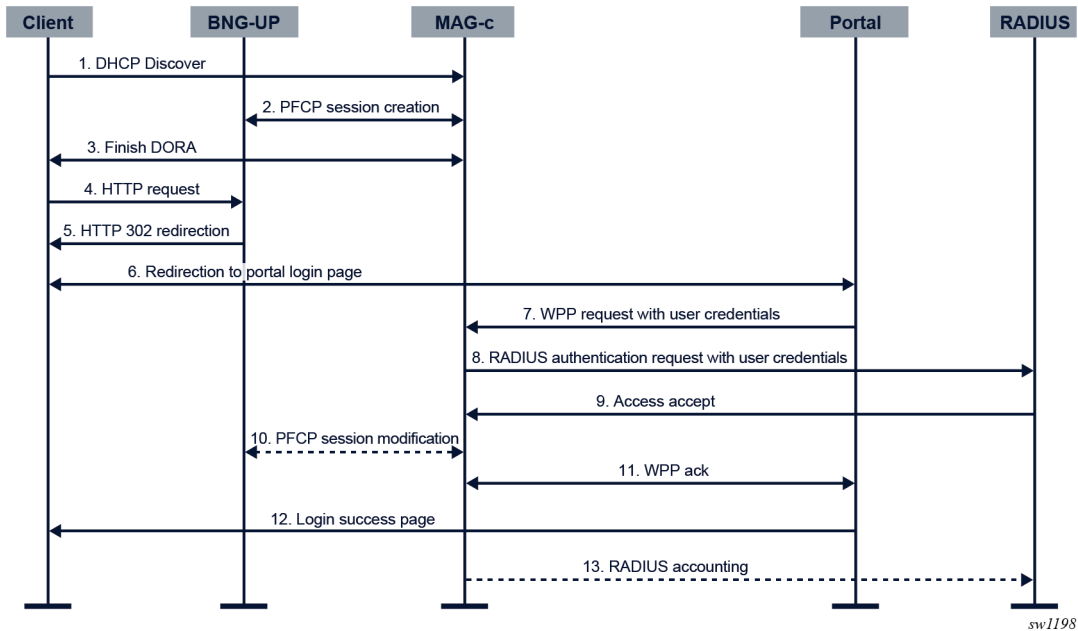
*The BNG supports the WPP protocol to authenticate broadband or WLAN users on a web portal.*

Web portal authentication uses the WPP protocol between a BNG and a web portal server for broadband or WLAN users. The user provides a username and password on the portal page and the web portal triggers the BNG to perform RADIUS authentication with the specified credentials.

### WPP call flow example

The following figure shows a high-level call flow example of WPP authentication for a client using DHCPv4 on a BNG CUPS system.

Figure 11: WPP on BNG CUPS call flow example



1. The client sends a DHCP discovery packet. The BNG-UP forwards it to the MAG-c. The MAG-c authenticates the packet via the authentication flow. The authentication returns two sets of SLA profile and subscriber profile, that is, initial profiles and after-auth profiles.
  - Initial profiles are used for the client before portal authentication. The SLA profile includes a filter that contains an HTTP redirection entry to the portal.
  - After-auth profiles are used for the client after portal authentication.
 The authentication returns the following key WPP attributes:
  - a WPP RADIUS authentication profile, used in step 8.
  - an HTTP redirection URL
2. The MAG-c creates a session on the BNG-UP with the initial profiles and the redirection URL via PFCP.
3. The client finishes DORA with the MAG-c.
4. The client sends an HTTP request.
5. Because of the redirection filter in the initial SLA profile, the BNG-UP intercepts the HTTP request and sends an HTTP 302 redirection response.
6. The client is redirected to the portal. The user provides user credentials (for example, username and password).
7. The portal sends the user credentials in a WPP request message to the MAG-c.
8. The MAG-c sends a RADIUS authentication request to the RADIUS server with the user credentials received in the WPP request.
9. The RADIUS server successfully authenticates the request and sends an Access Accept message. Optionally, the RADIUS server includes an SLA profile and a subscriber profile in the Access Accept message. Those profiles are after-auth profiles that override the after-auth profiles of step 1.

10. If the RADIUS server sent after-auth profiles in the Access Accept message of step 9, the MAG-c modifies the session with those profiles.
11. The MAG-c and the web portal exchange WPP Ack messages.
12. The web portal returns a login successful page to the client.
13. Optionally, the MAG-c starts RADIUS accounting.

### Variables in the HTTP redirection URL

The initial authentication returns an HTTP redirection URL that optionally contains one or multiple variables. The variables are replaced with session-specific values. The supported variables are:

- \$IP — the IPv4 or IPv6 address that triggers the WPP authentication
- \$MAC — the MAC address
- \$URL — the original requested URL
- \$SAP — the Layer 2 access ID and the VLAN tags
- \$SUB — the subscriber ID as a string
- \$CID — the circuit ID, or the interface ID of the subscriber host (in hexadecimal format)
- \$RID — the remote ID of the subscriber host (hexadecimal format)
- \$SYSNAME — CPF\_SYSTEM\_NAME:UPF\_IP

For example, when the initial authentication returns `http://www.example.com/login?sub=$SUB`, and the subscriber ID for the session equals `sub1`, the actual URL is `http://www.example.com/login?sub=sub1`

### HTTP redirection URL override

If the initial authentication contains RADIUS authentication, and the Access Accept message contains the `Alc-Portal-Url` RADIUS attribute, the value of the RADIUS attribute is used for HTTP redirection. The URL in the `Alc-Portal-Url` RADIUS attribute overrides the locally configured HTTP redirection URL.

### Portal group

You can configure up to eight WPP portals in a portal group. The MAG-c can receive WPP requests from any of the configured portals in the portal group.

When the BNG initiates a WPP `NTF_LOGOUT` message, it sends the `NTF_LOGOUT` message to all configured portals in the portal group. The first received `ACK_LOGOUT` stops the retransmission of the `NTF_LOGOUT` message.

A WPP portal group can be used to achieve WPP portal redundancy taking into account the following.

- A portal can only be configured in one portal group.
- A portal can be in a portal group and at the same time be used as an individual portal.
- Portals supporting different WPP versions (version 1 and version 2) are allowed in the same portal group.

### WPP port attribute

The MAG-c uses the WPP port attribute in the WPP protocol messages to identify the port of the session.

The format of the WPP port attribute is CPF\_SYSTEM\_NAME#UPF\_ADDR#L2\_ACCESS\_ID:VLAN1.VLAN2. If the length of the result string exceeds 35 chars, the system truncates it to the first 35 chars.

For example, CUPSBNG1#2.2.2.2#1/2/10:100.200, where

- SAP (the Layer 2 access ID and the VLAN tags) = 1/2/10:100.200
- BNG-UP address = 2.2.2.2
- MAG-c system name = CUPSBNG1

### 4.10.1 Configuring WPP

To configure a minimal WPP configuration, define a WPP listening interface, a WPP portal, a portal group, a RADIUS authentication profile, and configure WPP in the ADB entry.

#### About this task

The steps in this procedure define a minimal WPP configuration.

#### Procedure

**Step 1.** Configure a WPP listening IP interface.

Use the **interface** command in the **config>mobile>pdn>bng>wpp** context.

**Step 2.** Define a WPP portal.

Use the **portal** command in the **config>mobile>profile>bng>wpp** context.

**Step 3.** Define a portal group and include a reference to the portal defined in the previous step.

Use the **portal-group** command in the **config>mobile>profile>bng>wpp** context.

**Step 4.** Define a RADIUS authentication profile.

Use the **radius-authentication-profile** in the **config>mobile>profile>bng** context.

**Step 5.** Configure the WPP context in an ADB entry.

Use the following commands in the **config>mobile>profile>adb>entry>wpp** context.

- Use the **portal-group** command to reference the portal group defined in step 3.
- Use the **wpp-radius-authentication** command to reference the RADIUS authentication profile defined in step 4.
- Use the **initial-profiles** command to specify the names of the initial SLA and subscriber profiles.
- Use the **no shutdown** command to enable the WPP entity in the ADB entry.

**Step 6.** Configure the HTTP redirection URL in an ADB entry.

Use the **http-redirect url** command in the **config>mobile>profile>adb>entry** context.

**Step 7.** Configure the after-auth profiles in an ADB entry.

Use the **subscriber-mgmt** command in the **config>mobile>profile>adb>entry** context to define the after-auth SLA and subscriber profiles.

#### Example

```
*A:BNG-CPF>config>mobile>pdn>bng>wpp# info
-----
interface router "Base" name "system"
```

```
-----
*A:BNG-CPF>config>mobile>profile>bng>wpp# info
-----
        portal "p1"
          address 2001:beef::1
          router "Base"
          source-address 2001:dead::1
          no shutdown
        exit
        portal-group "g1"
          realm "mybngvrf"
          portal "p1"
          no shutdown
        exit
-----
*A:BNG-CPF>config>mobile>profile>adb>entry# info
-----
        apn "mybngvrf"
        dhcp-profile "mydefault"
        http-redirect
          url "http://www.exampleportal.com"
        exit
        address-assignment
          local-dynamic
            ipv4-pool "p1"
          exit
        exit
        interface
          group-interface-template "defaultgrp"
          sap-template "defaultsap"
        exit
        subscriber-mgmt
          sla-profile "base"
          sub-profile "base"
        exit
        wpp
          portal-group "g1"
          wpp-radius-authentication "wpp-rad"
          initial-profiles
            sla-profile "ini-sla"
            sub-profile "ini-sub"
          exit
          no shutdown
        exit
        no shutdown
-----
```



## 5 Accounting and charging

*Learn about the statistics collection from the BNG-UP, time-based and volume-based charging, RADIUS accounting configuration, and buffering of the RADIUS accounting messages for later retransmission.*

### 5.1 BNG charging profiles

*BNG charging profiles define the charging interfaces for a session. A session can be associated with multiple BNG charging profiles.*

BNG charging profiles contain the configuration of the charging interfaces for a session; for example, the RADIUS accounting interface. Charging profiles are assigned to sessions during authentication.

Multiple charging profiles can be provisioned per session. The following example use cases enable the same charging interface in different contexts.

- Two charging profiles support duplicate RADIUS accounting to a main and a backup accounting server. The two charging profiles are identical, except for the target servers.
- Multiple charging profiles support multiple distinct logging systems using the same interface. For example, one profile for RADIUS packet/octet accounting, one for NAT port block logging, and one for session create/delete logging.

To configure BNG charging profiles, use the **bng-charging** command in the **config>mobile>profile>charging** context. To use a BNG charging profile for a specific set of sessions, use the **bng-charging-profile** command in the **config>mobile>profile>adb>entry>charging** context.

Except for the communication with the BNG-UP, there is no interaction between multiple charging profiles for the same session. Each charging profile uses the session data and the triggers (periodic interval or trigger events) to act according its configuration. For example, a messaging failure for one profile does not affect retransmits in another profile.

Concerning the communication with the BNG-UP, the MAG-c tries to optimize the number of messages sent. For example, if the same event triggers multiple charging profiles to fetch BNG-UP data, the MAG-c attempts to send a single request to the BNG-UP instead of a request per charging profile.



**Caution:** The MAG-c guarantees unique charging identifiers (such as Acct-Session-Id) per session but not per profile and session. If the same servers are used in two profiles, the servers may not be able to distinguish between log events which leads to unpredictable behavior. Nokia recommends that you do not define multiple charging profiles with interfaces to the same set of servers; for example, two profiles using the same RADIUS group.

#### Related topics

[BNG EP and ADB lookup](#)

## 5.2 Statistics collection from the BNG-UP

Configure the pull or push model to fetch mid-session PFCP usage reports from the BNG-UP.


A PFCP Usage Report IE contains incremental BNG-UP statistics per session. A usage report only includes statistics that are collected since the previous usage report. The MAG-c aggregates the statistics for the session. This incremental method is robust and can handle a BNG-UP failure. When a BNG-UP failure occurs, the statistics collected since the previous report are lost. However, the MAG-c statistics remain correct and increase monotonically. Similarly, when the BNG-UP resiliency is used, the MAG-c can add counters of both BNG-UPs to calculate a correct aggregate.

The MAG-c supports two models to fetch mid-session PFCP usage reports:

- **pull model**

The MAG-c explicitly requests a usage report in a PFCP Session Modification Request message when it needs up-to-date counters. The MAG-c does not send periodic unsolicited pull requests. For example, it requests a Usage Report for a RADIUS Accounting Request Interim Update message.

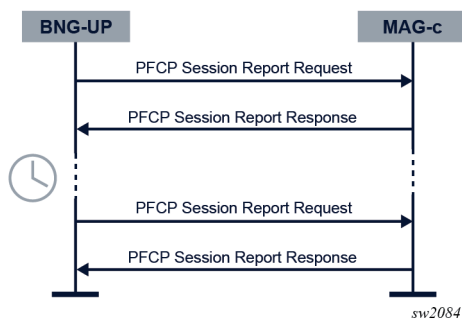
Figure 12: Statistics collection using the pull model



- **push model**

The BNG-UP periodically sends an unsolicited PFCP Usage Report IE in a PFCP Session Report Request message. On reception of the message, the MAG-c records and stores the statistics but does not take any further action.

Figure 13: Statistics collection using the push model



To enable the push model, configure the push interval using the following command:

```
configure mobile-gateway profile authentication-database entry charging statistics-collection-interval
```

A small push interval can reduce counter loss in case of BNG-UP failures, but increases the load on the MAG-c. The system must be dimensioned accordingly.

Push and pull modes can work in parallel. For example, the RADIUS accounting is enabled with an interval of 1 h (pull mode) and the BNG-UP reporting is enabled with an interval of 5 min (push mode). Every hour, the RADIUS accounting message triggers an explicit pull from the BNG-UP to fetch the latest counters, while the BNG-UP sends unsolicited usage reports every 5 min. This allows for a frequent counter update to avoid loss without overloading the AAA server and without the AAA server having outdated statistics.



**Note:**

If the pull interval is similar to the push interval, Nokia recommends to disable the push mode. Having similar intervals for the push and the pull mode increases the load without a direct benefit.

For resilient sessions, the MAG-c maintains one set of counters that are aggregated from all PFCP sessions that have been set up over the lifetime of the encompassing session. The incremental nature of the PFCP reports allow for this without risking duplicate counters. The MAG-c can present these aggregated counters monotonically increasing to other interfaces such as RADIUS accounting, MAG-c charging, and **show** commands. Because of BNG-UP resiliency, back-end systems such as accounting servers do not need to account for sudden counter resets.

In case of hot standby BNG-UP resiliency, there are two PFCP sessions. When performing pull requests, the MAG-c pulls statistics from the active BNG-UP in steady state, because it does not expect traffic from the standby BNG-UP. During switchovers, the MAG-c pulls from both BNG-UPs to fetch the optionally still available final statistics from a failed BNG-UP. The MAG-c installs the push mode on both BNG-UPs. The Nokia BNG-UP, however, is optimized to only send push reports if the reported statistics contain non-zero values. The statistics of a standby BNG-UP in steady state normally contain only zero values.

Independent of mid-session statistics collection, final statistics are always fetched when the session is removed in a PFCP session deletion procedure.

The following statistics are supported:

- aggregate number of bytes upstream and downstream
- aggregate number of packets upstream and downstream, if the BNG-UP signals the MNOP function feature in the PFCP association procedure
- detailed statistics as collected by the Nokia BNG-UP

Examples of detailed statistics are per queue statistics, per policer statistics, and separate IPv4 and IPv6 counters. The content of the detailed statistics depends on the BNG-UP QoS **stat-mode** configuration.

The MAG-c only collects statistics for sessions that are linked to a charging profile with at least one charging interface enabled (for example, RADIUS accounting). A charging profile can be assigned to a session during authentication.

Detailed statistics are collected and available for reporting if the detailed statistics are enabled in the ADB using the following command:

```
configure mobile-gateway profile authentication-database entry charging detailed-statistics
```

When enabled, the MAG-c requests the BNG-UP to send detailed statistics.

If detailed statistics are collected and available for reporting, they can be included in the RADIUS accounting messages. Use the following command to send detailed statistics in the RADIUS accounting messages:

```
configure mobile-gateway profile charging bng-charging radius session include-attribute detailed-statistics
```

When the MAG-c detects a new **stat-mode** or SLA Profile, the detailed statistics are reset. The MAG-c sends the final detailed statistics for the previous **stat-mode** or SLA profile in RADIUS accounting messages if enabled. Because aggregated statistics are not dependent on the **stat-mode** nor on the SLA profile, they are not reset.

**Note:**

On a Nokia BNG-UP, both aggregate and detailed statistics are based on the QoS model. If multiple sessions of the same subscriber share QoS resources, the statistics are collected on a per-session basis, but they do not provide real usage of a specific session. The aggregate of the session statistics are correct for the shared QoS resource. If real usage of per session statistics are required in a multiple sessions per subscriber model, Nokia recommends enabling an SPI per session model on the BNG-UP.

**Related topics**

[Authentication](#)

## 5.3 MAG-c-based charging

*Learn about time-based and volume-based charging.*

### Time-based charging

The MAG-c provides time-based charging using the session timeout mechanism. This session timeout starts after successful authentication. The MAG-c deletes the session when the timer expires.

### Volume-based charging

For basic volume-based charging, the MAG-c compares the statistics to provisioned thresholds. The MAG-c compares the statistics with the thresholds for every received usage report.

The following thresholds are supported:

- total number of upstream bytes
- total number of downstream bytes
- total number of upstream and downstream bytes

To configure the thresholds in the ADB, use the **cp-volume-tracking** command in the **config>mobile>profile>adb>entry>charging** context.

The threshold values can also be provided via RADIUS. For more information, see *CMG BNG CUPS RADIUS Attributes*.

As soon as one of the provisioned thresholds is reached, the MAG-c deletes the session.

For deterministic behavior, Nokia recommends combining the volume-based charging with periodic statistics collection.

**Related topics**

[Statistics collection from the BNG-UP](#)

## 5.4 RADIUS accounting

Learn about the RADIUS accounting configuration in the BNG charging profile and how to enable RADIUS accounting.

The MAG-c supports RADIUS accounting as defined in RFC 2866.

To enable RADIUS accounting, perform the steps in [Enabling RADIUS accounting](#).

The RADIUS accounting configuration in the BNG charging profile includes the following parameters.

- **RADIUS group**

The RADIUS group provides the list of RADIUS servers and load-balancing parameters. It also defines the default interim interval used for sending periodic RADIUS Accounting Request Interim Update messages. Any configuration that defines attribute content such as configuration by the **radius-avp-options** command and the **calling-station-id** command is ignored. Failure handling configuration is not supported for BNG RADIUS charging and is also ignored. For more information about generic RADIUS support, see *7750 SR MG and CMG Configuration Guide*, section *RADIUS Interface*.

Use the **radius-group** command in the **config>mobile>profile>charging>bng>radius** context to reference the RADIUS group in the BNG charging profile.

Use the **radius-group** command in the **config>mobile>profile** context to define and configure the RADIUS group parameters.

- **session accounting parameters**

The accounting parameters for sessions include configuration of attributes to include in accounting messages, and triggers to send the RADIUS Accounting Request Interim Update message.

Use the **session** command in the **config>mobile>profile>charging>bng>radius** context for the session accounting parameters.

For more information about the accounting attributes, their content, the associated include-attribute configuration, and the messages they can appear in, see *CMG BNG CUPS RADIUS Attributes*.

### 5.4.1 Enabling RADIUS accounting

#### Procedure

**Step 1.** Define a BNG charging profile.

Use the **bng-charging** command in the **config>mobile>profile>charging** context.

**Step 2.** Configure RADIUS accounting for the BNG charging profile

Use the **radius** command in the **config>mobile>profile>charging>bng-charging** context.

**Step 3.** Reference the BNG charging profile.

Use the **bng-charging-profile** command in the **config>mobile>profile>adb>entry>charging** context.

**Step 4.** Define match criteria for the ADB so that the BNG charging profile gets assigned to a session during authentication.

#### Related topics

[BNG EP and ADB lookup](#)

## 5.4.2 Session accounting

The MAG-c sends RADIUS accounting messages to start and stop session accounting. Interim update messages contain updates of the accounting data. The interim update messages can be periodic or triggered by an event.

### Start and stop messages

When session accounting is enabled, the MAG-c sends an Accounting Request Start message to the RADIUS accounting server. The exact time when the message is sent in the session setup procedure depends on the session type. Sending the Accounting Request Start message is linked to the data plane creation on the BNG-UP and to the IP address assignment protocols.

The server selection for the Accounting Request Start message follows the generic load-balancing configured in the following context:

```
configure mobile-gateway profile radius-group
```

Any subsequent messages are sent to the same server. Only when the selected server fails, is a new server selected.

When the accounting is started, the MAG-c sends Accounting Request Interim Update (IU) messages. The MAG-c sends the IU messages periodically or based on triggers.

When the session is removed, the MAG-c sends an Accounting Request Stop message, including the final counters.

You can enable or disable session-level charging on a live system, but the new configuration affects only new sessions. Existing sessions continue with or without charging, based on the previous configuration.

### Periodic interim updates

Periodic sending of Accounting Request Interim Update messages needs to be explicitly enabled using the following command:

```
configure mobile-gateway profile charging bng-charging radius session update-triggers periodic
```

When the periodic sending is enabled, the interval for the Accounting Request Interim Update messages can be provisioned as follows, in order of precedence:

1. during the session authentication; for example, using the RADIUS Acct-Interim-Interval attribute
2. using the following command:

```
configure mobile-gateway profile charging bng-charging radius interim-update-interval
```



**Note:** When the value is changed using this command, existing sessions use the changed value after the next scheduled Accounting Request Interim Update message.

3. using the following command:

```
configure mobile-gateway profile radius-group interim-update-interval
```

The interval can be changed during the lifetime of a session by sending a RADIUS CoA with the Acct-Interim-Interval attribute. In this case, a session sends an immediate Interim Update message with the reason Interval-Changed and starts a timer with the new interval.



**Note:** If during RADIUS authentication or CoA an Acct-Interim-Interval attribute with value 0 is sent, periodic interim updates are explicitly disabled. In this case, it is not possible to re-enable periodic interim updates for the session.

When multiple BNG charging profiles are configured, the periodic interim update interval can be provisioned or changed per BNG charging profile using the RADIUS Alc-Charging-Profile-Interim-Interval attribute.



**Note:** Nokia recommends to provision the same interval for all charging profiles with periodic interim updates enabled. In this case, the MAG-c runs a common interval timer, and sends only one message per periodic interim update interval to fetch the statistics from the BNG-UP for all the periodic Accounting Request Interim Update messages.

### Triggered interim updates

The MAG-c sends a triggered Accounting Request Interim Update (IU) message when it detects changes in the session or subscriber data, or when an external system instructs to send an IU message.

The vendor-specific Alc-Acct-Triggered-Reason attribute in the IU message indicates the type of trigger.

See the *MAG-c RADIUS Attributes and IU Triggers* for information about the supported trigger events.

When several trigger events occur at the same time, the MAG-c sends a single IU message with multiple Alc-Acct-Triggered-Reason attributes to include all trigger reasons.

In case one event triggers multiple IU messages, such as an SLA profile change, other simultaneous trigger events are included in the first IU message.

#### Related topics

[PPPoE](#)

[IPoE](#)

[Message retransmission and buffering](#)

[BNG charging profiles](#)

## 5.4.3 Message retransmission and buffering

*RADIUS accounting messages are retried, but can also be buffered for later retransmission. Learn how to enable and configure the buffering of the different RADIUS accounting messages.*

The MAG-c retries RADIUS accounting messages using the configuration of the **acct-retry-count**, **acct-retry-timeout**, and **max-peer-reselections** commands in the **config>mobile>profile>radius** context. If no accounting response is received after the retries, the MAG-c buffers the accounting messages to retransmit them later, if buffering is enabled.

To enable buffering of accounting messages, use the **accounting-buffer** command in the **config>mobile>profile>radius-group** context.

If buffering is enabled, you must not enable support of session deletion when the configurable amount of accounting interim retries is exhausted, that is, do not execute the **delete-session-acct-interim-exh** command in the **config>mobile>pdn>radius** context.

When buffering is enabled, the MAG-c can buffer one Accounting Stop message per session. Optionally one Accounting Start, and up to five Interim Update messages can be buffered.

To enable buffering of the Accounting Interim Update messages, use the **interim-update** command in the **config>mobile>profile>radius-group>accounting-buffer** context.

To enable buffering of one Accounting Start message per session, use the **start** command in the **config>mobile>profile>radius-group>accounting-buffer** context.

The following restrictions apply to buffering of the accounting messages.

- The lifetime of a buffered accounting message is configurable. The default lifetime is 24 hours. Because of final retransmission attempts, the message can be kept longer than the configured lifetime. To configure the lifetime of the buffered accounting messages, use the **lifetime** command in the **config>mobile>profile>radius-group>accounting-buffer** context.
- The system limits the maximum number of buffered accounting messages, and generates a trap (with name `tmnxMobGwAcctBuffResourceProblem`) when the limit is crossed.

The MAG-c classifies the Accounting Interim Update messages as critical or non-critical depending on the trigger event.

- Non-critical messages do not reflect a significant state change and contain data that is present either in the subsequent Accounting Interim Update messages, or in the Accounting Stop message. For example, a periodic Interim Update message contains only updated cumulative counters that are also present in a subsequent Interim Update or Stop message. When buffering of the Interim Update messages is enabled, the following rules apply.
  - Only the last non-critical Interim Update message for a session is buffered. If there was a previous non-critical message buffered, this previous message is discarded and overwritten.
  - When the session terminates, the optionally stored non-critical Intermediate Update message for that session is discarded and overwritten with the Stop message.
- Critical messages reflect a significant state change, and can contain data that is lost if not sent; for example, a stop of service and the final statistics related to that service. When buffering of the Interim Update messages is enabled, up to four critical Interim Update messages per session are buffered to prevent loss of data.

Periodic Interim Update messages are non-critical messages. Triggered Interim Update messages can be critical or non-critical depending on the trigger reason.

A non-configurable timer triggers a periodical retransmit of the buffered messages. The timer value changes depending on the load of the system. It uses exponential back-off when a server is unavailable and can increase to 1 hour.

#### **Related topics**

[Session accounting](#)



## 6 Residential NAT

Get an explanation of NAT on BNG CUPS, learn about the functional split between MAG-c and BNG-UP, NAT configuration, logging, and operational commands.

### 6.1 NAT terminology and references

<b>private side or inside NAT</b>	The terms private side or inside NAT are interchangeable in the context of NAT. They both refer to the side of NAT where the device being translated resides, before translation takes place. The source IP address and protocol port of the devices on the inside are translated to a global IP address and protocol port on the outside. In the scope of this topic, the term inside is used.
<b>public side or outside NAT</b>	The terms public side or outside NAT are interchangeable in the context of NAT. They both refer to the side of NAT after the translation takes place. On the outside, the devices are represented by their translated IP addresses and protocol ports. In the scope of this topic, the term outside is used.
<b>NAT pool</b>	A NAT pool is a collection of outside prefixes attached to an outside realm and shared by a group of subscribers. The NAT type (1:1, NAPT) is a property of a NAT pool. Multiple NAT pools can be associated with an outside realm.
<b>NAT flow</b>	<p>NAT flows result from translations for which states are maintained in NAT. The following fields represent a NAT flow:</p> <ul style="list-style-type: none"> <li>• source IP address</li> <li>• source port</li> <li>• translated IP address</li> <li>• translated port</li> <li>• destination port</li> <li>• destination IP address</li> <li>• protocol</li> </ul> <p>The NAT CLI and standard documents sometimes refer to NAT flows as sessions. However, a NAT flow must not be confused with a BNG CUPS session. In this topic, NAT flows are referred to as flows.</p>
<b>NAT and NAT44</b>	The terms NAT and NAT44 are used interchangeably in this topic.

The following guides are related references:

- *7450 ESS, 7750 SR, and VSR Multiservice ISA and ESA Guide*  
This guide describes many concepts of residential NAT on BNG CUPS that are borrowed from L2-aware NAT on SR OS.
- *7750 SR and VSR BNG CUPS User Plane Function Guide*  
This guide describes the supported NAT functionality on the BNG-UP.

- *7450 ESS, 7750 SR, and VSR RADIUS Attributes Reference Guide*  
This guide describes the supported RADIUS attributes for SR and VSR.
- *MAG-c RADIUS Attributes and IU Triggers*  
This guide describes the supported RADIUS attributes for MAG-c.

### Related topics

[Session management](#)

## 6.2 Residential NAT44 on BNG CUPS

Traditionally, a NAT binding represents a mapping between an IP address with all of its protocol ports on the inside and an IP address with a specific protocol port range on the outside. This allows sharing of a single IP address on the outside by multiple devices on the inside.

The traditional NAT concept can be extended to a residence or a home, where a NAT binding can represent a mapping between a subscriber residence (or home) and an outside IP address with a specific port block. Regardless of whether the residence is bridged with the devices' IP addresses exposed or routed with a single IP address, an entire residence can be mapped to a single outside IP address and a number of port blocks. The advantage of such aggregation of devices in bridged home environments helps to conserve NAT resources. This type of NAT is on BNG CUPS referred to as residential NAT.

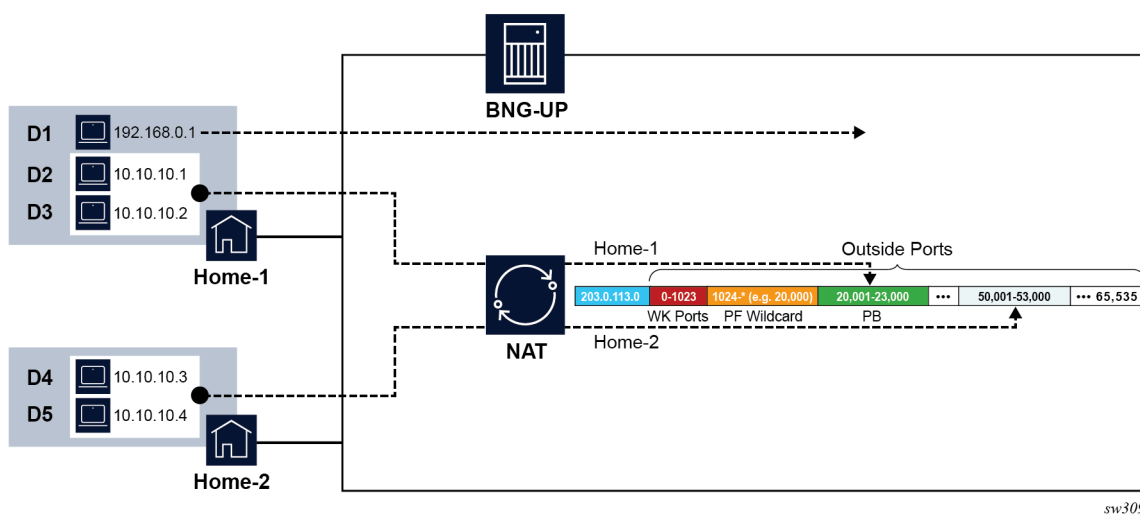
A subscriber can have a mix of sessions that are going through NAT processing and sessions that bypass NAT processing. The term "NAT enabled subscriber" used throughout this document refers to a subscriber that has all or some sessions going through NAT processing.



**Note:** On the integrated BNG in SR OS, this type of NAT is referred to as L2-aware NAT.

The figure shows a residential NAT example with two residences.

Figure 14: Residential NAT example



- In the first residence (Home-1), one device (D1) bypasses NAT and two devices (D2 and D3) go through NAT.

- In the second residence (Home-2), both devices (D4 and D5) go through NAT.

Instead of allocating 4 port blocks, one for each device going through NAT, residential NAT allocates only two port blocks, one per residence.

WK ports are the well-known ports such as HTML and SMTP. The PF wildcard range represents the static Port Forward range.

## 6.3 Functional split between MAG-c and BNG-UP

The distribution of NAT related functionality between the MAG-c and the BNG-UP on a BNG CUPS system is as follows.

### MAG-c

- During the session authentication phase, the MAG-c determines if a session needs to be associated with NAT.
- If the session is to be associated with NAT, the MAG-c selects the NAT outside prefix (in a NAT pool), an outside IP address and the initial (or the first) port block (in NAPT), and the NAT related policy which determines NAT operational parameters (ALG, protocol timers, and so on).
- The MAG-c logs the NAT resources (outside IP address and port block) via a CUPS session account (RADIUS based). In this way, NAT logging becomes an integral part of the session accounting.
- The MAG-c submits the selected NAT resources to the BNG-UP.

### BNG-UP

- Based on the received NAT parameters for the session, the BNG-UP creates a binding and performs NAT translations for data traffic without any further help from the MAG-c.
- BNG-UP can allocate additional extended port blocks for a subscriber and notify MAG-c about port block allocations and deallocations to properly integrate them in the subscriber management logging and accounting on the MAG-c.
- Optionally, flow-based logging can be enabled on the BNG-UP.

Maintaining the management of outside IP addressing on the MAG-c works in favor of multi-chassis redundancy, where existing outside IP addresses and port blocks can be preserved between switchovers.



**Note:** Although residential NAT is tightly coupled with subscriber management, it is not the only mode of operation on BNG CUPS. An alternative to residential NAT is to enable NAT only on the BNG-UP. In such mode, the BNG-UP performs traditional NAT (CGN) independent of the MAG-c, where bindings are created per device (not per residence) and logging is performed by the BNG-UP. In other words, CGN is decoupled from subscriber management and works as an independent function on the BNG-UP. For more information about this MAG-c independent version of NAT, see *7450 ESS*, *7750 SR*, and *VSR Multiservice ISA and ESA Guide*.

## 6.4 Management of NAT outside prefixes

Outside NAT prefixes are allocated on demand. This concept follows the On Demand Subnet Allocation (ODSA) approach, in which smaller subnets (or prefixes) are allocated from a pre-configured prefix (supernet). For more information about NAT prefix allocation, see [ODSA and local address assignment](#).

## 6.5 CP NAT profile

The CP NAT profile defines a set of NAT parameters related to the outside addressing and the type of NAT. The MAG-c uses the parameters in the CP NAT profile, with exception of the **up-nat-policy** parameter which is a reference to a NAT policy on the BNG-UP. When the MAG-c does not provide a UP NAT policy, the system uses the UP NAT policy with name default configured on the BNG-UP.

Use the **cp-nat-profile** command in the **config>mobile>profile>bng** context to configure a CP NAT profile.

The following are parameters provided via the NAT pool in the CP NAT profile:

- **laa-pool network-realm**

The network realm defines the outside routing context (VPRN, Base).

- **laa-pool name**

The name points to the ODSA pool name from which the NAT prefix is allocated. The ODSA pool may contain multiple address ranges (supernets).

- **mode**

The mode defines the NAT mode (1:1 or NAPT) which is necessary to properly allocate outside IP addresses, port blocks, number of subscriber per outside IP address and static port range from the pool.

- **up-nat-policy**

The UP NAT policy is a reference to the NAT policy that resides in the BNG-UP. This is an optional parameter. When the MAG-c does not provide a UP NAT policy, the system uses the UP NAT policy with name default configured on the BNG-UP.

The inside network realm (or routing context) is determined from the APN. For more information, see [Service selection](#).

The following are characteristics of a CP NAT profile.

- A CP NAT profile is associated with a IPoE or PPPoE session during the session authentication phase. If a session is not associated with a CP NAT profile during the authentication phase, NAT is not performed for that session, and the traffic of that session bypasses NAT.
- In residential NAT, all NAT enabled sessions of a specific subscriber must share the same CP NAT profile. The session setup fails for a session that is associated with a different CP NAT profile than the profile that is already assigned to existing sessions of the same subscriber.
- A CP NAT profile cannot be removed from a session via CoA.
- A CP NAT profile cannot be added to a session that was instantiated without NAT.

## 6.6 Port forwards

Port forwards are a session concept that allows devices on the outside to initiate traffic toward a configured port on the inside through an open NAT pinhole (a fixed mapping between an inside and an outside port). Port forwards can be allocated dynamically via UPnP or statically via RADIUS Access-Accept or CoA messages.

The UPnP policy is configured on the BNG-UP. A UPnP request from the client is forwarded to and served by the ISA or ESA. In UPnP, ports are allocated from the port-block that is allocated to the subscriber, not from the wildcard port forwarding range.

Static port forward requests sent via RADIUS CoA can be addressed to a session or a subscriber. In case of a subscriber, the port forward is accepted only if the subscriber has a single NAT enabled session.

The Alc-Static-Port-Forward VSA is used for allocation or deletion of static port forwards. For more information about the VSA, see the *CMG BNG CUPS RADIUS Attributes*.

## 6.7 Extended port blocks

### Multiple port blocks per subscriber

Residential NAT supports allocations of multiple port blocks (PBs) for each subscriber, or more accurately for a set of NAT-enabled sessions within a subscriber. The PB space of an outside IP address in a NAT pool is divided into two partitions. The first partition is reserved for the first (or initial) PB of a subscriber. The second partition is dedicated to the extended PBs, which are allocated dynamically on an as-needed basis in case a subscriber needs more ports. The two occupy the port space of an IP address consecutively, where the second port partition extends from the end of the first partition to the end of the port space of an IP address (port 65535).

Although the NAT resource are allocated and deallocated in the BNG-UP, the MAG-c, controls the allocation of the outside IP addresses, the first PBs, and the division of the port space in the BNG-UP. The BNG-UP controls the allocation of the extended PBs for each subscriber.

The BNG-UP notifies the MAG-c, of the allocation and deallocation of the extended PBs. In this way, logging of extended PBs is integrated into the accounting logic on the MAG-c, where the newly allocated and deallocated PBs are reported in triggered RADIUS Interim-Update messages.

### 6.7.1 Port space division

MAG-c programs the BNG-UP with the first port of the PB space used for extended PB allocations. This first port divides the port space of an outside IP address into two. The first part is reserved for well-known (WK) ports, port forwards and the ports reserved for the initial port blocks of each subscriber. This space is managed by the MAG-c. The second partition that follows the first partition to the end of the entire port space (port 65,535) is reserved for the extended PBs. This port range is managed by BNG-UP.

The following three configuration options determine the first port of the partition that is used for extended PBs on the MAG-c:

- maximum number of subscribers per outside IP address (**subscriber-limit** command)
- size of the first PB for each NAT subscriber (**port-reservation port** command)
- last port of shared port forwarding range (**port-forwarding-range** command)

All three parameters are configured in the CP NAT profile on the MAG-c:

```
configure mobile-gateway profile bng cp-nat-profile nat-pool laa-pool mode
```

Configuring the **subscriber-limit** command enables allocation of extended PBs. The extended PB port partition starts at the port determined by the follow formula:

subscriber limit per outside IP address [subscriber-limit] \* size of the first PB [port-reservation ports] + port forwarding range end [port-forwarding-range] + 1

While these parameters are configured on the MAG-c, the size of the extended PBs and the maximum number of PBs per subscriber are configured in the UP NAT policy on the BNG-UP; see the *7750 SR and VSR BNG CUPS User Plane Function Guide*, "Guidelines for configuring extended port blocks".

## 6.7.2 Managing port block space

Both the initial and extended port partitions are served on a first-come, first-serve basis. The initial port partition guarantees at least one port block (PB) for each of the preconfigured number of subscribers per outside IP address (subscriber-limit in the pool). If there are more subscribers in the network than the preconfigured number of NAT subscribers, this space becomes oversubscribed.

The extended port partition does not guarantee that each of the existing NAT subscribers receive additional PBs. Each subscriber can allocate additional free PBs only if they are available, up to the maximum combined limit (initial and extended) set in the UP NAT policy (**block-limit** parameter) configured on the BNG-UP.

For optimized NAT pool management and correct capacity planning, it is essential to understand the following configuration elements in the user's network, which determine the average PBs per subscriber:

- IP address compression ratio – how many subscribers share one outside IP address
- subscriber over subscription ratio – how many NAT subscribers are active simultaneously
- statistical port usage for subscribers – what percentage of subscribers are heavy, medium, and light port users
- PB sizes

After the average PBs per subscriber is determined, the following NAT parameters can be configured:

- the subscriber-limit per outside IP address in the CP NAT profile on MAG-c
- the size of the initial PB in the CP NAT profile on MAG-c
- the size of the extended PB in the UP NAT policy on the BNG-UP
- the maximum number of PBs per subscriber in the UP NAT policy on the BNG-UP
- the outside IP address range as part of the NAT prefix in the ODSA pool in MAG-c

### Guidelines for determining traffic patterns and port usage

The following guidelines and examples can serve as an initial configuration for administrators who are unsure of their traffic patterns in terms of port usage for their subscribers. The calculations are based on the following assumptions:

- There are 10,000 subscribers that require NAT, however only 8,000 of them are active simultaneously. This means that over subscription of outside (NAT) IP address is allowed.
- The subscriber's port usage is on average:
  - 60% light users with less than 1000 ports
  - 30% medium users with less than 2000 ports
  - 10% heavy users with less than 4000 ports

The following calculations are based on the stated assumptions:

1. There are 12,800,000 ports in total.

$$8,000 \text{ active subscribers} * (0.6 * 1000 + 0.3 * 2,000 + 0.1 * 4,000) = 12,800,000 \text{ ports}$$

2. One outside IP address can accommodate approximately 50,000 (64K ports less the static port forwards and well known ports), which yields 256 outside IP addresses (/24) in a pool.

$$12,800,000 / 50,000 = 256$$

3. Based on the compression ratio that follows from the preceding calculations, the subscriber limit is 32 (32 subscribers share one outside IP address).

$$8,000/256 = \sim 32$$

4. Based on the calculations, a reasonable size for the initial port block is 1000 ports and for the extended port block is 335 ports.
5. To accommodate heavy users with 4,000 ports, the maximum number of port blocks per subscriber is set to 10.

$$(1*1000 + 9*335 = 4015)$$

Based on the calculations, and assuming the subscribers are well load-balanced over ISAs or ESAs, configure the following to achieve the required port usage:

- 32 for the subscriber limit in a pool
- 1,000 initial and 335 extended for the PB sizes
- 10 for the PBs maximum per subscriber
- /24 address range in the pool

The following examples show the provisioning for the MAG-c and BNG-UP.

### Example: MAG-c BNG profile configuration

```
A:MAG-c>config>mobile>profile>bng# info
-----
      cp-nat-profile "demo-profile"
      nat-pool "demo-pool"
      laa-pool network-realm "demo-realm" name "laa-pool-1"
      mode napt
      port-reservation ports 1000
      port-forwarding-range 15000
      subscriber-limit 32
      exit
      exit
      exit
-----
```

### Example: MAG-c PDN configuration

```
A:MAG-c>config>mobile>pdn# info
-----
      local-address-assignment network-realm "demo"
      pool "laa-pool-1"
      dedicated
      ipv4
      prefix 10.10.10.0/24
      exit
      exit
-----
```

### Example: BNG-UP NAT policy configuration

```

A:node-2>config>service>nat>up-nat-policy# info
-----
      block-limit 10
      port-block-extensions
        ports 335
      exit
-----

```

See the 7750 SR and VSR BNG CUPS User Plane Function Guide, "Guidelines for configuring extended port blocks", for information about PB configuration in the UP NAT policy on the BNG-UP.

## 6.8 NAT logging

Residential NAT on BNG CUPS supports the following logging methods:

- Outside IP address, port-blocks, and realm via RADIUS logging, which is integrated with the BNG CUPS session accounting on the MAG-c
- IPFIX based flow logging on the BNG-UP
- Outside IP address, port-block, and realm via SYSLOG on the BNG-UP

For description of the principles of IPFIX logging, see *7450 ESS, 7750 SR, and VSR Multiservice ISA and ESA Guide* and *7750 SR and VSR BNG CUPS User Plane Function Guide*.

RADIUS based logging for residential NAT on BNG CUPS is integrated in the subscriber accounting. The logging uses the same RADIUS infrastructure for NAT as for subscriber management.

The relevant VSA used for NAT logging is Alc-Nat-Port-Range.

For a description of the VSA, see *CMG BNG CUPS RADIUS Attributes*.

The table describes the relation between the accounting message type, the NAT events, and the content of the Alc-Nat-Port-Range VSA.

*Table 4: RADIUS based logging for residential NAT*

Accounting Message Type	NAT event	Alc-Nat-Port-Range VSA
Start	Initial IP and PB creation	Includes info about the initial allocation
Periodic Interim Update	Resources in use	Includes info about the in use NAT resources
Triggered Interim Update	Allocated or de-allocated extended PBs Timestamp precision is 1 second Only deltas are reported	Includes information about the in use extended PB
Stop	Session closed, all PBs freed	Includes info about the released NAT resources



## 6.8.1 RADIUS-based logging

RADIUS-based logging for residential NAT on BNG-UP is integrated in the subscriber accounting. The logging uses the same RADIUS infrastructure for NAT as for subscriber management. The relevant VSAs used for NAT logging are:

- Alc-Nat-Port-Range
- Alc-ISA-Event-Timestamp
- Alc-Acct-Triggered-Reason
  - NAT-MAP
  - NAT-FREE

For a description of the RADIUS accounting attributes, see *MAG-c RADIUS Attributes and IU Triggers*.

The accounting START message carries the RADIUS Event-Timestamp (type 55) attribute, which correctly reflects the creation of the initial port block (PB) and outside IP address. The initial PB and outside IP address allocation is triggered by the MAG-c at the time when the first session is created. In other words, the initial PB and outside IP address creation in the ISA or ESA is not triggered by data traffic. However, the allocation of extended (non-initial) PBs is triggered by data traffic on BNG-UP.

The Interim-Updates and STOP accounting message carry the following two timestamps:

- The RADIUS Event-Timestamp with a 1second resolution is updated by the MAG-c to reflect the time when the Interim-Update message is generated on the MAG-c.
- The Nokia Alc-ISA-Event-Timestamp is updated only when an event on the ISA or ESA occurs; for example, an extension PB is allocated or deallocated. This timestamp has the same format and resolution as the Event-Timestamp.

The following table describes integrated subscriber management and RADIUS logging attributes relevant to NAT.

Table 5: Integrated subscriber management and NAT RADIUS accounting

Subscriber management and NAT integrated RADIUS accounting and logging			
Acct msg type	Subscriber accounting	Session accounting	Comments
START	<p>The accounting START message is generated for each subscriber, when the subscriber's first session is instantiated.</p> <p>The message carries the NAT-related information (outside IP address and the initial NAT PB), if the subscriber's first session is NAT enabled (associated with the CP NAT profile during authentication).</p> <p>NAT-related information is carried in the following VSA:</p> <p>Alc-Nat-Port-Range (26.6527.121)</p>	<p>The accounting START message is generated for every new session of a subscriber.</p> <p>For NAT-enabled sessions, the message carries:</p> <ul style="list-style-type: none"> <li>• the outside IP address and initial port for the subscriber's first NAT-enabled session</li> <li>• the outside IP address, the initial PB, and the extended PBs for any additional NAT-enabled sessions of the subscriber.</li> </ul>	<p>Subscribers are not NAT aware, only sessions are. However, IP address and PBs are allocated per subscriber. In other words, all NAT-enabled sessions within a subscriber share the same outside IP address and PBs.</p>

<b>Subscriber management and NAT integrated RADIUS accounting and logging</b>			
<b>Acct msg type</b>	<b>Subscriber accounting</b>	<b>Session accounting</b>	<b>Comments</b>
	<p>This attribute includes the outside IP address, newly allocated initial PB, outside realm, and NAT pool.</p> <p>If the first session is not NAT enabled, the Alc-Nat-Port-Range VSA is not present in accounting START message.</p>	<p>The NAT-related information is carried in the following RADIUS attribute:</p> <p>Alc-Nat-Port-Range (26.6527.121)</p> <p>This attribute includes the outside IP address, PBs, outside realm, and NAT pool.</p>	.
Regular Interim-Update	<p>In the NAT context, the regular Interim-Update message is used to periodically report the allocated NAT resources (cumulative update) for each subscriber, if the subscriber has at least one NAT enabled session. The following VSA carries the specified NAT-related information:</p> <p>Alc-Nat-Port-Range (26.6527.121)</p> <p>This attribute includes the outside IP address, all existing PBs, outside realm, and NAT pool.</p> <p>Alc-ISA-Event-Timestamp (241.26.6527.86)</p> <p>This attribute includes the time of the last extended PB allocation or deallocation on the ISA or ESA on the BNG-UP.</p> <p>Event-Timestamp (55)</p> <p>This attribute includes the time when the RADIUS message is generated on the MAG-c.</p>	<p>In the NAT context, this message is used to periodically report allocated NAT resources (cumulative update) for each NAT enabled session. NAT-related information is carried in the following VSA:</p> <p>Alc-Nat-Port-Range (26.6527.121)</p> <p>This attribute includes the outside IP address, all existing PBs, outside router ID, and NAT policy.</p> <p>Alc-ISA-Event-Timestamp(241.26.6527.86)</p> <p>This attribute includes the time of the last extended PB allocation or deallocation on the ISA or ESA on the BNG-UP.</p> <p>Event-Timestamp (55)</p> <p>This attribute includes the time when the RADIUS message is generated on the CPM</p> <p>This is repeated for all NAT-enabled sessions or hosts of an ESM subscriber.</p>	—
Triggered Interim-Update	<p>The triggered Interim-Update message carries the following:</p> <ul style="list-style-type: none"> <li>outside IP address and initial PB for the subscriber's first NAT-enabled session, which is not the first session of the subscriber (the initial PB already missed the accounting START message sent when the subscriber's first non-NAT-enabled session was established)</li> </ul>	<p>This message carries differential updates tracking changes for extended PB allocations and deallocations.</p> <p>Alc-Nat-Port-Range (26.6527.121)</p> <p>This attribute includes the IP address, newly allocated or deallocated extended PB, outside realm , and NAT pool.</p> <p>Alc-Acct-Triggered-Reason (26.6527.163)</p>	—

Subscriber management and NAT integrated RADIUS accounting and logging			
Acct msg type	Subscriber accounting	Session accounting	Comments
	<ul style="list-style-type: none"> <li>differential updates tracking changes for extended PBs (momentary allocations/deallocations).</li> <li>initial and extended PBs when the last NAT-enabled session leaves the subscriber, while other non-NAT-enabled session continue to be present.</li> </ul> <p>Alc-Nat-Port-Range (26.6527.121)</p> <p>This attribute includes the outside IP address, newly allocated or deallocated PBs, outside realm, and NAT pool.</p> <p>Alc-Acct-Triggered-Reason (26.6527.163)</p> <ul style="list-style-type: none"> <li>NAT-MAP (20)</li> <li>NAT-FREE (19)</li> </ul> <p>The reason for this message is an extended PB is allocated (MAP) or de-allocated (FREE).</p> <p>Alc-ISA-Event-Timestamp (241.26.6527.86)</p> <p>This attribute includes the time of the extended port-block allocation or deallocation on the ISA or ESA on the BNG-UP.</p> <p>Event-Timestamp (55)</p> <p>This attribute includes the time when the RADIUS message is generated on the MAG-c.</p>	<ul style="list-style-type: none"> <li>NAT-MAP (20)</li> <li>NAT-FREE (19)</li> </ul> <p>This attribute includes the reason for this triggered Interim-Update message, which is an extended PB is allocated (MAP) or de-allocated (FREE).</p> <p>Alc-ISA-Event-Timestamp (241.26.6527.86)</p> <p>This attribute includes the time of the extended port-block allocation or deallocation on the ISA or ESA on the BNG-UP .</p> <p>Event-Timestamp (55)</p> <p>This attribute includes the time when the RADIUS message is generated on the MAG-c.</p> <p>This is repeated for all the subscriber's NAT-enabled sessions. For example, a single extended port-block allocation can trigger multiple triggered Interim-Updates (one for each existing NAT-enabled session).</p> <p>.</p>	
STOP	<p>An accounting STOP message is sent when a subscriber is terminated (the last session associated with the subscriber terminates). If the subscriber's last session is NAT-enabled, the accounting STOP message carries NAT information related to the resources being released when the last session was terminated (initial and extended PBs).</p>	<p>An accounting STOP message is sent when a session of a NAT-enabled subscriber terminates. The message reports the initial end of extended PBs, regardless of whether this NAT-enabled session is last for the subscriber. In other words, it reports all PBs currently in use by any of the subscriber's NAT-enabled sessions and indicates that this particular</p>	<p>Each accounting stream (START, I-U, or STOP) with the same accounting session ID is treated as a separate entity. In the case of session accounting, the streams may contain NAT</p>

Subscriber management and NAT integrated RADIUS accounting and logging			
Acct msg type	Subscriber accounting	Session accounting	Comments
	<p>Alc-Nat-Port-Range (26.6527.121) This attribute includes the outside IP address, initial and extended PBs, outside realm, and NAT pool.</p> <p>Alc-ISA-Event-Timestamp (241.26.6527.86) This attribute includes the time of the last extended port-block allocation or deallocation on the ISA or ESA on the BNG-UP.</p> <p>Event-Timestamp (55) This attribute includes the time when the RADIUS message is generated on the MAG-c.</p> <p>If the terminated NAT enabled session is not the last for the subscriber, the accounting STOP message does not carry NAT-related information because it was already reported in the triggered Interim-Update message when the last NAT-enabled sessions was released.</p>	<p>session is dissociated from any NAT resources.</p> <p>Alc-Nat-Port-Range (26.6527.121) This attribute includes outside IP address, initial and extended PBs, outside realm, and NAT pool.</p> <p>Alc-ISA-Event-Timestamp (241.26.6527.86) This attribute includes the time of the last extended PB allocation or deallocation on the ISA or ESA on the BNG-UP.</p> <p>Event-Timestamp (55) This attribute includes the time when the RADIUS message is generated on the on the MAG-c.</p> <p>If the terminated session is not NAT enabled, the STOP message does not carry any NAT-related information.</p>	<p>information that overlaps other accounting streams of the same subscriber.</p> <p>PB allocations and deallocations for NAT-enabled sessions within a subscriber are accounted for in each accounting stream. That is, if a port-block allocation (through accounting START or MAP I-U) is present in an accounting stream, the deallocation of the same PB is also present in the same stream (through accounting STOP or FREE I-U). Similarly, if an allocation is missing, the deallocation is also missing in the same stream.</p>

### NAT-related attributes for subscriber-based accounting

The following example describes only relevant NAT-related attributes for subscriber accounting.

#### Example

1. The first session of a subscriber is a NAT-enabled session. At the time of session instantiation, the following RADIUS accounting START messages is generated. Outside IP address 192.168.20.2 and initial PB [2001-2004] are allocated at time T1 in MAG-c.

```
Alc-Nat-Port-Range = "192.168.20.2 2001-2024 realm realm-1 nat-pool pool-1"
Event-Timestamp = T1
```

2. Allocation of a new extended PB follows. Differential data is carried in a triggered Interim-Update message.

Only the newly allocated PBs are present in this update with the triggered reason Nat-Map (20). This PB is allocated on the ISA or ESA in BNG-UP at time T2, which may be different than time T3 at which the Interim-Update from MAG-c is sent to the RADIUS server.

```
Alc-Nat-Port-Range = "192.168.20.2 3000-3023 realm realm-1 nat-pool pool-1"
Alc-Acct-Triggered-Reason = Nat-Map (20)
Event-Timestamp = T3
Alc-ISA-Event-Timestamp = T2
```

- The periodic Interim-Update message is triggered at regular intervals to carry cumulative (or absolute) data. This update carries previously allocated PBs, the initial PB, and the extended PB. T4 in the Event-Timestamp reflects the time when the message is generated, while the **Alc-ISA-Event-Timestamp** is unchanged from the previous update because no new event occurred on the ISA or ESA in BNG-UP.

```
Alc-Nat-Port-Range = "192.168.20.2 2001-2024, 3000-3023 realm realm-1 nat-pool pool-1"
Event-Timestamp = T4
Alc-ISA-Event-Timestamp = T2
```

- Deallocation of an existing extended PB follows. Differential data is carried in the triggered Interim-Update message. Only the deallocated PB is present in this update with the triggered reason Nat-Free (19). This PB was deallocated on the ISA or ESA in the BNG-UP at time T5, which may be different than time T6 at which the Interim-Update is sent to the RADIUS server.

```
Alc-Acct-Triggered-Reason = Nat-Free (19)
Alc-Nat-Port-Range = "192.168.20.2 3000-3023 realm realm-1 nat-pool pool-1"
Event-Timestamp = T6
Alc-ISA-Event-Timestamp = T5
```

- At session termination, a RADIUS accounting STOP message with initial PB is generated. This final update for the session carries the initial PB that the session no longer uses. Although this session is terminated, the initial PB may be used by other sessions still present under the same subscriber. T7 in the Event-Timestamp reflects the time when the message is generated, while the Alc-ISA-Event-Timestamp is always the same as in the previous triggered accounting Interim-Update message.

```
Alc-Nat-Port-Range = "192.168.20.2 2001-2024 realm realm-1 nat-pool pool-1"
Event-Timestamp = T7
Alc-ISA-Event-Timestamp = T5
```

### 6.8.1.1 Enabling RADIUS logging on MAG-c

#### Procedure

Use the **nat-port-range** and the **acct-triggered-reason** commands in the following contexts to enable subscriber and session accounting with NAT-related information.

```
configure mobile-gateway charging bng charging radius subscriber include-attribute
configure mobile-gateway charging bng charging radius session include-attribute
```

The **nat-port-range** command enables sending the Alc-Nat-Port-Range and Alc-ISA-Event-Timestamp VSAs for the subscriber and session accounting.

The **acct-triggered-reason** command together with the triggered Interim-Update message conveys information about the event itself .

### 6.8.1.2 Timestamp interpretation

The extended port block functionality uses an additional NAT-related timestamp in the logging framework, in addition to the standard Event-Timestamp that is carried in every RADIUS accounting message. This additional timestamp is introduced in the accounting stream when the first extended port block is allocated for the subscriber, and thereafter it is present in every accounting message in the stream. It represents the time of the most recent extended port-block allocation or deallocation, as recoded by the ISA or ESA in the BNG-UP.

The following are the interpretations of the two timestamps:

- Event-Timestamp (55) – records the time when the accounting message was generated on the MAG-c
- Alc-ISA-Event-Timestamp (241.26.6527.86) – records the time of the most recent NAT-related event (extended port block allocation or deallocation)

#### Example: Timestamp interpretation

As an example, the following periodic Interim-Update message with the specified NAT-related attributes indicates that at time 1000, a subscriber has two port blocks allocated, [2001-2024] and [3000-3023], and the most recent change to extended port blocks is at time 500.

```
Alc-Nat-Port-Range = "192.168.20.2 2001-2024,3000-3023 realm realm-1 nat-pool pool-1"
Event-Timestamp = 1000
Alc-ISA-Event-Timestamp = 500
```

Consider that the following scenario occurs:

- The extended port block [3000-3023] is released a few milliseconds before the previous periodic Interim-Update message is sent.
- The notification from the ISA or ESA on BNG-UP about this event does not reach MAG-c in time to include the event in the periodic Interim-Update message.

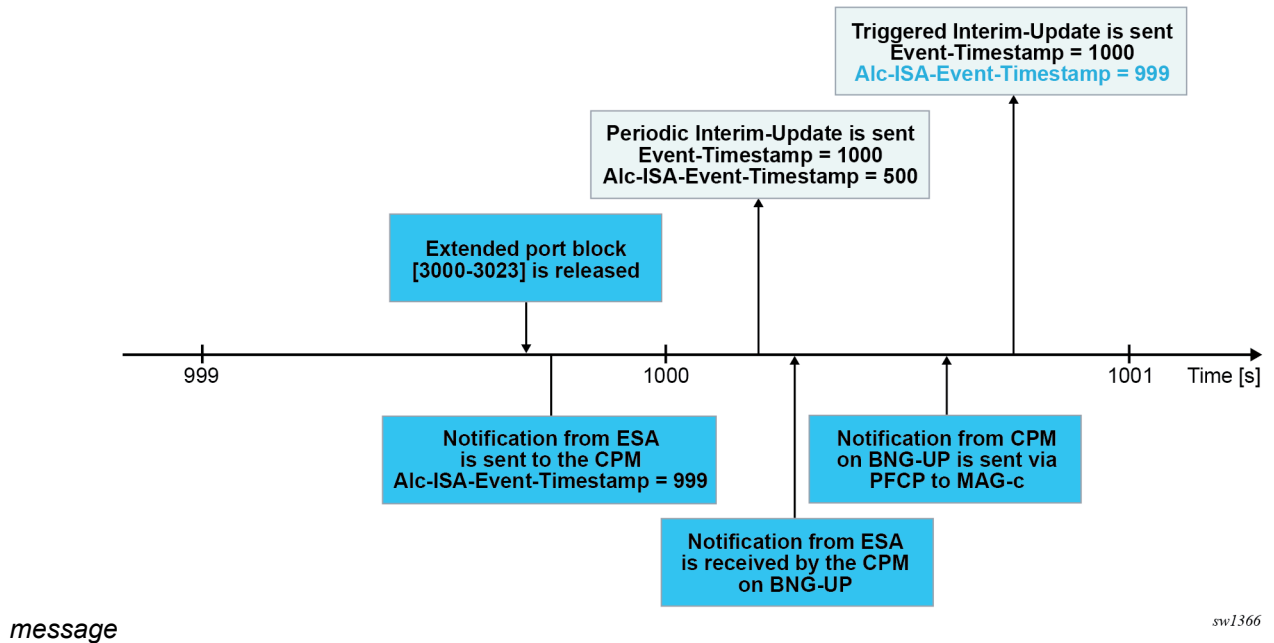
In this scenario, the following triggered Interim-Update message immediately follows the previous periodic Interim-Update message, with the following relevant NAT-related attributes:

```
Alc-Nat-Port-Range = "192.168.20.2 3000-3023 realm realm-1 nat-pool pool-1"
Alc-Acct-Triggered-Reason = Nat-Free
Event-Timestamp = 1000
Alc-ISA-Event-Timestamp = 999
```

Both messages have the same Event-Timestamp of 1000 because the timestamp resolution is 1 second. However, the port block [3000-3023] is released at time 999 indicated by the Alc-ISA-Event-Timestamp triggered Interim-Update message.

The following figure shows this scenario.

Figure 15: AIC-ISA-Event-Timestamp triggered Interim-Update



### 6.8.1.3 High logging rates

A system with on-demand port-block allocation is dynamic and possibly generates a high volume of logs. The transport of NAT logs through RADIUS accounting relies on the generic RADIUS accounting infrastructure implemented in MAG-c, which supports multiple RADIUS servers and failover mechanisms. If the rate of accounting messages exceeds the capacity of the entire accounting system, the queue of accounting message toward the RADIUS servers in MAG-c starts filling up. The cause of this could be an internal condition in the CUPS system or slow or even unresponsive RADIUS servers. Considering that NAT is only a contributor to the accounting messages in the larger accounting framework that includes subscriber management, the rate of the allocation and deallocations of extended PBs is internally limited. Although this does not prevent the loss of accounting messages in an overloaded accounting system (for example, because the RADIUS server is slow), it reduces the possibility that the system becomes overloaded in the first place.

### 6.8.1.4 Buffering during RADIUS failure

MAG-c provides a mechanism whereby the system can buffer accounting/logging packets for longer periods of time while the accounting servers are unreachable. When the server connections recover, the messages from the buffer are transmitted to the servers, preserving the information during the downtime.

This functionality is not supported with logging of extended PBs. The reason for this is the buffering logic overrides the older messages for the same stream and type with the new ones. For example, the current Interim-Update message (for a specific session) that is in the buffer is overridden by the next one. This is acceptable because the periodic Interim-Update messages carry cumulative information (bytes/octets) and consequently the information is preserved in the most recent message. However, this is not the case for Triggered-Interim-Update messages for extended PBs in NAT, where only the new information (allocation

and deallocation) is carried. This means that every message would need to be preserved in the buffer, in which case the higher rate of logs in NAT would overrun the buffer too quickly.

## 6.9 Watermarks

On the MAG-c, a threshold can be configured to monitor the availability of micro-nets. The threshold is set for the minimal number of free micro-nets. When the number of free micro-nets reaches this threshold, a log is generated, alerting the operator about this condition. See [ODSA](#) for more information.

In addition to this threshold on MAG-c level, a number of watermarks can be defined on the BNG-UP level. The BNG-UP reports threshold crossing of the watermarks on the BNG-UP level. See *7750 SR and VSR BNG CUPS User Plane Function Guide* for more information.

## 6.10 Minimum configuration steps

*Learn what minimum configuration residential NAT on MAG-c needs to be operational.*

### About this task

This procedure defines the minimum configuration steps that are necessary to operationalize residential NAT on MAG-c.

### Procedure

**Step 1.** Configure a local address assignment (ODSA) with an outside NAT prefix, so a **cp-nat-profile** can point to it.

To configure a local address assignment pool with an outside NAT prefix, use the **pool** command in the **config>mobile>pdn>laa>network-realm** context

#### Example

```
configure mobile-gateway pdn
  local-address-assignment
    network-realm "realm-1"
    pool "laa-pool-1"
      ipv4
        prefix 198.51.100.0/24
        prefix 198.51.101.0/24
        micro-net-length 28
```

**Step 2.** Configure a **cp-nat-profile** on the MAG-c, so the ADB or RADIUS can point to it.

To configure a CP NAT profile, use the **cp-nat-profile** command in the **config>mobile>profile>bng** context. The minimal configuration of **cp-nat-profile** consists of a NAT pool with a reference to the local address assignment pool (ODSA), the outside realm, the mode of operation, and a reference to the **up-nat-policy**.

#### Example

```
configure mobile-gateway profile bng
  cp-nat-profile "profile-1"
    nat-pool "pool-1"
      laa-pool network-realm "realm-1" name "laa-pool-1"
    mode napt
    port-reservation ports 2000
```



```
up-nat-policy "up-pol-1"
```

**Step 3.** A new NAT enabled session is associated with a **cp-nat-profile** during the authentication phase. Make a reference to the profile locally in the ADB or have it returned from an external AAA server.

To reference the CP NAT profile in the ADB, use the **cp-nat-profile** command in the **config>mobile>profile>adb>entry** context.

A RADIUS server must return the Alc-Cp-Nat-Profile VSA for the session in the Access-Accept message.

**Step 4.** Integrate NAT Logging in the subscriber accounting.

Use the **nat-port-range** command in the **config>mobile>profile>charging>bng>radius>session>include-attribute** context to explicitly enable the Alc-Nat-Port-Range VSA.

#### Example

```
configure mobile-gateway profile charging bng-charging "charging prof" radius session
include-attribute
nat-port-range
```

**Step 5.** Configure the parameters on the BNG-UP.

The following parameters must be configured on the BNG-UP:

- **nat-group** including the ISA redundancy mode
- **up-nat-policy** (When the MAG-c does not provide a UP NAT policy, the system uses the UP NAT policy with name default.)
- **pfcp association** with the **nat-group**

For more information, see *7750 SR and VSR BNG CUPS User Plane Function Guide*.

## 6.11 Operational commands

Get an overview of the CLI commands in the **show** context to obtain NAT information.

To obtain information about the operational NAT state, use the following CLI commands.

- To get information about the session association with the **cp-nat-profile** and the inside routing context, use the **session** command in the **show>mobile>bng** context.
- To get information about the NAT pool ranges and the outside routing context, use the **nat** command in the **show>mobile>bng>session** context.
- To displays information of a specific CP NAT profile, use the **cp-nat-profile** command in the **show>mobile>profile>bng** context.

## 7 Geo-redundancy

Two MAG-c systems can be deployed in a geo-redundant configuration. If one system fails, the other system ensures uninterrupted service and minimizes the impact of failure for broadband clients.

### 7.1 Geo-redundancy overview

Geo-redundancy is the deployment of two MAG-c systems, a primary and secondary system, in a redundant configuration. If the primary system fails, the secondary system continues the service and minimizes the impact of failure for broadband clients.


The session states are synchronized between the two systems so that the secondary system can operate at full state after switchover.

The user configures the administrative primary or secondary role for each system. The operational active and standby roles are determined at run time. The algorithm ensures that only one system has the active role at a specific time.

Both systems get the same service address (for example, the PFCP address) and advertise their route with different metrics. Geo-redundancy is integrated in the routing protocols so that the active system always attracts the CP traffic.

The following figure shows a geo-redundant MAG-c deployment.

Figure 16: Geo-redundant MAG-c deployment



sw1381

## 7.2 Operational and administrative roles

Each system in a geo-redundant deployment has one of the following operational roles at runtime:

- active – attracts and processes the traffic
- standby – takes over from the active system when required

Only one system is active and only the active system processes traffic.

Each system also has one of the following configured administrative roles:

- primary
- secondary

The primary system is preferred over the secondary when the active role is assigned.

The administrative and operational roles represent two different concepts; that is, a primary system can be standby, and a secondary system can be active.

The following events trigger a switchover of the operational roles:

- failure of the active system
- execution of the following command for a manual switchover:

```
admin redundancy mc-mobile-switchover
```

The mc-mobile protocol runs between the two systems to select the active system. The standby system detects a failure of the active system when mc-mobile goes down. Optionally, a BFD session can be bound to mc-mobile to speed up the failure detection.

Use the following command to display the administrative and operational roles:

```
show redundancy multi-chassis mc-mobile peer
```

## 7.3 Traffic detection

A network issue that brings down the mc-mobile protocol triggers an unwanted switchover when the active system has not failed. To avoid the actual switchover in this scenario, the standby system listens for incoming traffic (traffic detection) before switching to the active role.

Use the following command to configure the traffic detection behavior:

```
configure redundancy multi-chassis peer mc-mobile traffic-detection
```

When the traffic detection is set to **relaxed**, the standby system changes to the active role only when it receives a PFCP or IBCP packet.

Use the following command to configure traffic detection on the active system:

```
configure redundancy multi-chassis peer mc-mobile traffic-detection-master
```

When the preceding command is enabled, the active system performs does traffic detection to avoid an active/active scenario when mc-mobile is down.



**Note:** Nokia recommends setting the **traffic-detection** to **relaxed** and the **traffic-detection-master** to **enable**.

If the number of BNG-UPs is small, ensure that the MAG-c receives PFCP or IBCP packets during traffic detection by decreasing the PFCP heartbeat timer on the BNG-UP and increasing the traffic detection time on the MAG-c.

Use the following command to decrease the PFCP heartbeat timer on the BNG-UP:

- MD-CLI

```
configure subscriber-mgmt pfcp association heartbeat
```

- classic CLI

```
configure subscriber-mgmt pfcp-association heartbeat
```

Use the following command to increase the traffic detection timer on the MAG-c:

```
configure redundancy multi-chassis peer mc-mobile traffic-detection-poll-timer
```

## 7.4 State synchronization

The session states are synchronized between the active and standby system to ensure continuity of service after a switchover.

The system can be in one of the following synchronization states:

- hot – all session states are synchronized
- warm – synchronization is ongoing
- cold – no session states are synchronized

Use the following command to configure the session state synchronization:

```
configure redundancy multi-chassis peer mc-mobile mc-complete-ue-sync
```

Use the following command to display the synchronization state:


```
show redundancy multi-chassis mc-mobile
```

## 7.5 Routing

The active system attracts traffic by advertising its corresponding route with a better metric than the standby system. The active and standby systems advertise the same route with different metrics based on their operational state. To achieve routes with different metrics, use the following command to configure a router policy to export the route and configure different entries with different metrics:

```
configure router policy-options policy-statement entry
```

The options for the **from state** command in the preceding context include the following:

- **mobile-master**  
Routes associated with an active system match this entry.
  - **mobile-slave**  
Routes associated with the following systems match this entry:
    - standby system when mc-mobile is down or shunting is down (if shunting is configured) or shunting is not configured (see [Shunting](#))
-  **Note:** When mc-mobile is up and shunting is up, the route on the standby system does not match this entry.
- active system during manual switchover after the active system switches to standby, but before the route is withdrawn from the active system
- **mobile-pre-slave**  
Routes associated with an active system during a manual switchover before the active system switches to standby match this entry.

The configured metrics on the primary and secondary systems for each state must meet the following requirements:

- The metric values for each state must be assigned from best to worst in the following order:
  1. primary active (best metric value)
  2. secondary active
  3. primary standby
  4. secondary standby (worst metric value)
- The **mobile-pre-slave** and the **mobile-master** metric values must be identical.

## 7.6 Shunting

When the standby system receives a packet (for example, before the routing finishes convergence after a switchover), it can forward the packet to the active system. This behavior, which is called shunting, is configurable.

Use the following command to enable shunting:

```
configure redundancy multi-chassis peer mc-mobile mc-redirect
```

Shunting is supported for VPRN over generic routing encapsulation (GRE) service destination point (SDP) for the following types of traffic:

- PFCP
- IBCP
- RADIUS CoA

The standby system performs shunting when a received and resolved multiprotocol BGP (MPBGP) VPN-IPv4 or VPN-IPv6 route matches the local route for supported traffic. The MPBGP route is preferred over the local route.

## 7.7 Manual switchover

Use the following command to execute a manual switchover:

```
admin redundancy mc-mobile-switchover
```

A manual switchover can only be triggered on the active system.

The process of a manual switchover is as follows:

1. A user executes the **mc-mobile-switchover** command on the active system.
2. The active system starts synchronizing its session states with the standby system and changes its state to **mobile-pre-slave**, which triggers a routing metric update.
3. When the synchronization is complete, the active system changes its state to **mobile-slave**, which triggers a next routing metric update.
4. If configured, shunting is enabled on the active system.
5. The standby system becomes the active system and advertises its route with state **mobile-master**.
6. The previously active system (now standby) withdraws its routes.

## 7.8 Deploying and configuring geo-redundancy

### Deployment guidelines

When using geo-redundancy, configure the PFCP path timers so that it takes longer to terminate a PFCP path than to complete a geo-redundancy switchover, including the detection time and the convergence of the routing. Otherwise, the BNG-UPs may terminate a PFCP path and remove all corresponding sessions during a geo-redundancy switchover.

The PFCP headless mode can be used to achieve the above configuration. See [PFCP connectivity failure](#) for more information about the PFCP path timers.

### Configuration commands

The following commands and their leaf commands configure geo-redundancy:

```
configure redundancy multi-chassis peer mc-mobile
configure router policy-options policy-statement
```

### Example: Geo-redundancy configuration on a primary system

```
config>redundancy>multi-chassis# info
-----
peer 46.46.46.46 create
  mc-mobile
  mc-redirect
  mc-complete-ue-sync
  master-traffic-detection enable
  traffic-detection relaxed
  mobile-gateway 1 role primary
```

```

        no shutdown
        exit
    exit
    no shutdown
    exit
-----
config>router>policy-options>policy-statement# info
-----
    entry 10
        from
            prefix-list "prefix-list-1"
            state mobile-slave
        exit
        action accept
            community add "vprn100"
            metric set 30
        exit
    exit
    entry 20
        from
            prefix-list "prefix-list-1"
            state mobile-master
        exit
        action accept
            community add "vprn100"
            metric set 10
        exit
    exit
    entry 30
        from
            prefix-list "prefix-list-1"
            state mobile-pre-slave
        exit
        action accept
            community add "vprn100"
            metric set 10
        exit
    exit
exit

```

### Example: Geo-redundancy configuration on a secondary system

```

config>redundancy>multi-chassis# info
-----
    peer 45.45.45.45 create
        mc-mobile
        mc-redirect
        mc-complete-ue-sync
        master-traffic-detection enable
        traffic-detection relaxed
        mobile-gateway 1 role secondary
        no shutdown
        exit
    exit
    no shutdown
    exit
-----
config>router>policy-options>policy-statement# info
    policy-statement "mcred"
        entry 10
            from
                prefix-list "prefix-list-1"
                state mobile-slave
            exit

```

```
        action accept
            community add "vprn100"
            metric set 40
        exit
    exit
    entry 20
        from
            prefix-list "prefix-list-1"
            state mobile-master
        exit
        action accept
            community add "vprn100"
            metric set 20
        exit
    exit
    entry 30
        from
            prefix-list "prefix-list-1"
            state mobile-pre-slave
        exit
        action accept
            community add "vprn100"
            metric set 20
        exit
    exit
exit
```



## 8 Python support

*A user-defined Python script can customize the MAG-c behavior.*

### Python scripts

Sending or receiving specific control protocol packets can trigger a user-defined Python script. The packet is the input of the script. Using a set of Nokia API calls, the script can inspect and modify the packet. The output of the script is the modified packet.

The direction of the triggering protocol message defines when the Python script runs:

- ingress – before the subscriber management processing
- egress – after the subscriber management processing

For example, when the MAG-c receives a RADIUS Access-Accept message, a user-defined Python script can update the Alc-SLA-Prof-Str attribute in the message to a new SLA profile name. The system processes the modified packet and creates the session with the new SLA profile.

### Python version and libraries

MAG-c Python support is based on MicroPython version 3.4. The SW includes Nokia-provided APIs and the following standard libraries:

- MicroPython libraries
  - sys
  - uarray
  - ubinascii
  - ucollections
  - uhashlib
  - uio
  - ure
  - ustruct
  - utime



**Note:** Nokia modified the implementation of this module. For more information about the use of this module, see *7750 SR pySROS API documentation*.

- standard libraries
  - datetime
  - ipaddress

For more information about the Nokia-provided APIs, see *CUPS BNG TPSDA Python3 API documentation*.

### Supported protocol messages

The following tables list the supported protocol message types and direction.

Table 6: Supported direction for RADIUS messages

Message type	Ingress	Egress
Access-Request	—	✓
Access-Accept	✓	—
Access-Reject	✓	—
Account-Request	—	✓
Account-Response	✓	—
Access-Challenge	✓	—

Table 7: Supported direction for RADIUS CoA messages

Message type	Ingress	Egress
CoA Request	✓	—
DM Request	✓	—
CoA/DM Reply	—	✓

Table 8: Supported direction for PPPoE messages

Message type	Ingress	Egress
PADI	✓	—
PADO	—	✓
PADR	✓	—
PADS	—	✓
PADT	✓	✓
LCP	✓	✓
PAP	✓	✓
CHAP	✓	✓
IPCP	✓	✓
IPv6CP	✓	✓

### Operational commands

To check if a Python script is in service, use the **python-script** command in the **show>python** context.

To enable debugging for Python, enable the **bng** command in the **debug>mobile>call-insight** context and the **python-script** command in the **debug>python>python-script** context.

To bring a modified script in service, reload it using one of the following options:

- the **reload** command in the **tools>perform>python-script** context
- the **shutdown** and **no shutdown** commands for the modified Python script



**Note:** When you modify a signed script, you also need to sign it again using the **protect** command in the **tools>perform>python-script** context.

## 8.1 Configuring a Python script

You can customize the MAG-c behavior with a Python script.

### Procedure

**Step 1.** Create a Python script file. Save the file on local storage or a FTP server.

**Step 2.** Configure the URL of the script file.

Use the **python-script** command in the **config>python** context.



**Note:** Because the MAG-c supports only Python 3, you need to specify **version python3** when you configure the Python script.

**Step 3.** Specify the trigger packet type, the direction, and the corresponding Python script in a Python policy.

Use the **python-policy** command in the **config>python** context.

**Step 4.** Reference the Python policy in the corresponding protocol configuration; for example, inside the RADIUS group for the RADIUS messages.

### Example

The following example configures to run the `cf3:/test.py` Python script file upon sending the RADIUS Access-Request message.

```
config>python# info
-----
python-script "test" create version python3
  primary-url "cf3:/test.py"
  no shutdown
exit
python-policy "test" create
  radius access-request direction egress script "test"
exit
-----
config>mobile>profile>radius-group
-----
server-type both
interface "toRADIUS"
radius-profile "default"
python-policy "test"
peer 172.16.20.100
  secret "KrbVPnF6Dg13PM/biw6ErJsxP6jP" hash2
  no shutdown
exit
supported-features
exit
-----
```

## 8.2 Protecting a Python script file

*You can use a password to protect a Python script file against unauthorized changes. Only a user with the password can load the Python script file.*

### About this task

This procedure provides integrity protection for a Python script. It does not provide confidentiality, that is, a signed file is not encrypted.

### Procedure

**Step 1.** Create a plain Python script file.

**Step 2.** Sign the Python script file using the HMAC-SHA-256 algorithm.

Use the **protect** command in the **tools>perform>python-script** context.



#### Note:

- Remember the chosen password for the following step and to update the script later.
- If you later modify the script, you must sign the updated script again using this command, and you need to reload the script to bring it in service.

**Step 3.** Use the signed script file.

Use the **protection** command in the **config>python>python-script** context with the output file and password of the preceding step.

### Related topics

[Python support](#)

## 9 BNG-UP resiliency

*An overview of the BNG-UP resiliency function and capabilities, resiliency handling, and deployment use cases.*

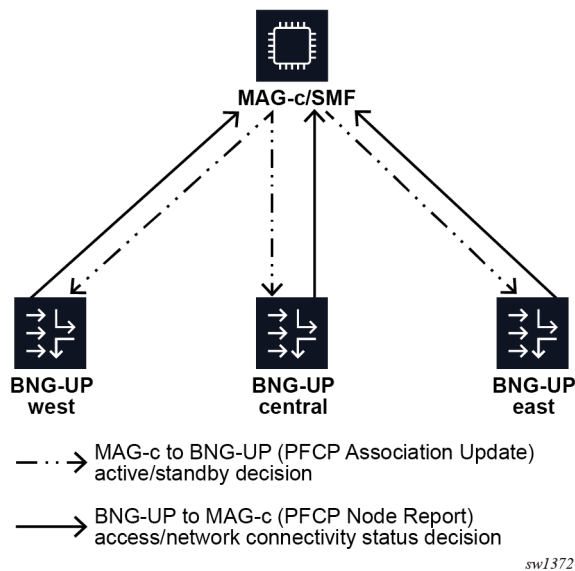
### 9.1 Terminology for BNG-UP resiliency

<b>fate sharing group (FSG)</b>	An FSG is a group of sessions that stay together when moved between BNG-UPs. This guarantees that any associated resources, such as ODSA allocated prefixes, are moved together with the sessions.
<b>active BNG-UP</b>	In the scope of a single FSG, the active BNG-UP is the BNG-UP on which the sessions are created and that actively forwards traffic for those sessions.
<b>standby BNG-UP</b>	In the scope of a single FSG, the standby BNG-UP indicates the BNG-UP that is ready to install sessions and forward traffic upon failure of the active BNG-UP. Whether sessions are proactively created on this BNG-UP depends on the chosen resiliency model.
<b>hot standby</b>	In the hot standby resiliency model, sessions are proactively created on a standby BNG-UP. The standby BNG-UP does not attract traffic but is ready to start forwarding as soon as the MAG-c instructs it to do so.
<b>warm standby</b>	In the warm standby resiliency model, sessions are created solely on the active BNG-UP. Sessions on the standby (new active) BNG-UP are only created after the active BNG-UP fails.

### 9.2 Introduction to MAG-c-driven BNG-UP resiliency


The Nokia MAG-c supports a MAG-c-driven BNG-UP resiliency scheme. In this scheme, the MAG-c selects the active and standby BNG-UPs and the BNG-UPs must follow this decision. The BNG-UPs do not communicate directly to negotiate the active or standby role or to synchronize session state. Instead, each BNG-UP sends its local status indicators to the MAG-c ; for example, whether it has full connectivity to the access network. The MAG-c aggregates these status indicators from all BNG-UPs and makes an informed decision that is sent to the BNG-UPs. The PFCP node messages of the PFCP association between the BNG-UP and MAG-c that are already in place for session management carry the status indicators and informed decisions.

Figure 17: High-level overview of communication for BNG-UP resiliency



It is possible and often wanted that a BNG-UP is active for a subset of the sessions and standby for another subset of the sessions. For example, when two BNG-UPs are fully available, making both BNG-UPs active for half of the sessions and standby for the other half of the sessions may be preferred. Similarly, two Layer 2 access IDs (ports) on the same BNG-UP can be backed up by two different BNG-UPs. The following figure shows the use case where the BNG-UP "central" is backed up by both the BNG-UPs "west" and "east" for two different Layer 2 access IDs.

Figure 18: Multiple backup BNG-UPs



To support all use cases, the MAG-c assigns sessions to a FSG. The MAG-c assigns the active or standby state to each FSG. The state applies to all sessions of the FSG, but not to any other session on the same BNG-UPs. ODSA is also FSG-aware and allocates micro-nets on an FSG basis, instead of a BNG-UP basis, to account for FSGs moving between BNG-UPs.

## 9.3 Modeling a resilient BNG-UP deployment using UP groups

The UP group configuration is a key component of the CUPS BNG-UP resiliency. This configuration serves as a high-level description of the BNG-UP access network so that the MAG-c knows which BNG-UPs are interconnected for BNG-UP resiliency. Based on the UP group configuration, the MAG-c automatically generates FSGs for the resiliency functionality. The UP group contains parameters to create the FSGs.

Use the following command to configure the UP group:

```
configure mobile-gateway pdn bng up-group
```

At the core of the UP group configuration is a list of BNG-UPs. The PFCP Node ID IE as signaled during the PFCP association setup procedure identifies each BNG-UP. The identifier can be either a name or an IP address. The BNG-UPs that form the UP group are interconnected and BNG-UP resiliency can occur between them.

**Note:**

If the identifier of the BNG-UP is an IP address, it must match the PFCP Node ID IE. It does not fall back to the PFCP source IP address.

If a BNG-UP uses a name in the PFCP Node ID IE, the MAG-c must be configured to use the name and not the PFCP source IP address.

**Related topics**

[Fate sharing groups](#)

### 9.3.1 Fate sharing group creation

The MAG-c creates a single FSG per configured UP group. The following configuration for the FSG is provisioned via the UP group:

- **reference to an FSG profile**

Use the following command to configure the FSG profile to reference to:

```
configure mobile-gateway profile bng fsg-profile
```

The profile contains detailed parameters on the resiliency behavior; for example, health calculation for each BNG-UP.

- **preferred indicator**

Per BNG-UP, a flag indicates whether the BNG-UP is active by preference. When the flag is set for a BNG-UP, the FSG prefers this BNG-UP to be active if all other parameters are equal.

- **drain indicator**

Per BNG-UP, a flag indicates whether the BNG-UP is in drain mode. When the flag is set for a BNG-UP, the FSG avoids selecting this BNG-UP as active. For example, this flag can be used before upgrading a BNG-UP to achieve a graceful switchover.



**Note:** Changing the drain flag for an active BNG-UP acts as a BNG-UP reselection trigger for the linked FSGs. The MAG-c moves the sessions after changing the configuration.

## Related topics

[Fate sharing groups](#)


### 9.3.2 Fixed access

Fixed access sessions require the Layer 2 circuit (Layer 2 access ID and VLAN parameters) that is learned from incoming IBCP packets. In a resilient setting, the Layer 2 circuits can differ between the BNG-UPs. For example, in [Figure 18: Multiple backup BNG-UPs](#), Layer 2 access ID "central-A" on BNG-UP "central" is backed up by Layer 2 access ID "west-A" on BNG-UP "west". Because the MAG-c cannot rely on the initial IBCP messages to learn all the Layer 2 access IDs, the IDs must be configured manually.

A single Layer 2 access ID can be configured per BNG-UP in a UP group.. When setting up a new session for this UP group, the MAG-c learns the initial Layer 2 access ID from the incoming IBCP packet, but derives the Layer 2 access IDs for the other BNG-UPs from the configuration. A UP group-level default can be configured to simplify cases where the Layer 2 access IDs are identically named. See [Example for a 1:1 hot standby resiliency with an S-tag per access node](#) for this use case.

When all BNG-UPs use identical Layer 2 access IDs, it is possible to list multiple Layer 2 access IDs per UP group at the group level to avoid creating multiple UP groups for each Layer 2 access ID. When this is configured, the BNG-UP assumes that each Layer 2 access ID is backed up by the identically named Layer 2 access ID on other BNG-UPs. The MAG-c does not assume that there is one big broadcast domain shared between all ports. The following figure shows a UP group that covers two Layer 2 access IDs, named "link-1" and "link2". The sessions on "link-1" cannot be backed up on "link-2" because "link-2" connects to another access node.

Figure 19: Multiple Layer 2 access IDs per UP group



```
up-group "demo"
  fsg-profile "empty"
  l2-access-id link-1 link-2
  up "north"
  exit
  up "south"
  exit
  no shutdown
```



```
exit
```

Similarly, a VLAN range can be configured per BNG-UP for both S-tags and C-tags. A UP group-level default is also available. The VLAN range configuration serves the following purposes:

- Split a single Layer 2 access ID in multiple FSGs and set a different preferred status on different BNG-UPs. In stable conditions, this achieves active-active behavior where some sessions are active on one BNG-UP while others are active on another BNG-UP. See [Example for a 1:1 hot standby resiliency with an S-tag per access node](#) for this use case.
- Set different VLAN ranges on several BNG-UPs in more complex aggregation requirements. The MAG-c automatically adjusts the VLANs learned from IBCP for each UP based on the difference between the start values of the VLAN ranges of each BNG-UP. For example, if UP A is configured with range 100 to 200, and UP B with range 500 to 600, a session with VLAN 150 on UP A automatically uses VLAN 550 on UP B. While the start values of the VLAN range can be different, all ranges must have an equal size. For example, it is not possible to configure a range of 100 to 200 on one BNG-UP, and 100 to 300 on another BNG-UP in the same UP group.



**WARNING:** VLAN ranges with a different offset over more BNG-UPs are an advanced use case and should be carefully validated against the deployed aggregation network. To avoid accidentally enabling different offsets when this functionality is not required, Nokia recommends only configuring a VLAN range on the UP group level.

The following subsections provide deployment use cases and example UP group configurations for the BNG-UP resiliency concepts.

### Example for a 1:1 hot standby resiliency with an S-tag per access node

Four access nodes are connected to a pair of BNG-UPs using a shared broadcast domain. To simplify Layer 2 forwarding, each access node is assigned a unique S-tag. The broadcast domain is connected to each BNG-UP through an identically-named Layer 2 access ID on both BNG-UPs. The MAG-c makes abstraction of whether this connection is a port, LAG, BGP-VPLS, EVPN, or any similar construct.



**Note:** To achieve identical naming on a Nokia BNG-UP, provision an Layer 2 access ID alias using the following command:

- MD-CLI

```
configure service vpls capture-sap pfcpl2-access-id-alias
```

- classic CLI

```
configure service vpls sap pfcpl2-access-id-alias
```

The goal is to have hot standby resiliency, in stable conditions (both BNG-UPs are healthy), such that the active sessions are split between the two BNG-UPs. The following configurations achieve this goal:

- Split the Layer 2 access IDs based on S-tag ranges in two UP groups, each serving half of the access nodes.
- Configure a different BNG-UP as preferred in each group to make the associated FSG active on the preferred BNG-UP as long as that BNG-UP is healthy.


- Link to an empty FSG profile to make the default session standby mode equal to hot (see [Warm and hot standby](#) for more information).

```

up-group "prefer-east"
  fsg-profile "empty"
  s-tag-range start 1 end 2
  l2-access-id to-access
  up "east"
    preferred
  exit
  up "west"
  exit
  no shutdown
up-group "prefer-west"
  fsg-profile "empty"
  s-tag-range start 3 end 4
  l2-access-id to-access
  up "east"
  exit
  up "west"
    preferred
  exit
  no shutdown

```

Figure 20: 1:1 hot standby resiliency example



### Example for a per S-tag 1:1 hot standby resiliency with an S-tag per access node

This example extends the model in [Example for a 1:1 hot standby resiliency with an S-tag per access node](#) with two access nodes and two BNG-UPs.


Instead of splitting the BNG-UPs such that there are two pairs of 1:1 BNG-UPs, each S-tag range gets a different pair of standby BNG-UPs as follows:

- S-tag 1 is backed by BNG-UP "north" and "east"
- S-tag 2 is backed by BNG-UP "east" and "west"
- S-tag 3 is backed by BNG-UP "west" and "south"
- S-tag 4 is backed by BNG-UP "south" and "north"
- S-tag 5 is backed by BNG-UP "north" and "west"

- S-tag 6 is backed by BNG-UP "east" and "south"

```
up-group "s-tag-1"
  fsg-profile "empty"
  s-tag-range start 1 end 1
  l2-access-id to-access
  up "north"
  exit
  up "east"
  exit
  no shutdown
up-group "s-tag-2"
  fsg-profile "empty"
  s-tag-range start 2 end 2
  l2-access-id to-access
  up "east"
  exit
  up "west"
  exit
  no shutdown
up-group "s-tag-3"
  fsg-profile "empty"
  s-tag-range start 3 end 3
  l2-access-id to-access
  up "west"
  exit
  up "south"
  exit
  no shutdown
up-group "s-tag-4"
  fsg-profile "empty"
  s-tag-range start 4 end 4
  l2-access-id to-access
  up "south"
  exit
  up "north"
  exit
  no shutdown
up-group "s-tag-5"
  fsg-profile "empty"
  s-tag-range start 5 end 5
  l2-access-id to-access
  up "north"
  exit
  up "west"
  exit
  no shutdown
up-group "s-tag-6"
  fsg-profile "empty"
  s-tag-range start 6 end 6
  l2-access-id to-access
  up "east"
  exit
  up "south"
  exit
  no shutdown
```

Figure 21: Per S-tag 1:1 hot standby resiliency example



When using default FSGs, the MAG-c load-balances the FSGs and sessions as equal as possible by default:

- Two BNG-UPs have two active FSGs.
- Two BNG-UPs have one active FSG.

To improve the balance, you can add more S-tags or more BNG-UPs or both. For example, using 12 S-tags with a UP group each leads to a balance where each BNG-UP has three active FSGs.

The difference between a BNG-UP-level 1:1 model and an S-tag-level 1:1 model lies in the impact of multiple BNG-UP failures. For example, compare the deployment where "north" and "south" back up each other and "east" and "west" back up each other without overlap. We assume each S-tag range is responsible for about 1/6th of the traffic.

- When two BNG-UPs fail in the per S-tag mode, it always impacts 1/6th of traffic because each pair of BNG-UPs is always uniquely responsible for one S-tag out of six. For example, if "north" and "south" fail, S-tag 4 completely fails.
- When two BNG-UPs fail in the per-BNG-UP mode, the impact depends on which nodes fail and that can either impact 0% or 50% of the traffic. For example, if both "north" and "west" fail, there is no lasting traffic impact because they do not back up each other. If both "south" and "north" fail, all traffic of the two S-tags covered by these BNG-UPs fails.

This effect becomes stronger with more BNG-UPs and S-tags to distribute. For example, in a model with 10 BNG-UPs, the configuration can limit a failure of two BNG-UPs to only affect about 2% of the traffic versus potentially 20% of the traffic if five 1:1 pairs are used.

This model makes the following assumptions on the aggregation model:

- A shared L2 broadcast domain must be available for all BNG-UPs.
- A suitable granularity to differentiate UP groups must be available, such as S-tags in the example above.

- The BNG-UP failures are unrelated. If the BNG-UP failures happen in bulk (for example, because they are co-located), it can be better to make sure no co-located BNG-UPs back up each other instead of to distribute resiliency as much as possible.

**Related topics**

[In-band control plane and BNG-UP selection](#)

[Session keys and anti-spoofing](#)

## 9.4 Fate sharing groups

Fate sharing groups (FSGs) are groups of sessions on which resiliency operations are performed. FSGs are automatically created based on configured UP groups. The FSGs are provisioned via the UP group.

When an FSG is created, the MAG-c performs the following operations:

- Map new sessions to the FSG (see [Session-to-FSG mapping](#)).
- Determine traffic management parameters to attract traffic only to the BNG-UP that serves the specific FSG (see [Traffic steering parameters](#)).
- Determine an aggregated health value for each BNG-UP in the FSG.
- Upon BNG-UP state and health changes, reselect an active and standby BNG-UP for the FSG. Any change triggers this reselection, which guarantees that no state change is lost. In many cases, the MAG-c selects the same active and standby BNG-UP as before.
- Upon any active/standby change, update the FSG state on the BNG-UP and, if necessary, update the session state on the BNG-UP.

FSGs follow an intent-based processing model. The configuration specifies the edge conditions of resiliency behavior, expressing its intent. For example, the configuration specifies whether switchovers should be revertive and whether there is a preferred BNG-UP. The MAG-c monitors multiple parameters and, if necessary, changes active/standby decisions to better match the intent.

**Related topics**

[Fate sharing group creation](#)

### 9.4.1 Session-to-FSG mapping

When setting up a fixed access session, the MAG-c uses the BNG-UP ID, the Layer 2 access ID, and the VLAN ranges of the triggering IBCP packet to look up a UP group. If a UP group contains this set of parameters, the MAG-c links the session automatically to the FSG created for that UP group.

### 9.4.2 Traffic steering parameters

FSGs specify the granularity for the session switchover from one BNG-UP to another. A BNG-UP must uniquely attract traffic for a specific FSG in both the uplink and downlink direction without affecting other FSGs. To achieve this, the MAG-c:

- associates unique uplink and downlink parameters with each FSG

- signals those parameters to the BNG-UP as part of creating the FSG when that BNG-UP is selected as active or standby BNG-UP for that specific FSG

ODSA allocates a unique set of per-FSG subnets (micro-nets). Because the subnets are unique per FSG, the active BNG-UP can announce these subnets. To achieve the uniqueness, a session that is linked to an FSG passes the FSG as an allocation context to ODSA. ODSA automatically makes the micro-nets unique in that context.



**Note:** A standby BNG-UP can also announce the subnet in routing messages but it should make sure that the subnet has lower priority. To achieve this, the standby BNG-UP appropriately sets metrics or preference values in the used routing protocol.

For fixed access sessions, the MAG-c generates a unique MAC address per FSG. When receiving ARP or ND requests in the scope of sessions or subnets linked to a specific FSG, only the active BNG-UP can respond to the requests with the unique MAC address. This makes sure that any MAC forwarding databases in the Layer 2 aggregation point to the correct active gateway. Each time the MAG-c signals a BNG-UP to become active, the BNG-UP can generate GARPs with the unique MAC address to expedite traffic convergence to the new active BNG-UP. The MAG-c bases the generation of the MAC addresses on a /32 prefix configuration. Use the following command to configure the prefix:

```
configure mobile-gateway profile bng fsg-profile mac-prefix
```


The default 02-00-5e-00 prefix is based on the MAC prefix used for VRRP, with the L bit flipped to remove its globally unique significance.

[Example of the relationship between FSGs, MAC addresses, and subnets](#) shows the MAC addressing for 6 FSGs with 2 subnets each, distributed over 3 BNG-UPs. The relationship between the FSGs, MAC addresses, and subnets is as follows:

- FSG 1
  - MAC 02-00-5e-00-00-01
  - session subnet 10.1.1.0/24
  - session subnet 10.1.2.0/24
- FSG 2
  - MAC 02-00-5e-00-00-02
  - session subnet 10.2.1.0/24
  - session subnet 10.2.2.0/24
- FSG 3
  - MAC 02-00-5e-00-00-03
  - session subnet 10.3.1.0/24
  - session subnet 10.3.2.0/24
- FSG 4
  - MAC 02-00-5e-00-00-04
  - session subnet 10.4.1.0/24
  - session subnet 10.4.2.0/24
- FSG 5

- MAC 02-00-5e-00-00-05
- session subnet 10.5.1.0/24
- session subnet 10.5.2.0/24
- FSG 6
- MAC 02-00-5e-00-00-06
- session subnet 10.6.1.0/24
- session subnet 10.6.2.0/24

Figure 23: Example of the relationship between FSGs, MAC addresses, and subnets



## Related topics

[ODSA](#)

### 9.4.3 BNG-UP health determination

The BNG-UP health is the main criterion that the MAG-c uses to determine the active and standby BNG-UP. Health is a value between 0% and 100%; the -1 value indicates BNG-UP unavailability. The following rules determine the BNG-UP health:

- When the PFCP path between the MAG-c and the BNG-UP is down or in headless mode, the health value is -1 (unavailable).



**Note:** If a PFCP association is not set up, the BNG-UP is operationally not part of the UP group and has no health.

- When the following command is set to true, the health value is -1 (unavailable).

```
configure mobile-gateway pdn bng up-group up drain
```

- In all other cases, the health value is based on an aggregation of the operational statuses received from the BNG-UP.

The BNG-UP can signal the following operational status values to the MAG-c:

- **per Layer 2 access ID**  
A percentage value per Layer 2 access ID indicates the current forwarding capacity compared to the full forwarding capacity. For example, if the Layer 2 access ID represents a LAG with five members where one member failed, the expected capacity is 80%.
- **per Layer 3 service (also known as network instance or network realm)**  
A binary connectivity status per Layer 3 service indicates whether the Layer 3 core network is reachable or not (connected or isolated). A Nokia BNG-UP additionally augments this value with a percentage value to cover partial failures. The MAG-c uses the more detailed percentage value if available; otherwise, the MAG-c interprets the binary connectivity status as 100% for the connected state and 0% for the isolated state.

Not all status values of a single BNG-UP apply to a specific FSG. For example, a UP group that only covers a single Layer 2 access ID is not impacted by any other Layer 2 access ID status. The MAG-c determines the applicable status values as follows:

- By default, the MAG-c uses for the aggregation all Layer 2 access IDs configured for the BNG-UP in the UP group. The following commands configure the L2 access IDs:

```
configure mobile-gateway pdn bng up-group l2-access-id
configure mobile-gateway pdn bng up-group up l2-access-id
```

- The MAG-c can exclude configured Layer 2 access IDs from the health calculation. This prevents the MAG-c from automatically setting the health value to 0 if the BNG-UP does not or cannot provide a status value for Layer 2 access IDs. The following command specifies whether to include L2 access IDs and is enabled by default:

```
configure mobile-gateway profile bng fsg-profile health-calculation include-l2-access-ids
```

- The MAG-c tracks a list of configured network realms for health aggregation. The following command configures the tracked network realms:

```
configure mobile-gateway profile bng fsg-profile health-calculation network-realm
```

To calculate a single health value from the set of status values, the MAG-c applies an aggregation calculation that is configured using the following command:

```
configure mobile-gateway profile bng fsg-profile health-calculation aggregation-mode
```

The options for the aggregation mode are:

- **lowest**  
This mode sets the per-BNG-UP health to the lowest value of any Layer 2 access ID and network realm value. A single failure aggressively decreases the health.
- **average**  
This option sets the per-BNG-UP health to the arithmetic mean of all Layer 2 access ID and network realm values. A single failure less aggressively impacts the health.

If the BNG-UP does not signal a status value for a Layer 2 access ID or network realm that is configured to be tracked, the MAG-c sets the status value for the respective Layer 2 access ID or network realm to 0%.



Because the MAG-c uses those values in the aggregation calculation, any missing status value sets the BNG-UP health to 0% for an aggregation mode that is equal to **lowest**.

Next to the BNG-UP health ranging from 0% to 100%, the MAG-c maintains a simplified BNG-UP failure state. A BNG-UP is considered failed if its health is below the failure threshold. To configure the failure threshold, use the following command:

```
configure mobile-gateway profile bng fsg-profile health-calculation failure-threshold
```

By default, the failure threshold is set to 1% , meaning that only a BNG-UP with a health value equal to 0% or -1 (unavailable) is considered failed.

The MAG-c maintains a special not-ready indicator for the current standby BNG-UP. This indicator is set in the following conditions:

- The BNG-UP changes to standby, independent of its previous state or health.
- The BNG-UP health becomes unavailable (-1).

The MAG-c removes the not-ready indicator each time an FSG change successfully completes (see [Active/standby change or switchover](#)) and the health of the BNG-UP at that time is 0% or higher.

The MAG-c avoids making a standby BNG-UP with the not-ready indicator active unless it has no other choice; for example, when the PFCP association for the active BNG-UP is released. This mechanism gives a failed or new standby BNG-UP a chance to go through one FSG change sequence to reinstall all the hot standby sessions before it can be made active.

The MAG-c can put a BNG-UP in a lockout state for an FSG. When a BNG-UP is in the lockout state, it cannot be made active or standby. Contrary to the other health values, the lockout state is intended to recover from hard failures where it is important that all FSG and related session state is removed from the BNG-UP before it is considered active or standby again. See [UP Lockout](#) for more information.

[Table 1](#) provides an overview of the states that are kept for BNG-UPs that have an active association and that are linked to at least one FSG.

*Table 9: Summary of BNG-UP states*

State	Description	Sources
health	A value between 0% and 100% or the special value -1 (unavailable)  Indicates the health of the BNG-UP	Aggregation of the per-logical-port and per-network-realm health reports from the BNG-UP  PFCP path management state (for example, headless)  The drain mode configured with the following command:  <pre>configure mobile-gateway pdn bng up-group up drain</pre>
failed indicator	An indicator that considers the BNG-UP failed if its health is less than the failure threshold	Based on the health state and the threshold configured with the following command:  <pre>configure mobile-gateway profile bng fsg-profile</pre>

State	Description	Sources
	Enables switchovers in more restrictive (for example, non-revertive) scenarios	health-calculation failure-threshold
not-ready indicator	An indicator on the standby BNG-UP that does not have all hot standby sessions installed  Kept until the standby BNG-UP has installed the hot standby sessions	Set for each new standby BNG-UP or a standby BNG-UP whose health becomes unavailable (-1)  Removed after the first successful FSG change when the health is 0% or higher
lockout	A failure state in which the BNG-UP cannot be made active or standby  Kept until the BNG-UP is no longer active or standby and a lockout timer has expired	Applied automatically for multiple failure scenarios, see <a href="#">UP Lockout</a> for more information.

#### 9.4.4 Active/standby selection triggers

The MAG-c monitors multiple triggers that can impact the active/standby selection and trigger a potential switchover. Most events are classified as one of the following:

- recovery (for example, health up)
- degradation (for example, health down)

When a trigger occurs, the MAG-c:

- starts a hold timer
- waits for the hold timer expiry
- triggers the active/standby selection

A different hold time can be set for recovery and degradation using the following commands respectively:

```
configure mobile-gateway profile bng fsg-profile active-standby-selection hold-off-on-recovery
configure mobile-gateway profile bng fsg-profile active-standby-selection hold-off-on-degradation
```

By default, the degradation hold timer is disabled (0 ms) to immediately execute potential switchovers because of failure.

When a trigger occurs while the hold timer is running, the new hold timer is only applied if it is shorter than the one already running. For example, suppose the following events occur with 2 s in between:

- A health increase triggers a recovery hold timer of 5 s.
- A health decrease triggers the default degradation hold timer of 0 ms.

Because the second hold timer is shorter than the first one, the MAG-c immediately triggers the active/standby selection for the degradation.

When a trigger occurs while an active/standby change is in progress, the MAG-c ignores the hold timer of the new trigger and re-evaluates the active/standby selection as soon as the in-progress change completes.

The MAG-c treats the following events as a trigger:

- Any health increase acts as a recovery trigger. The cause of the health increase is irrelevant and may be because of headless recovery, change of the **drain** configuration of the BNG-UP, or a BNG-UP health report.
- Any health decrease acts as a degradation trigger.
- A PFCP association setup acts as a recovery trigger, except if it is the first BNG-UP set up for the FSG.
- A PFCP association release acts as a degradation trigger, except if it is already the active or standby BNG-UP.
- A UP lockout acts as a degradation trigger.
- A UP lockout removal acts as a recovery trigger.
- The intended FSG state not matching the current FSG state after an FSG event acts as a recovery trigger (see [Active/standby change or switchover](#)).

The following exceptional triggers bypass the normal reselection mechanism because of their big impact:

- The setup of the first PFCP association for an FSG triggers an immediate reselection. The MAG-c does not wait for the expiry of the recovery hold timer. If the PFCP association being set up is not the first association, it acts as a health increase and the MAG-c starts the recovery hold timer.
- A PFCP association release for the active or standby BNG-UP triggers an immediate reselection, bypassing any hold timers. If an active/standby change is already in progress, the ongoing change is completed first. A PFCP association release for any other BNG-UP acts as a health decrease and the MAG-c starts the degradation hold timer.
- If all BNG-UPs become headless, the MAG-c does not trigger any reselection. As soon as the first BNG-UP recovers from headless, the MAG-c ignores the recovery hold timer but starts a timer based on the configured path-management heartbeat intervals. The MAG-c triggers reselection of all BNG-UPs when one of the following occurs:
  - the timer based on the configured path-management heartbeat intervals expires
  - 5 s have passed after the last BNG-UP recovered



**Note:** This mechanism ensures that after a full connectivity failure, all BNG-UPs have time to recover the PFCP communication. It makes sure that the MAG-c makes decisions based on the full set of recovered BNG-UPs and not on the first recovered BNG-UPs.

### 9.4.5 Active/standby selection

When an active/standby selection trigger occurs, the MAG-c re-evaluates the selection of the active and standby BNG-UPs for an FSG. If only one BNG-UP with an active association is available, that specific BNG-UP is always selected as the active BNG-UP. Otherwise, both the active and standby BNG-UP can be reselected.

Replacing the active BNG-UP with the current standby BNG-UP works in one of the following basic modes:

- **revertive**

The current standby BNG-UP can be selected as the active BNG-UP even if the active BNG-UP did not fail. The conditions in which the standby BNG-UP can become the active BNG-UP are the same as the conditions to select the standby BNG-UP. Additionally, the standby BNG-UP cannot have the not-ready indicator set.

- **non-revertive**

The current standby BNG-UP can only be selected as the active BNG-UP if the PFCP association of the current active BNG-UP is removed or if the BNG-UP is considered failed (see [BNG-UP health determination](#)), or if the BNG-UP is in lockout state (see [UP Lockout](#)). Otherwise, the current active BNG-UP is always reselected as the active BNG-UP.

To configure the mode, use the following command:

```
configure mobile-gateway profile bng fsg-profile active-standby-selection active-change-without-failure
```

The following command options are available:

- **always**

The MAG-c always uses the revertive mode.

- **never**

The MAG-c always uses the non-revertive mode.

- **initial-only**

The MAG-c uses the revertive behavior for a short period after the first BNG-UP PFCP association for the FSG was set up. After that short period, the MAG-c automatically switches to the non-revertive mode. This option is useful when the non-revertive mode is required but a predictable active/standby BNG-UP is expected during start-up of the BNG-UP and MAG-c; for example, to select the preferred BNG-UP at startup. When the **never** option is set, the first BNG-UP to come up is always selected as active (and that does not change), independent of its preferred state.

If the standby BNG-UP becomes active, the active BNG-UP automatically becomes standby. The MAG-c takes no further action.

The MAG-c selects a standby BNG-UP independent of the revertive mode configuration.

Both the revertive active BNG-UP and the standby BNG-UP are selected using the following criteria. This is a fall-through list that stops as soon as there is only one BNG-UP that meets all the criteria. Any BNG-UP for which the PFCP association is down or which is in lockout is not considered.

1. the BNG-UP with the highest health (see [BNG-UP health determination](#))
2. the preferred BNG-UP
3. the BNG-UP with the lowest number of sessions, simulated as if the FSG would move to that BNG-UP.



**Note:** To avoid unnecessary FSG changes when the number of sessions on several BNG-UPs is very similar, the MAG-c applies a weight multiplier to the FSG session count when it simulates a move to a different BNG-UP than the current one.

4. the BNG-UP with the lowest amount of FSGs, excluding the current FSG, with the goal to provide initial load-balancing when no sessions are set up.
5. The current state of the BNG-UP, where the current active BNG-UP has priority over the current standby BNG-UP which then has priority on any backup BNG-UP. This avoids any unnecessary active or standby changes if all else is equal.

6. the BNG-UP with the lowest IP used in PFCP signaling, with no specific goal other than to have a deterministic tiebreaker when all else is equal

If the result of the active/standby selection differs from the current active/standby selection, the MAG-c initiates an active/standby change.

If the result of the active/standby selection is the same as the current active/standby selection, but the health of any BNG-UP has changed from unavailable (-1) to 0% or higher, the MAG-c initiates an active/standby change.

Otherwise, the MAG-c takes no further action.



**Note:** The trigger to change the FSG for a recovered BNG-UP (even without an active/standby change) is to guarantee that a BNG-UP has all the PFCP state information after a potential communication failure between the BNG-UP and the MAG-c. The FSG change procedure guarantees that all the FSG states and PFCP session states are correctly downloaded if necessary. For example, when a standby BNG-UP becomes headless, it may miss the FSG updates and session installations and modifications for hot standby sessions. When the BNG-UP is recovered from headless, it becomes not ready (see section [BNG-UP health determination](#)). The active/standby state does not change, but the MAG-c triggers an FSG change procedure so that the latest FSG and session state are installed on the BNG-UP. After the FSG change, the MAG-c removes the not-ready indicator from the BNG-UP and the standby BNG-UP is again ready to fully take over.

#### 9.4.6 Active/standby change or switchover

If the active/standby selection results in a new active or new standby BNG-UP, the MAG-c executes the change on the BNG-UPs as follows:

1. The MAG-c updates the PFCP FSG state on all involved BNG-UPs.

The change procedure ends if the active BNG-UP does not positively confirm. If the active BNG-UP change times out or explicitly returns an error, the MAG-c rolls back the changed FSG states and stops the active/standby change procedure.

Changes to other BNG-UPs (for example, standby BNG-UPs) may fail. This is even expected in some cases; for example, in 1:1 deployments where the previously active BNG-UP has failed and becomes standby, the failed BNG-UP is not expected to respond.

A BNG-UP that explicitly rejects an explicit FSG update is put into lockout. This triggers a degradation reselection, which is handled as soon as the change is completed. See [UP Lockout](#) for more information.

2. When the active BNG-UP confirms the FSG change, the MAG-c starts updating the PFCP session states. The exact update for each session depends on the change and the session resiliency model as follows:
  - Warm standby, active/standby switch  
The MAG-c establishes the session on the new active BNG-UP and deletes it from the previous active BNG-UP.
  - Warm standby, new standby BNG-UP  
No updates to the BNG-UPs are needed.
  - Warm standby, health change only  
No updates to the BNG-UPs are needed.

- Hot standby, active/standby switch  
No updates to the BNG-UPs are needed.
  - Hot standby, new standby BNG-UP  
The MAG-c establishes the session on the new standby BNG-UP and deletes it from the previous standby BNG-UP if there was one.
  - Hot standby, health change only  
This acts as a trigger to reinstall missing standby sessions on the standby BNG-UP.
3. When the standby BNG-UP confirms the FSG change, the MAG-c sends a second FSG update message to the active BNG-UP without changing anything. This can be done in parallel with the previous step. The second FSG update message may seem redundant, but is required to resolve a rare race condition in the GARP/ARP signaling for fixed access connections.
  4. When the change is completed, the MAG-c evaluates whether the current active/standby state matches the expected active/standby state by running the selection logic again (see [Active/standby selection](#)). If the states do not match, the MAG-c automatically triggers a recovery reselection and starts the recovery hold timer (see [Active/standby selection triggers](#)).

### GARP/ARP race conditions

Fixed access connections use per-FSG MAC addresses to attract traffic (see [Traffic steering parameters](#)). Most Layer 2 aggregation switches keep a forwarding database (FDB) that points each gateway MAC address to the correct BNG-UP to avoid broadcasting traffic. The FDBs are populated by snooping ARP and ND messages. To expedite updates of the FDBs during active/standby switchovers, the Nokia BNG-UP generates a gratuitous ARP (GARP) message with the FSG MAC address when the FSG is signaled to become active. However, in a very exceptional case, a single GARP is not enough when the following conditions apply:

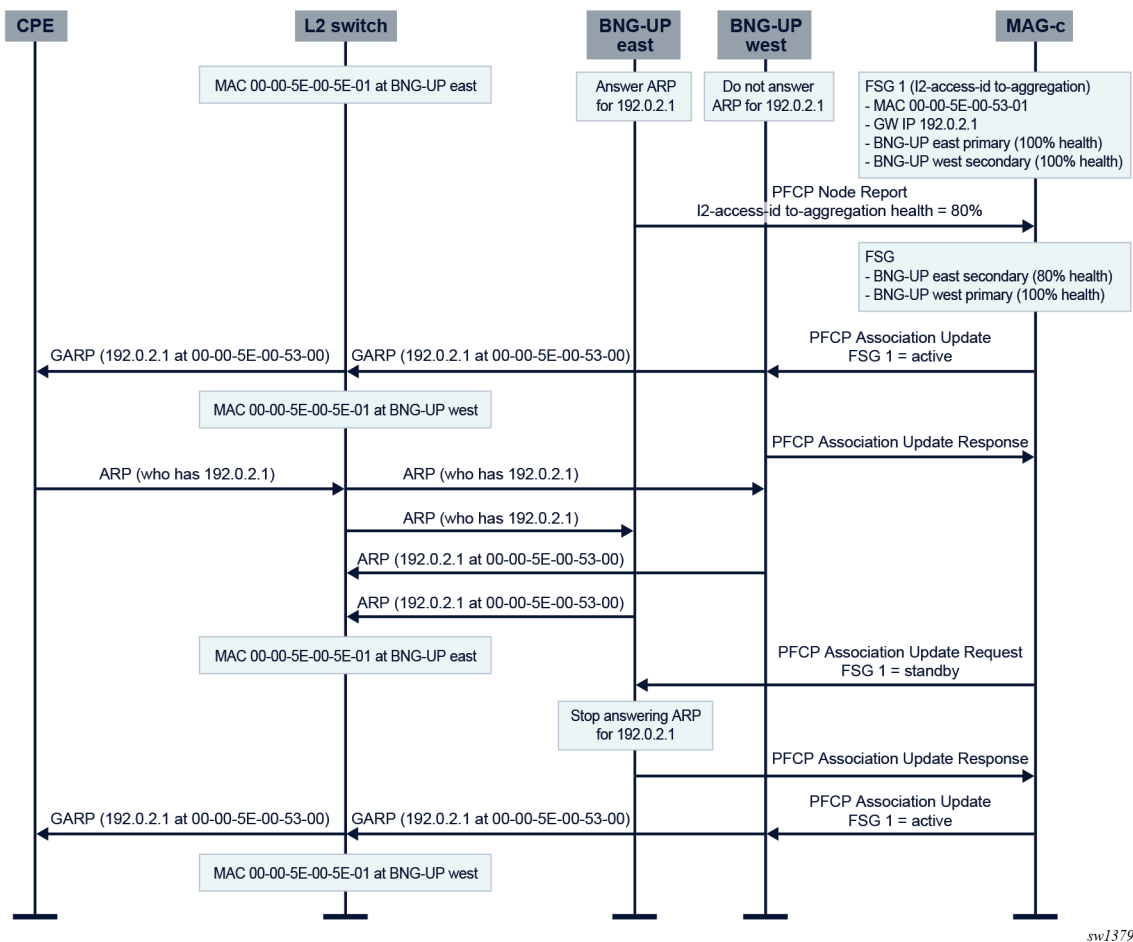
- The new standby BNG-UP has not yet processed the message that asks it to become standby.
- A regular ARP is sent and broadcast as normal.
- Both BNG-UPs answer, and the ARP response from the new standby BNG-UP comes later than the ARP response of the new active BNG-UP.

If the preceding conditions apply, the Layer 2 aggregation switch has a wrong FDB entry. Sending a second update to the new active BNG-UP can act as a new GARP trigger to correct the situation. [Figure 24: GARP race conditions](#) shows this case.



**Note:** The second update is a very lightweight operation as no actual FSG changes need to occur. It only acts as a GARP trigger. The BNG-UP may not have any action to perform if it does not need to send GARPs; for example, on aggregation networks where the FDBs are populated out-of-band such as EVPN networks.

Figure 24: GARP race conditions



## 9.4.7 UP Lockout

To handle FSG failure scenarios, the MAG-c can put a specific BNG-UP in lockout for that FSG. The following example scenarios trigger lockout:

- an explicit FSG error from the BNG-UP when signaling an FSG create, modify, or delete
- the path of a BNG-UP goes down, in addition to setting its health to -1 (unavailable) (see [BNG-UP health determination](#))

The MAG-c sees putting a BNG-UP in lockout as a degradation trigger for the FSG (see [Active/standby selection triggers](#)). The MAG-c attempts to remove the locked out BNG-UP from being selected as either active or standby (see [Active/standby selection](#)).

Because many failure scenarios do not have an automatic recovery signal, the lockout is subject to a timer. For explicit FSG errors, use the following command to configure the lockout timer:

```
configure mobile-gateway profile bng fsg-profile active-standby-selection failure-lockout
```

For other scenarios, the lockout timer is set to a fixed value, typically equal to the minimal configurable value. When the lockout timer expires, the MAG-c takes one of the following actions:

- If the BNG-UP is not active or standby for the FSG, the MAG-c removes the lockout state and triggers a recovery reselection for the FSG.
- Otherwise, the MAG-c restarts the lockout timer with a fixed value and takes no further action. This guarantees that the BNG-UP is removed from the FSG at least once and starts from a clean slate before it can be made active or standby again.

## 9.5 Warm and hot standby

Warm and hot standby in BNG-UP resiliency is a per-session concept that defines how a session is handled on the standby BNG-UP:

- Warm standby sessions are created on the standby BNG-UP when the BNG-UP becomes active. The sessions are not precreated on the standby BNG-UP. This saves resources on the standby BNG-UP, but it takes a significantly longer time during which there is no forwarding capability for those sessions.
- Hot standby sessions are precreated on the standby BNG-UP. As soon as the BNG-UP becomes active, it can start forwarding traffic for those sessions. While this consumes more resources on the standby BNG-UP, it can offer significantly reduced forwarding loss during switchovers. Depending on the capabilities of the aggregation network, it may even be possible to achieve non-loss planned switchovers; for example, to seamlessly handle BNG-UP upgrades.

For hot standby, any procedure that interacts with a BNG-UP change (for example, a CoA with a QoS update) first applies the change on the active BNG-UP. If the change succeeds, the procedure continues as usual and updates the standby BNG-UP in parallel. In the unlikely event that only the standby BNG-UP update fails, the MAG-c does not fail the triggering procedure. Instead, it tries to reapply the update periodically in the background until the standby BNG-UP is realigned with the active BNG-UP. If this realignment is not resolved when the standby BNG-UP becomes active, the MAG-c does one final attempt to update the session state and if not successful, locally removes the full session.

By default, the MAG-c creates a session in an FSG scope always in the hot standby mode. To change the default at a per-FSG level, use the following command:

```
configure mobile-gateway profile bng fsg-profile default-standby-mode
```

To overwrite the default at a per-session level, use the following command:

```
configure mobile-gateway profile authentication-database entry resiliency standby-mode
```

See [BNG EP and ADB lookup](#) for more information about how the MAG-c chooses an ADB entry.



**WARNING:** At large scale and depending on the install rate of the involved BNG-UPs, it can take a long time for warm standby sessions to switch over. Timers such as the PPP keepalive, DHCP lease times, and RA lifetimes may time out before the switchover is completed if they are set too short.



## 9.6 Interaction with headless mode

BNG-UP resiliency is supported in combination with the BNG-UP headless mode (see [Headless mode](#)). When a BNG-UP becomes headless, its health becomes unavailable (-1) because the MAG-c cannot differentiate between a BNG-UP toward which communication failed (headless) or a BNG-UP that completely failed. See [BNG-UP health determination](#) for more information.

A BNG-UP becoming headless acts as a trigger to perform a potential switchover from active to standby. A switchover cannot be signaled to the headless BNG-UP, which operates on stale data. The Nokia BNG-UP, by default, keeps its FSG state from before becoming headless. As a result, it is possible that there is an active/active forwarding situation in which both the headless and non-headless BNG-UPs of an FSG have an active state. In this scenario, the following applies:

- The MAG-c does not act on any LCP keepalive failure reports coming from the BNG-UP because it is possible that the headless BNG-UP is handling the keepalives. After the headless recovery, failures are again handled as normal.
- QoS cannot always be guaranteed because traffic may switch from one BNG-UP to the other at any time. After headless recovery, the active/standby situation stabilizes and traffic flows through only one BNG-UP with normal QoS guarantees.
- During headless, accounting reports may be off because traffic on the headless BNG-UP is not counted. After headless recovery, the MAG-c can fetch the missing statistics and the accounting is corrected.
- If there is unicast replication in the access network, these packets may end up being replicated also in the data network.

If the consequences of the active/active state are not wanted, the Nokia BNG-UP can be configured to always automatically make any FSG standby when the headless conditions occur. This configuration avoids an active/active state, and one of following scenarios occurs:

- When a single BNG-UP is headless, that BNG-UP makes its FSGs standby and the MAG-c makes the other BNG-UP active. This results in an active/standby state as expected.
- When both BNG-UPs are headless; for example, because of a networking issue at the MAG-c, the FSG becomes standby on all BNG-UPs and all traffic is dropped.

As an exception, in case all BNG-UPs go headless simultaneously, the MAG-c does not take any action on the FSGs. This may occur when the common network infrastructure between all BNG-UPs and the MAG-c fails. To avoid moving all FSGs to the first BNG-UP to recover, the MAG-c initiates a custom hold timer after the first BNG-UP recovers. Only after this timer expires or after all BNG-UP paths recover and send their health, the MAG-c restarts the FSG selection procedure. This mechanism avoids needless FSG switches because BNG-UP failure is unlikely in this scenario.



**Note:** Because of the consequences of headless mode as described above, Nokia recommends reducing the path restoration time to the absolute minimum that is acceptable for the planned deployment.

## 9.7 Operational commands

The same operational commands can be used for resilient sessions as for non-resilient sessions. See [Operational commands and debugging](#).

The MAG-c supports the following commands to display information about BNG-UPs, UP groups, FSGs, and the relationship between them. The commands are in the following context:


```
show mobile-gateway pdn bng
```

- **up-group**  
This command provides an overview of all configured UP groups, and the number of sessions and BNG-UPs that are currently active in the UP groups.
- **up-group *group-name***  
This command provides an overview of the specific UP group, the associated list of FSGs, and the associated BNG-UPs.
- **fsg**  
This command provides an overview of the UP group of the specific FSG and the current active/standby BNG-UPs. This is useful to quickly retrieve information when a specific FSG ID is known. The following are examples to get a specific FSG ID:
  - from a command in the following context:

```
show mobile-gateway bng session
```
  - from a BNG-UP-specific operational command
- **up**  
This command provides an overview of all BNG-UPs and their UP group participation. The applicable Layer 2 access ID, the applicable VLAN range, and the active/standby state for the BNG-UP in the UP group are displayed.



# Customer document and product support



## Customer documentation

[Customer documentation welcome page](#)



## Technical support

[Product support portal](#)



## Documentation feedback

[Customer documentation feedback](#)