



Multi-Access Gateway – controller

Release 26.3

Control Plane Function Guide

3HE 22200 AAAA TQZZA

Edition: 01

March 2026

© 2026 Nokia.

Use subject to Terms available at: www.nokia.com/terms.

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.

© 2026 Nokia.

Table of contents

| | |
|---|-----------|
| List of tables | 11 |
| List of figures | 13 |
| 1 Getting started | 15 |
| 1.1 About this guide..... | 15 |
| 1.2 Conventions..... | 15 |
| 1.2.1 Precautionary and information messages..... | 15 |
| 1.2.2 Options or substeps in procedures and sequential workflows..... | 16 |
| 2 MAG-c configuration quick start | 17 |
| 2.1 Configuring MAG-c for subscriber management..... | 17 |
| 2.1.1 Configuration overview..... | 17 |
| 2.1.2 Configuring MAG-c..... | 18 |
| 2.2 Configuring Service Based Interfaces..... | 24 |
| 3 Session management | 28 |
| 3.1 PFCP session..... | 28 |
| 3.2 General session functionality..... | 28 |
| 3.2.1 PFCP protocol..... | 28 |
| 3.2.1.1 PFCP association and path..... | 28 |
| 3.2.1.2 BNG-UP selection..... | 29 |
| 3.2.1.3 PFCP session state..... | 29 |
| 3.2.1.4 PFCP provisioning..... | 29 |
| 3.2.1.5 Session message priority..... | 32 |
| 3.2.1.6 UP overload protection..... | 32 |
| 3.2.2 PFCP connectivity failure..... | 33 |
| 3.2.2.1 Headless mode..... | 33 |
| 3.2.2.2 Session timer alignment..... | 34 |
| 3.2.3 Subscribers..... | 35 |
| 3.2.4 QoS..... | 36 |
| 3.2.5 Service selection..... | 37 |
| 3.2.6 Operational commands and debugging..... | 37 |
| 3.2.6.1 Call trace..... | 38 |

| | | |
|---------|---|----|
| 3.2.6.2 | Simulating data triggers for IPoE sessions..... | 41 |
| 3.2.7 | CP session state and VM resilience..... | 41 |
| 3.2.8 | Prefix delegation as a framed route..... | 42 |
| 3.2.9 | Configuring subscriber information in PFCP IEs to the UPF..... | 43 |
| 3.3 | MAG-c lawful intercept..... | 44 |
| 3.3.1 | MAG-c LI solution for wireline application..... | 44 |
| 3.3.2 | MAG-c LI solution for FWA applications..... | 45 |
| 3.3.2.1 | Configuring FWA LI IRI for 5G RGs..... | 45 |
| 3.3.2.2 | Configuring FWA LI IRI for 4G RGs..... | 46 |
| 3.3.2.3 | Configuring FWA LI CC for 4G and 5G RGs..... | 46 |
| 3.3.2.4 | Additional information for 3GPP IRI messages..... | 47 |
| 3.3.3 | Alternative MAG-c LI solution through the UP..... | 48 |
| 3.3.3.1 | Subscriber ID and IP address notifications for LI meditation devices..... | 48 |
| 3.4 | Fixed access sessions..... | 49 |
| 3.4.1 | Layer 2 circuit..... | 49 |
| 3.4.2 | In-band control plane and BNG-UP selection..... | 50 |
| 3.4.3 | Session keys and anti-spoofing..... | 51 |
| 3.4.4 | Subscriber identification..... | 52 |
| 3.4.5 | Session limits..... | 52 |
| 3.4.6 | Session lockout..... | 54 |
| 3.4.6.1 | Enabling session lockout..... | 54 |
| 3.4.7 | IPoE..... | 55 |
| 3.4.7.1 | Static IPoE sessions..... | 56 |
| 3.4.8 | PPPoE..... | 56 |
| 3.5 | Fixed wireless access sessions..... | 66 |
| 3.5.1 | Introduction to FWA sessions and terminology..... | 66 |
| 3.5.2 | Residential Gateway models..... | 68 |
| 3.5.3 | Selected and real APN..... | 68 |
| 3.5.4 | Session identification, subscriber identification, and multi-APN support..... | 70 |
| 3.5.5 | FWA-UP selection..... | 70 |
| 3.5.5.1 | Configuring policy-based FWA-UP selection..... | 71 |
| 3.5.5.2 | Static FWA-UP selection per IMSI..... | 75 |
| 3.5.5.3 | Session rebalancing and revert operations..... | 76 |
| 3.5.6 | Address signaling methods and deferred allocation..... | 77 |
| 3.5.7 | QoS..... | 78 |
| 3.5.7.1 | 4G QoS attributes..... | 79 |

| | | |
|----------|--|-----|
| 3.5.7.2 | 5G QoS attributes..... | 80 |
| 3.5.7.3 | Dynamic QoS based on PCC rules..... | 81 |
| 3.5.7.4 | DSCP marking..... | 84 |
| 3.5.8 | PAP/CHAP authentication..... | 85 |
| 3.5.9 | Session lifetime and held addresses..... | 85 |
| 3.5.10 | PDN type selection..... | 86 |
| 3.5.11 | Mobility and idling..... | 87 |
| 3.5.11.1 | Location tracking..... | 87 |
| 3.5.12 | Headless mode for FWA..... | 88 |
| 3.5.13 | 4G and 5G NSA option 3 sessions..... | 88 |
| 3.5.14 | 5G standalone sessions..... | 91 |
| 3.5.14.1 | SMF PDUSession service server configuration..... | 93 |
| 3.5.14.2 | AMF client communication service configuration..... | 94 |
| 3.5.14.3 | UDM function configuration..... | 95 |
| 3.5.14.4 | FWA session setup for 5G..... | 97 |
| 3.5.14.5 | 5G slicing..... | 100 |
| 3.5.14.6 | 5G-4G interworking..... | 100 |
| 3.5.14.7 | AMF set support..... | 102 |
| 3.5.15 | Signaling IPv4 MTU to UE..... | 102 |
| 3.5.16 | Load control, overload control, and congestion mitigation..... | 103 |
| 3.5.16.1 | Load control configuration guidelines..... | 103 |
| 3.5.16.2 | Overload control configuration guidelines..... | 104 |
| 3.5.16.3 | Congestion mitigation..... | 106 |
| 3.5.17 | Wireless Priority Service..... | 107 |
| 3.5.17.1 | WPS overview..... | 107 |
| 3.5.17.2 | WPS configuration guidelines..... | 107 |
| 3.5.18 | Radio UP security profiles..... | 109 |
| 3.5.19 | Aggregate statistics collection..... | 110 |
| 3.5.20 | DNS-based content filtering..... | 111 |
| 3.5.21 | Enabling multiple sessions per APN or DNN for each subscriber..... | 112 |
| 3.6 | Address assignment protocols..... | 112 |
| 3.6.1 | DHCP..... | 113 |
| 3.6.2 | IPCP..... | 114 |
| 3.6.3 | ICMPv6 Router Advertisements and SLAAC..... | 114 |
| 3.6.4 | DHCPv6..... | 116 |
| 3.7 | YANG state..... | 117 |

| | | |
|----------|---|------------|
| 4 | Service Based Interfaces | 120 |
| 4.1 | Introduction to Service Based Architecture | 120 |
| 4.2 | NF peer discovery | 121 |
| 4.2.1 | NRF-based discovery | 122 |
| 4.2.2 | Enabling NRF-based discovery | 124 |
| 4.2.3 | Enabling discovery based on a local NF ID list | 124 |
| 4.2.4 | Enabling DNS-based IP lookup | 126 |
| 4.3 | URI construction | 127 |
| 4.4 | NF registration | 127 |
| 4.5 | HTTP/2, OpenAPI, and Python | 132 |
| 4.6 | Failure handling | 133 |
| 4.6.1 | NRF failure handling | 135 |
| 4.6.2 | Operational state and CLI commands | 136 |
| 4.7 | SCP model D communication | 137 |
| 4.7.1 | SCP client configuration | 137 |
| 4.7.2 | SCP failure handling | 138 |
| 4.7.3 | SCP load distribution | 138 |
| 5 | Address assignment | 139 |
| 5.1 | Overview of address assignment | 139 |
| 5.2 | ODSA and local address assignment | 139 |
| 5.2.1 | ODSA | 139 |
| 5.2.2 | Variable prefix length and micro-net length | 141 |
| 5.2.3 | Local address assignment | 142 |
| 5.3 | AAA-based address assignment from RADIUS, HSS, or UDM | 143 |
| 5.4 | AAA framed routes | 144 |
| 5.5 | Non-provisioned address assignment | 145 |
| 5.6 | Local static address assignment via authentication database | 146 |
| 5.7 | Address assignment for data-triggered IPoE sessions | 146 |
| 5.8 | External DHCPv4 and DHCPv6 server address assignment | 147 |
| 5.8.1 | Local ODSA pool tracking | 147 |
| 5.8.2 | DHCPv4 relay | 147 |
| 5.8.3 | DHCPv6 relay | 148 |
| 5.8.4 | DHCP and DHCPv6 for data-triggered IPoE sessions | 151 |
| 5.8.5 | Dual-stack relay | 151 |

| | | |
|----------|---|------------|
| 5.8.6 | DHCP options..... | 152 |
| 5.8.7 | DHCP relay and BNG-UP redundancy..... | 152 |
| 6 | Authentication..... | 154 |
| 6.1 | Overview of the authentication process..... | 154 |
| 6.2 | BNG entry point..... | 154 |
| 6.3 | Authentication database..... | 155 |
| 6.4 | Authentication flow..... | 157 |
| 6.5 | BNG EP and ADB lookup..... | 158 |
| 6.6 | Required minimal configuration for a session creation..... | 161 |
| 6.7 | RADIUS authentication profile..... | 162 |
| 6.8 | RADIUS CoA and DM..... | 163 |
| 6.9 | Example configuration..... | 164 |
| 6.10 | Web portal authentication..... | 166 |
| 6.10.1 | Configuring WPP..... | 168 |
| 7 | PCF-based policy management..... | 171 |
| 7.1 | Enabling PCF communication..... | 173 |
| 7.2 | PCC rule retrieval..... | 174 |
| 7.3 | Predefined PCC rules..... | 175 |
| 7.3.1 | Evaluation of predefined PCC rules during session setup..... | 176 |
| 7.3.2 | Evaluation of predefined PCC rules during session update..... | 176 |
| 7.3.3 | Predefined PCC rules configuration example..... | 177 |
| 7.3.4 | Use cases for predefined PCC rules..... | 178 |
| 7.4 | PCF failure handling..... | 180 |
| 8 | Accounting and charging..... | 181 |
| 8.1 | BNG charging profiles..... | 181 |
| 8.2 | Statistics collection from the BNG-UP..... | 182 |
| 8.3 | MAG-c-based charging..... | 184 |
| 8.4 | RADIUS accounting..... | 185 |
| 8.4.1 | Enabling RADIUS accounting..... | 186 |
| 8.4.2 | Session accounting..... | 186 |
| 8.4.3 | Message retransmission and buffering..... | 188 |
| 8.5 | CHF charging..... | 189 |
| 8.5.1 | Enabling CHF charging..... | 189 |

| | | |
|----------|---|------------|
| 8.5.2 | CHF peer discovery..... | 190 |
| 8.5.2.1 | Selecting predefined CHF peers based on charging characteristics..... | 191 |
| 8.5.3 | Charging sessions and rating groups..... | 191 |
| 8.5.3.1 | Starting a charging session toward the CHF..... | 193 |
| 8.5.4 | Session charging using a session RG..... | 194 |
| 8.5.5 | Custom charging using a static RG..... | 194 |
| 8.5.6 | Offline Charging..... | 194 |
| 8.5.7 | Online Charging..... | 195 |
| 8.5.8 | Charging update triggers..... | 197 |
| 8.5.9 | CHF failure handling..... | 199 |
| 9 | Residential NAT..... | 203 |
| 9.1 | NAT terminology and references..... | 203 |
| 9.2 | Residential NAT44 on BNG CUPS..... | 204 |
| 9.3 | Functional split between MAG-c and BNG-UP..... | 205 |
| 9.4 | Management of NAT outside prefixes..... | 205 |
| 9.5 | CP NAT profile..... | 206 |
| 9.6 | Port forwards..... | 207 |
| 9.7 | Extended port blocks..... | 207 |
| 9.7.1 | Port space division..... | 208 |
| 9.7.2 | Managing port block space..... | 208 |
| 9.8 | NAT logging..... | 210 |
| 9.8.1 | RADIUS-based logging..... | 211 |
| 9.8.1.1 | Enabling RADIUS logging on MAG-c..... | 215 |
| 9.8.1.2 | Timestamp interpretation..... | 216 |
| 9.8.1.3 | High logging rates..... | 217 |
| 9.8.1.4 | Buffering during RADIUS failure..... | 217 |
| 9.9 | Watermarks..... | 218 |
| 9.10 | Minimum configuration steps..... | 218 |
| 9.11 | Monitoring NAT resources..... | 219 |
| 9.11.1 | Terminology overview..... | 220 |
| 9.11.2 | NAT pool resources in ODSA..... | 221 |
| 9.11.3 | Number of subscribers and extended PBs..... | 224 |
| 9.11.4 | Subscriber sessions and CP NAT profile..... | 226 |
| 9.11.5 | Histograms..... | 229 |

| | | |
|-----------|---|------------|
| 10 | Geo-redundancy | 231 |
| 10.1 | Geo-redundancy overview | 231 |
| 10.2 | Operational and administrative roles | 232 |
| 10.3 | Traffic detection | 232 |
| 10.4 | State synchronization | 233 |
| 10.5 | Routing | 234 |
| 10.6 | Shunting | 234 |
| 10.7 | Manual switchover | 235 |
| 10.8 | Deploying and configuring geo-redundancy | 235 |
| 11 | Python support | 239 |
| 11.1 | Configuring a Python script | 242 |
| 11.2 | Protecting a Python script file | 243 |
| 12 | BNG-UP resiliency | 244 |
| 12.1 | Terminology for BNG-UP resiliency | 244 |
| 12.2 | Introduction to MAG-c-driven BNG-UP resiliency | 244 |
| 12.3 | Modeling a resilient BNG-UP deployment using UP groups | 246 |
| 12.3.1 | Fate sharing group creation | 246 |
| 12.3.2 | Fixed access with broadcast access | 247 |
| 12.3.3 | Fixed access with active/standby pseudowire access | 252 |
| 12.3.3.1 | Configuring an A/S PW port on the Nokia BNG-UP | 254 |
| 12.3.3.2 | Maintenance on a Nokia BNG-UP in an A/S PW setup | 254 |
| 12.3.3.3 | Bringing down the PW in headless conditions | 256 |
| 12.3.3.4 | Moving an FSG to standby in headless mode | 257 |
| 12.4 | FWA UPF mesh | 258 |
| 12.4.1 | Provisioning an FWA UPF mesh | 258 |
| 12.4.2 | Inter-site redundancy for FWA sessions | 262 |
| 12.4.2.1 | Creating sites with assigned MAG-u nodes | 262 |
| 12.4.3 | Anti-affinity to avoid FSG creation | 263 |
| 12.4.4 | Session rebalance and revert | 265 |
| 12.4.5 | Adding a MAG-u to an existing UP group for FWA sessions | 267 |
| 12.4.6 | Removing a MAG-u from an existing UP group | 267 |
| 12.5 | Maintenance on a Nokia BNG-UP | 268 |
| 12.6 | Fate sharing groups | 269 |

| | | |
|-----------|---|------------|
| 12.6.1 | Session-to-FSG mapping and load-balancing..... | 270 |
| 12.6.2 | Traffic steering parameters..... | 271 |
| 12.6.3 | BNG-UP health determination..... | 274 |
| 12.6.4 | Active/standby selection triggers..... | 277 |
| 12.6.5 | Active/standby selection..... | 279 |
| 12.6.6 | Active/standby change or switchover..... | 281 |
| 12.6.7 | UP lockout..... | 282 |
| 12.6.8 | Monitoring ongoing FSG changes..... | 283 |
| 12.7 | Hot standby..... | 285 |
| 12.8 | Interaction with headless mode..... | 285 |
| 12.9 | Operational commands..... | 286 |
| 13 | KPI and KCI performance monitoring..... | 287 |
| 13.1 | KPI and KCI XML file naming..... | 287 |
| 13.2 | Storage and retrieval of KPI and KCI XML files..... | 288 |
| 13.3 | Content and format of KPI and KCI XML files..... | 288 |
| 13.4 | Record types and counter module mappings..... | 289 |
| 13.5 | Measurement objects groups and indexes..... | 290 |

List of tables

| | |
|---|-----|
| Table 1: MAG-c configuration components and quick-start steps..... | 17 |
| Table 2: Message header for session-related PFCP messages..... | 32 |
| Table 3: IPID encoding..... | 47 |
| Table 4: IPID encoding example..... | 48 |
| Table 5: Use cases and context for the apn-format command..... | 69 |
| Table 6: Example overload reporting and local throttling..... | 105 |
| Table 7: Message throttling priority..... | 106 |
| Table 8: Address assignment protocols per session type..... | 113 |
| Table 9: URI scheme..... | 127 |
| Table 10: Attributes sent during NRF registration..... | 129 |
| Table 11: Minimal configuration for a session creation..... | 161 |
| Table 12: Update the app-profile to enable redirect..... | 178 |
| Table 13: Update policy to throttle all traffic..... | 179 |
| Table 14: Update policy to throttle selective traffic..... | 179 |
| Table 15: Time sequence for an RG with a time and volume limit..... | 195 |
| Table 16: CHF cause code to PCF creditManagementStatus code mappings..... | 199 |
| Table 17: CHF RG result code to PCF creditManagementStatus code mappings..... | 200 |
| Table 18: RADIUS based logging for residential NAT..... | 211 |
| Table 19: Subscriber management and NAT integrated RADIUS accounting and logging..... | 212 |
| Table 20: Supported direction for RADIUS messages..... | 240 |
| Table 21: Supported direction for RADIUS CoA messages..... | 240 |

| | |
|--|-----|
| Table 22: Supported direction for PPPoE messages..... | 240 |
| Table 23: Supported direction for DHCPv4 messages..... | 241 |
| Table 24: Supported direction for DHCPv6 messages..... | 241 |
| Table 25: Summary of BNG-UP states..... | 276 |

List of figures

| | |
|---|-----|
| Figure 1: HQoS example..... | 36 |
| Figure 2: IPoE session setup with RADIUS authentication..... | 55 |
| Figure 3: PPPoE session setup flow..... | 57 |
| Figure 4: Resiliency based on PADO delay..... | 62 |
| Figure 5: L2TP LAC network components..... | 63 |
| Figure 6: LAC-enabled PPPoE session setup flow..... | 64 |
| Figure 7: A separate model with PPP connectivity..... | 68 |
| Figure 8: Example of a bearer or QoS flow rate enforcement..... | 79 |
| Figure 9: Downlink QoS enforcement on a FWA-UP..... | 83 |
| Figure 10: Uplink QoS enforcement on a FWA-UP..... | 83 |
| Figure 11: 4G/5G QoS aware versus non-QoS aware NEs..... | 84 |
| Figure 12: Basic FWA network..... | 89 |
| Figure 13: 4G FWA session setup..... | 90 |
| Figure 14: Basic 5G FWA network..... | 92 |
| Figure 15: 5G FWA session setup..... | 98 |
| Figure 16: Relationship between service and interface..... | 120 |
| Figure 17: Basic 5G network with service representation..... | 121 |
| Figure 18: NF-instance discovery flowchart..... | 122 |
| Figure 19: HTTP/2 stack..... | 132 |
| Figure 20: DHCPv4 call flow..... | 148 |
| Figure 21: DHCPv6 call flow with RS message..... | 149 |

| | |
|---|-----|
| Figure 22: DHCPv6 call flow without RS message..... | 150 |
| Figure 23: WPP on BNG CUPS call flow example..... | 166 |
| Figure 24: PCF Policy Control Operations..... | 173 |
| Figure 25: PCC rule request initiated by the UE..... | 175 |
| Figure 26: Statistics collection using the pull model..... | 182 |
| Figure 27: Statistics collection using the push model..... | 183 |
| Figure 28: Charging Session operations..... | 193 |
| Figure 29: CHF-triggered update..... | 198 |
| Figure 30: Residential NAT example..... | 204 |
| Figure 31: Alc-ISA-Event-Timestamp triggered Interim-Update message..... | 217 |
| Figure 32: NAT pools in the MAG-c..... | 221 |
| Figure 33: Geo-redundant MAG-c deployment..... | 231 |
| Figure 34: High-level overview of communication for BNG-UP resiliency..... | 245 |
| Figure 35: Multiple backup BNG-UPs..... | 245 |
| Figure 36: Multiple Layer 2 access IDs per UP group..... | 248 |
| Figure 37: 1:1 hot standby resiliency example..... | 250 |
| Figure 38: Per S-tag 1:1 hot standby resiliency example..... | 251 |
| Figure 39: A/S PW deployment..... | 253 |
| Figure 40: FSG mesh basic overview..... | 258 |
| Figure 41: Mesh anti-affinity..... | 265 |
| Figure 42: Example of the relationship between FSGs, MAC addresses, and subnets..... | 272 |
| Figure 43: Example of the relationship between FSGs, GTP-u endpoint addresses, and subnets..... | 274 |
| Figure 44: GARP race conditions..... | 282 |

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 generically covers content for the release specified on the title page of the guide, and may also contain some 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 software release.



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

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 MAG-c configuration quick start

The quick start steps configure the main components required to prepare the MAG-c for subscriber management or Service Based Interface (SBI) communication.

2.1 Configuring MAG-c for subscriber management

Set up and configure basic components to prepare the MAG-c for subscriber management.

2.1.1 Configuration overview

This overview describes the components and steps required to configure the MAG-c for subscriber management.

The following table lists the main components and links to the specific configuration step that you must complete to prepare the MAG-c for subscriber management.

Table 1: MAG-c configuration components and quick-start steps

| Configuration component | Configuration task |
|---|--------------------|
| Configure MAG-c system resources. | Step 1 |
| Configure IP connectivity between MAG-c and BNG-UP. | Step 2 |
| Configure the PFCP association. | Step 3 |
| Configure the IBCP and network realm (default PFCP session) for fixed-access sessions. Configure control plane communication for fixed wireless access (FWA) sessions. | Step 4 |
| Configure the local IP address pool. | Step 5 |
| Configure the APN and network realm. | Step 6 |
| Configure the authentication database (ADB). | Step 7 |
| Configure the BNG entry point for fixed access sessions. Configure the incoming APN/DNN for FWA sessions. | Step 8 |
| Configure routing for non geo-redundant (singleton) deployment. | Step 9 |

2.1.2 Configuring MAG-c

These quick-start instructions provide the basic steps required to configure the MAG-c for subscriber management, including references to additional topics for users who want more information.

Prerequisites

- Identify the name of the IES or VPRN service used to host subscriber sessions on the BNG-UP; see [Service selection](#).
- Ensure the BNG-UP is configured; see the *7750 SR and VSR BNG CUPS User Plane Function Guide*, "PFCP association" and "Session management", for more information.

About this task

This procedure describes the minimal configuration required to create subscriber sessions on the MAG-c.

Procedure

Step 1. Configure the system resources.

The values specified in the following example are for a MAG-c with a single VM instance. The assumption is the resource pool, gateway, card, and group are already configured. See [CP session state and VM resilience](#) for information about different VM resiliency models.

Example

System resource configuration with a single VM instance

```
A:MAG-c>config>mobile>system# info
-----
resource-pool 1 redundancy many-to-many gateway 1
card 2
exit
group 1 resource-pool 1
no shutdown
exit
-----
```

Step 2. Configure IP connectivity between the MAG-c and BNG-UP.

The following example shows a MAG-c directly connected to a BNG-UP. Connectivity to the system interface is also required for basic MAG-c functionality. See the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR Interface Configuration Guide* for information about IP routing configuration.

Example

IP connectivity between the MAG-c and BNG-UP

```
A:MAG-c>config>router# info
-----
interface "system"
address 1.1.1.1/32
no shutdown
exit
interface "to-BNG-UP1"
address 192.168.1.1/24
port 1/2/1
no shutdown
exit
-----
```

- Step 3.** Create a PDN instance and under it, configure the PFCP association between the MAG-c and the BNG-UP.

The following example shows a minimal configuration required to form a PFCP association between the MAG-c and the BNG-UP. See [PFCP association and path](#) for more information.

Example

PFCP association between the MAG-c and the BNG-UP

```
A:MAG-c>config>mobile# info
-----
  profile
    pfc
      pfc-association-peer-list "peers"
        pfc-peer 192.168.50.84
        pfc-peer 192.168.50.83
        pfc-peer 192.168.50.82
      exit
      up-peer-list "ups"
        peer 192.168.50.84
        peer 192.168.50.83
        peer 192.164.50.82
      exit
    exit
  exit
  pdn 1
    instance-type control
    pfc-node-id-type ip
    epc-node 001.01.magc.nokia.1.1
    sx-n4 "default"
      interface
        pfc "to-BNG-UP1"
      exit
      pfc-association-list "peers"
      signaling
        pfc
          profile "default"
        exit
      exit
      up-peer-list "ups"
      no shutdown
    exit
  -----
```

- Step 4.** Configure a fixed-access or FWA session.

- For fixed-access sessions, configure IBCP to allow tunneling of IPoE sessions. The default IBCP tunnel supports tunneling of control packets from the BNG-UP to the MAG-c. The following example configures IPoE-based tunneling sessions. See [In-band control plane and BNG-UP selection](#) for more information about IBCP.

Example

IPoE-based tunneling configuration for the IBCP tunnel

```
A:MAG-c>config>mobile>pdn#info
-----
  sx-n4 "default"
    interface
      ibcp "to-BNG-UP1"
    exit
    signaling
      ibcp
```

```

        triggers ipoe-dhcp ipoe-dhcpv6 ipoe-router-solicit
    exit
exit
-----

```

- For 4G FWA sessions, configure the S11 and S5 endpoints to enable control plane communication toward the MME. The following example shows a 4G FWA configuration for control plane traffic toward the MME. You require the S5 interface, even though it is an internal interface over which no messages are sent.



Note: 5G FWA sessions follow similar steps but require more groundwork to enable Service Based Interface (SBI) communications. See [Configuring Service Based Interfaces](#) for the specific quick-start procedure.



Note: See [4G and 5G NSA option 3 sessions](#) and [5G standalone sessions](#) for more information about configuring 4G and 5G sessions.

Example

4G FWA configuration for control plane traffic toward the MME

```

A:MAG-c>config>mobile>pdn#info
-----
    s5 "default"
        interface
            gtp-c "system" interface-realm "to_ran"
        exit
    exit
    s11 "default"
        interface
            gtp-c "system" interface-realm "to_ran"
        exit
    exit
    sx-n4
-----

```

Step 5. Configure a local address pool for IP address assignment.

The following example shows a simple address pool configuration that uses the prefix 192.168.2.0/24 for an IP address assignment. Other methods can also be used. See [Local address assignment](#) for more information about IP address assignment.

Example

Local address pool configuration for IP address assignment

```

A:MAG-c>config>mobile>pdn# info
-----
    local-address-assignment
        network-realm "realm-01"
        pool "pool-01"
            ipv4
                default-gateway first-address
                micro-net-length 24
                prefix 192.168.0.0/16
            exit
            ipv6
                pd
                    micro-net-length 48
                    prefix 2001::/37

```

```

                prefix 192.168.2.0/24
            exit
        exit
    exit
exit
-----

```

Step 6. Configure a network realm for the APN.

Configure the network realm for the APN to match the service name of the IES or VPRN service used to host the subscriber session on the BNG-UP. See [Service selection](#) for more information about network-realm configuration.

Example

Network-realm configuration for the APN

```

A:MAG-c>config>mobile>pdn# info
-----
    apn "apn-01"
      network-realm "realm-01"
      no shutdown
-----

```

For FWA sessions, provision the APN in the UP peer list created in step 3, for the BNG-UPs that can provide service for this APN. The following example shows this configuration.

Example

APN configuration in the UP peer list for FWA sessions

```

A:MAG-c>config>mobile>profile>pfcp# info
-----
    up-peer-list "ups"
      peer 192.168.50.84
        apn "apn-01"
        exit
      exit
      peer 192.168.50.83
        apn "apn-01"
        exit
      exit
      peer 192.168.50.82
        apn "apn-01"
        exit
      exit
    exit
-----

```

Step 7. Configure the ADB.

The ADB configuration determines the authentication process and address assignment method, and also retrieves attributes associated with the subscriber session. The following example uses the ADB to provide the address assignment and the subscriber and SLA profiles for the subscriber session. The address pool refers to the local-address assignment configured in step 5. See [Authentication database](#) for more information.

Example

ADB configuration

```

A:MAG-c>config>mobile>profile# info

```

```

-----
authentication-database "auth-db-01"
  match 1 attribute username
    optional
  exit
  entry "default"
    address-assignment
      local-dynamic
        ipv4-pool "pool-01"
    apn "apn-01"
    ip-anti-spoof true
    subscriber-mgmt
      sla-profile "cup-prof"
      sub-profile "cup-prof"
    exit
  no shutdown
exit
no shutdown
exit
-----

```

Step 8. Configure a fixed-access or FWA session.

- For fixed-access sessions, configure the BNG entry point. The BNG entry point informs the MAG-c about the authentication flow and other aspects of fixed-subscriber session management and setup. See [BNG entry point](#) for more information. The following example shows a simple configuration that refers to the ADB configured in step [7](#).

Example

BNG entry point configuration

```

A:MAG-c>>config>mobile>profile#
-----
  bng
    entry-point "bng-entry-01"
      match 1 attribute up-ip
        optional
      exit
      entry "1"
        ipoe
          authentication-flow
            adb "auth-db-01"
        exit
      exit
    match
      up-ip 192.168.50.84
    exit
  no shutdown
  exit
  no shutdown
  exit
exit
-----

```

- For FWA sessions, configure the incoming APN/DNN, according to the GTP Session Establishment Request (4G) or Nsmf_PDUSession_CreateSMContextRequest (5G) messages. This APN informs the MAG-c about the authentication flow used for FWA sessions. The following example shows a simple configuration that refers to the ADB configured in step [7](#).



Note: If the incoming APN and service APN are the same, it is not necessary to provide the APN via the ADB.

Example

FWA incoming APN configuration

```
A:MAG-c>config>mobile>pdn#
-----
    apn "fwa.mnc01.mcc001.gprs"
      fixed-wireless-access
      authentication-flow
      adb "auth-db-01"
    exit
  no shutdown
  exit
no shutdown
exit
-----
```

Step 9. Configure routing for non geo-redundant (singleton) deployment.

The MAG-c attracts signaling traffic by advertising routes for locally-configured reference-point interfaces via BGP. If the MAG-c is deployed in singleton mode using only one node, the routing policies must be configured without dependency on the geo-redundant state of the node.



Note: When deploying the MAG-c in singleton mode, you must disable the configuration commands that are applicable for geo-redundant deployment. See [Geo-redundancy](#) for more information.

Example

Routing policy configuration to advertise reference-point interfaces

```
A:MAG-c>config>router>policy-options>policy-statement# info
-----
    entry 10
      from
        prefix-list "prefix-list-1"
      exit
    to
      protocol bgp
    exit
  action accept
  exit
exit
-----
```

2.2 Configuring Service Based Interfaces

These quick-start instructions provide the basic steps required to configure the MAG-c to support Service Based Interfaces (SBI) communication, with references to additional topics for users who want more information.

Prerequisites

- Configure the MAG-c for basic session management; see the [Configuring MAG-c](#) quick start topic for more information.
- Administratively disable the MAG-c PDN configuration using the following command.

```
configure mobile-gateway pdn shutdown
```

About this task

This procedure describes the minimal configuration required to set up SBI communication for a MAG-c. This includes setting up the MAG-c as an SMF SBI server and as an SBI client toward the NRF, AMF, UDM, PCF, and CHF peer network functions (NFs). This configuration assumes an NRF is used to discover any NFs and is used by the NFs to discover the MAG-c.

Procedure

Step 1. Configure an interface that is used for communication to SBI services.

The following example shows a single interface that is used throughout this quick-start procedure for all SBI services, but it is possible to specify a different interface for each service.

Example

Interface for SBI communication

```
A:MAG-c>config>mobile>pdn# /configure router interface "to_sbi"
A:MAG-c>config>router>if# info
-----
      address 192.0.2.1/24
      port 1/1/1
      ipv6
        address 2001:db8:1:1::1/64
      exit
      no shutdown
-----
```

Step 2. Configure the NF instance ID that the MAG-c uses to identify itself toward other NFs, and the slices that the MAG-c is serving.

The following example uses a randomly generated universally unique ID (UUID) as NF instance ID. Nokia recommends that you should generate a new and unique UUID when configuring a new system. See [NF registration](#) for more information about these parameters.

Example

MAG-c NF instance ID and slice configuration

```
A:MAG-c>config>mobile>profile>list# info
-----
      slice-instance-list "fwa_inst"
        slice-instance "fwa_inst"
      exit
      slice-list "fwa_slices"
-----
```

```

        slice "fwa_slice" sst 1 sd 000001
        slice-instances fwa_inst
    exit
exit
-----
A:MAG-c>config>mobile>pdn# info
-----
    slices
        slice-list "fwa_slices"
    exit
    nf-profile-attributes
        nf-instance-id "ce23c7c3-b710-49ac-9e3c-764849a012fa"
    exit
-----

```

Step 3. Configure the SMF SBI server service to allow the setup of 5G FWA sessions.

In the following example:

- the MAG-c communicates with the AMF using the common "to_sbi" interface
- the FWA-UP terminates GTP-u traffic using the "to_ran" service. This matches the service used on the S11 and S5 interface as described in [Configuring MAG-c](#).



Note: For more information about 5G Session management and the SMF service, see [5G standalone sessions](#).

Example

SMF server configuration

```

A:MAG-c>config>mobile>pdn>sba-server-services# info
-----
    nsmf-pdusession "fwa_smf"
        n3-interface-realm "to_ran"
        interface "to_sbi" port 80
    exit
-----

```

Step 4. Configure the NRF SBI client to allow the MAG-c to register itself with the NRF as an SMF NF, and discover other peer NFs.

The following example configures the MAG-c with a single NRF identified by the configured UUID, listening on a single IP. The NRF is the only SBI NF that always requires local configuration. Other NFs can subsequently be discovered through the NRF. See [NF peer discovery](#) for more information.

The configuration of the example allows the MAG-c to automatically register itself with the NRF as an SMF using the provided UUID and any configured APNs. For more information about registration parameters and their provisioning, see [NF registration](#).

Example

NRF client configuration

```

A:MAG-c>config>mobile>profile>list# info
-----
    prioritized-ip-address-list "nrfs"
        address 192.168.40.1
    exit
    nf-id-list "nrfs"
        nf-prof-id 1
        prioritized-address-list "nrfs"

```

```

        uuid AB61220B-F589-40D4-978B-D9C9F5AB742E
        enable
    exit
exit
A:MAG-c>config>mobile>pdn>sba-client-services>nrf-client# info
-----
        nrf-nfm "nrf_client"
        nf-id-list "nrfs"
        interface "to_sbi"
    exit
        nrf-disc "nrf_disc_client"
        nf-id-list "nrfs"
        interface "to_sbi"
    exit
-----

```

Step 5. Configure the SBI client for all other required SBI services.

The following example shows SBI client configuration for AMF, UDM, PCF, and CHF services using SBI-based discovery. For more information about each service, see the following topics:

- [AMF client communication service configuration](#)
- [UDM function configuration](#)
- [PCF-based policy management](#)
- [CHF charging](#)

Example

AMF, UDM, PCF, and CHF client configuration

```

A:MAG-c>config>mobile>pdn>sba-client-services# info
-----
        amf-client
        namf-comm "amf_comm_client"
        interface "to_sbi"
    exit
exit
        chf-client
        nchf-convergedcharging "chf_client"
        interface "to_sbi"
    exit
exit
        pcf-client
        npcfsmpolicycontrol "pcf_client"
        interface "to_sbi"
    exit
exit
        udm-client
        nudm-sdm "sdm_client"
        interface "to_sbi"
    exit
        nudm-uecm "uecm_client"
        interface "to_sbi"
    exit
exit
-----

```

Step 6. Configure a service realm and add all SBI server and client services to enable all services.

Example

SBA service realm configuration

```
A:MAG-c>config>mobile>pdn# info
    sba-service-realm "fwa_realm"
    server-service nsmf-pduseession service-instance "fwa_smf"
    client-service nrf-nfm service-instance "nrf_client"
    client-service nrf-disc service-instance "nrf_disc_client"
    client-service namf-comm service-instance "amf_comm_client"
    client-service npcf-smpolicycontrol service-instance "pcf_client"
    client-service nudm-sdm service-instance "sdm_client"
    client-service nudm-uecm service-instance "uecm_client"
    client-service nchf-convergedcharging service-instance "chf_client"
exit
```

Step 7. Enable the PDN to start using the SBI services.

```
configure mobile-gateway pdn no shutdown
```

3 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.

3.1 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 communication are all done in the scope of a single session. A PFCP association is needed to create a PFCP session.

3.2 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.

3.2.1 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.

3.2.1.1 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.

3.2.1.2 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.

3.2.1.3 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.

3.2.1.4 PFCP provisioning

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

Components for PFCP provisioning

To enable the PFCP protocol, provision and reference the following components in the PDN configuration:

- **IP interface**

Before provisioning the PFCP protocol, an IP interface is required on both the MAG-c and the BNG-UPs. The PFCP protocol and PFCP association reference the IP interface endpoint. This means that the IP endpoints use full IP connectivity, which may require a routing protocol.



Note: You can use a direct interface, loopback interface, or system interface as an IP interface endpoint.

- **Node ID**

The MAG-c requires a Node ID for the PFCP association. This can either be an IP address or an FQDN (default). Use the following command to configure the **ip** or **fqdn** option for the PFCP node ID.

```
configure mobile-gateway pdn pfcpc-node-id-type
```

To use an FQDN, use the following command to also configure an FQDN.

```
configure mobile-gateway pdn epc-node
```

- **PFCP association peer list**

Use the following command to configure the list of peer BNG-UP devices.

```
configure mobile-gateway profile pfcpc pfcpc-association-peer-list
```

The list can be 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 PFCP peer association list is enforced using the following command, all BNG-UP devices must already be provisioned in the PFCP peer association list.

```
configure mobile-gateway pdn sx-n4 enforced-pfcpc-association-list
```

If enforcement is enabled, PFCP association requests coming from a peer that is not configured in the PFCP peer association list are not accepted and the requesting peer is not dynamically added to the list.

- **UP peer list**

Use the following command to configure a list of UPs with the supported APNs and the UE IP address pools.

```
configure mobile-gateway profile pfcpc up-peer-list
```

This context creates a list of peer FWA-UP or BNG-UP devices 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.

- **PFCP profile**

Use the commands in the following context to configure PFCP profile options.

```
configure mobile-gateway profile pfcpc pfcpc-profile
```



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. See the *7750 SR and VSR BNG CUPS User Plane Function Guide* for more information about BNG-UP configuration.

Example: PFCP provisioning in PFCP profile and PDN contexts



Note: In this example the system interface is referenced. However, other interface types (for example, direct and loopbacks) are also allowed.

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

Verifying PFCP association setup

Use the following command to view the PFCP association establishment for the MAG-c.

```
show mobile-gateway pdn ref-point-peers sx-n4
```

Output example: PFCP association reference point peers

```
A:MAG-c# show mobile-gateway pdn ref-point-peers sx-n4
=====
PFCP reference point peers
=====
Peer address : 192.168.50.84
Node Id : 192.168.50.84
Router : Base
Path Mgmt State : up
Create Time : 06/21/2023 13:53:16 Gateway Id : 1
UP Features : FTUP TREU EMPU PDIU FRRT ADPDP MNOP IP6PL
```

```

UP BBF Features : PPPOE IPOE LCPK0
UP Nokia Features: Bulk-Audit LAC SSSG FSG
UP Association : up   Last Change Time : 06/21/2023 13:53:16
UP Selection : False
Enforced PFCP association list : No
-----
Number of peers : 1
=====

```

Related topics

[Headless mode](#)

3.2.1.5 Session message priority

To prioritize PFCP session-related messages toward the BNG-UP, the MAG-c uses the message priority field of the PFCP header. The value of the message priority ranges from 0 to 15, with 0 indicating the highest priority. For more information about how the BNG-UP enforces the priority, see the *7750 SR and VSR BNG CUPS User Plane Function Guide*.

Table 2: Message header for session-related PFCP messages

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

The MP bit, which can only be set for session-related PFCP messages, indicates whether the message priority field is set. If the MP bit is not set, the message priority field is ignored and the BNG-UP determines the priority. The Nokia BNG-UP usually uses priority 15 if the MP bit is not set. See the *7750 SR and VSR BNG CUPS User Plane Function Guide* for more information.

The message priority field indicates the actual priority. The two highest bits are based on the session priority classification using the priority-mapping described in [Wireless Priority Service](#). For example, PFCP messages for sessions with a **critical** priority have a message priority value of 0, 1, 2, or 3; PFCP messages for sessions with a **regular** priority have a message priority value of 12, 13, 14, or 15, so always a lower priority. The two lowest bits are used for prioritizing different types of procedures, see [UP overload protection](#) for more information.

3.2.1.6 UP overload protection

The MAG-c uses the following mechanisms to provide basic UP overload protection for the PFCP protocol:

- The MAG-c enforces a maximum number of outstanding PFCP transactions. When this limit is reached, the MAG-c automatically fails the PFCP transactions of new procedures until the outstanding PFCP transactions are lower than this maximum.



Note: The consequence of a failed PFCP transaction depends on the procedure type. Some failed procedures require a trigger to retry; for example, a session setup procedure. Other procedures continue without the up-to-date UP data; for example, a charging report may be sent with outdated statistics. Other procedures are retried automatically; for example, a session reinstallation in the case of a UP resiliency recovery.

- The MAG-c prioritizes messages based on the importance of the procedure. For example, FWA idle, reactivation, and handover procedures have higher priority than a session creation procedure. This ensures that setting up new sessions does not impact existing sessions.

3.2.2 PFCP connectivity failure

To protect against temporary PFCP connectivity failures, MAG-c supports a headless mode. Following the rules for configuring sessions timers ensures that the headless mode works as expected.

3.2.2.1 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 pfc pfc-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 pfc pfc-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.

Use the commands in the following context to configure message retransmit options.

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



Note: When using headless mode, Nokia recommends configuring the total message retransmit timeout for all other messages to be longer than the time to detect headless mode.

- When BFD is enabled, configure the message retransmit to be higher than the interval times the multiplier. It is important to use the operational, negotiated values and not the configured values because these could differ.

- When BFD is not enabled, configure the message retransmit to be higher than the value of the **interval** command plus the **retry-count** command times the **timeout** command, as configured in the **heart-beat** context.

When headless mode is enabled, the sessions are not removed when there is a heartbeat failure. Instead, the configured timer starts 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.

$$\text{total heartbeat timeout} = \text{heart-beat retry-count } N1 \times \text{heart-beat timeout } T1$$

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](#)

3.2.2.2 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 display information locally, use the following commands:

- **path restoration time**

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

- **DHCP lease time**

```
configure mobile-gateway profile bng dhcp-profile options option
```

with the value for the *option-number* variable set to 51

- **renew time**

```
configure mobile-gateway profile bng dhcp-profile options option
```

with the value for the *option-number* variable set to 58

- **router lifetime in RA messages**

```
configure mobile-gateway profile bng ra-profile options router-lifetime
```

- **maximum advertisement interval**

```
configure mobile-gateway profile bng ra-profile advertisement-interval max
```

- **IPv6 preferred lifetime**

```
configure mobile-gateway profile authentication-database entry address-assignment lifetimes preferred
```

3.2.3 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.


3.2.4 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.

The following figure shows an example of a HQoS model.

Figure 1: HQoS example



sw2080

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 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.

3.2.5 Service selection

The MAG-c model is based on 3GPP, which means that 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
-----
```

3.2.6 Operational commands and debugging

Use the commands in the **show**, **clear**, and **tools** contexts to display MAG-c sessions and subscribers, remove a session, and debug a failing session setup. Use the call trace feature for advanced debugging.

Removing, debugging, and displaying session information

Use the following command to display information about MAG-c sessions and subscribers.

```
show mobile-gateway bng session
```

- The **session** command displays basic data related to a session.
- The supported filter options for the **session** command display the data for specific sets of sessions.
- The **count** keyword displays only the number of matching sessions.
- Use other commands in the session context to display specific details for the session (for example, the **charging** command displays session charging data).

Alternatively, use the following command to display information about MAG-c sessions and subscribers based on state queries.

```
show mobile-gateway bng queries session
```

The input of this command is the name of a query configured using the following command.

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

See [YANG state](#) for more information on how to configure queries.

Depending on the configuration under the **output-options** context of the query, the displayed information is as follows:

- If the **count** option is configured, the command displays the number of sessions matching the query.

- If any other option is configured, or no option is configured, the command displays a short summary of basic session identifiers, such as MAC address, IMSI, SUPI, L2 access ID, or APN, together with the UP node ID and UP subscriber ID.

Use the following command to display an overview of the operational data related to a subscriber. The command has similar options as the session command.

```
show mobile-gateway bng subscriber
```

Use the **ibcp** keyword in the following context to display IBCP statistics for fixed-access sessions.

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

Use the following command to remove a session from the MAG-c. When you issue this command, the MAG-c sends a PFCP Session Deletion Request to the BNG-UP.

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

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



Tip: Always add filters to the **session** command in the **clear mobile-gateway bng** context to avoid accidentally clearing all sessions.

Use the following command to display a basic error log to debug a failing setup session.

```
tools dump mobile-gateway bng error-history
```

Related topics

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

[PFCP connectivity failure](#)

3.2.6.1 Call trace

For advanced debugging, the MAG-c supports the call trace feature, which allows tracing of control-plane packets or events during the lifetime of a session.

Call trace allows correlation of multiple signaling interfaces and detection of 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.

When call trace is enabled for a session on the MAG-c, the used call trace profile is signaled to the BNG-UP for that specific session. The BNG-UP can subsequently use that profile to start an automatic call trace session for that session, if automatic call trace is enabled on the BNG-UP.

To enable automatic call trace for CUPS sessions on the BNG-UP, one of the following requirements must apply:

- A call trace profile with the same name as the one provided by the MAG-c exists on the BNG-UP.
- A default call trace profile is configured on the BNG-UP.

On the BNG-UP, use the following MD-CLI command to configure a default call trace profile.

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



Note: The preceding command is a BNG-UP MD-CLI command; all other commands in this topic are MAG-c commands.

Use the following command to enable call trace debugging of fixed access sessions.

```
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.

Use the following command to enable call trace debugging of FWA sessions.

```
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. These 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 system call-insight location
```

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

- ongoing traces in the `calltrace/running/` folder
- 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 system performance. Therefore, Nokia recommends not to send call trace output to the debug logging system when a large number of messages are 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:

- Use the following command to create a profile and configure the profile parameters.

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

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

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

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 limit 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 specify conditions to automatically finish the tracing. By default, these limits are set to 10 MB and 1 day.

Use the following commands to display the details for an ongoing call trace capture.

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

Output example: Call trace detail for a MAC address

```
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.
=====
```

3.2.6.2 Simulating data triggers for IPoE sessions

The MAG-c supports configuration of data-trigger simulation options for IPoE sessions.

Use the following command to configure data-triggers for IPoE sessions. You can use the command options to define the details of the data-trigger packet.

```
tools perform mobile-gateway pdn bng data-triggered-host
```

The following message confirms if the data-trigger generation is successful.

```
INFO: CLI Data-trigger-host setup request successful,please monitor and validate further
session establishment.
```

3.2.7 CP session state and VM resilience

For scalability, the Nokia MAG-c distributes the session state functionality over multiple 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. 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. The SM-VMs must be configured in 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:MAG-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 SM-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 SM-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 failed SM-VM are irretrievably lost. The following example shows the database connectivity configuration.



Note:

The VNF deployment of the MAG-c does not support redundant DB-VM. However, you can manually configure two peers to deploy a redundant DB-VM.

The MAG-c appliance supports redundant DB-VM. Redundant DB-VM peers are configured automatically during the MAG-c appliance installation. See the *MAG-c Appliance Installation and Deployment Guide* for more information.

```
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
-----
```

3.2.8 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 following context to **true**.

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

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 following context is set to **false**.

```
configure mobile-gateway profile authentication-database entry
```

Related topics

[AAA-based address assignment from RADIUS, HSS, or UDM](#)

3.2.9 Configuring subscriber information in PFCP IEs to the UPF

To optionally send subscriber information in PFCP IEs to the UPF, you can configure a PFCP message translation profile with the subscriber information fields to send in the PFCP message, and associate the profile with an APN.

About this task

For the MAG-c to optionally send subscriber information in PFCP messages to the UPF, you can configure a PFCP message translation profile with the subscriber information fields to send in the IE and associate the profile with an APN.

The following subscriber information fields are supported in PFCP messages:

- **pcc-rule-name** – PCC rule name field
- **rat-type** – RAT type field
- **user-location** – user location field
- **default-qos-id** – QCI (4G) or 5QI (5G) field
- **serving-node-id** – MME or AMF IP address field

See the *MAG-c CLI Reference Guide* for more information about the CLI configuration commands.

Procedure

Step 1. Configure a PFCP message translation profile.

```
configure mobile-gateway profile pfc pfc-translate-profile
```

Example

```
A:MAG-c# configure mobile-gateway profile pfc pfc-translate-profile "pfcTranslate Prof"
```

Step 2. Configure the subscriber session information fields to include for the PFCP message translation profile.

```
configure mobile-gateway profile pfc pfc-translate-profile field
```

Example

```
A:MAG-c# configure mobile-gateway profile pfc pfc-translate-profile "pfcTranslate Prof" field rat-type
A:MAG-c# configure mobile-gateway profile pfc pfc-translate-profile "pfcTranslate Prof" field user-location
A:MAG-c# configure mobile-gateway profile pfc pfc-translate-profile "pfcTranslate Prof" field default-qos-id
```

```
A:MAG-c# configure mobile-gateway profile pfcpc pfcpc-translate-profile "pfcpcTranslateProf" field serving-node-id
```

Step 3. Associate the PFCPC message translation profile with an APN.

```
configure mobile-gateway pdn apn pfcpc-translate-tx
```

Example

```
A:MAG-c# configure mobile-gateway pdn apn pfcpc-translate-tx "pfcpcTranslateProf"
```

3.3 MAG-c lawful intercept

The lawful intercept (LI) solution is implemented on the MAG-c and on the User Plane (UP). The MAG-c and the UP share a private key to allow decryption of LI PFCPC IEs. This topic describes the LI implementation, the content of LI notifications, and how to configure LI on the MAG-c.

LI 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.

LI typically consists of the following interfaces, irrespective of the access technology:

- administrative interface – supports LI target provisioning
- information-related interface – provides event information related to subscribers
- contents-of-communications interface – sends mirrored packets to the LI gateway (LIG)

The MAG-c architecture supports administrative and information related interfaces on the MAG-c and the contents-of-communication interface on each UP.

The MAG-c provides a centralized location to provision all LI targets, and instructs the UP to perform LI for specific target subscribers by sending encrypted LI PFCPC IEs through the Sx interface. The MAG-c and the UP share a private key to allow decryption of LI PFCPC IEs.

To allow the LI target to remain anonymous, every subscriber PFCPC session includes encrypted LI PFCPC IEs.

3.3.1 MAG-c LI solution for wireline application

Understand the tools to use and guidelines to follow when configuring MAG-c LI for wireline applications.

For wireline (BNG) application, the following criteria apply for the MAG-c LI:

- Perform all target provisioning for LI on the MAG-c through SSH CLI.
- The MAG-c sends log events related to LI targets via the SNMPv3 interface.
- Each BNG-UP can be configured to send mirrored traffic according to the mirror destination type: SAP, SDP, or IP-UDP SHIM.

Use the following command on the MAG-c to activate an LI target.

```
configure li target
```

For wireline subscribers, use the following command with the **subscriber** keyword to configure the target source. The name (ID) must match the subscriber ID returned from RADIUS, which is VSA Alc-Subsc-ID-Str [11].

```
configure li target source id subscriber name
```

You can also use this command to configure other settings, including the ingress, egress, intercept ID, and session ID.

3.3.2 MAG-c LI solution for FWA applications

Get an overview of the guidelines and steps to configure MAG-c LI for wireless applications.

MAG-c configuration requirements

As defined in 3GPP, LI for fixed wireless application (FWA) can be IRI-only, CC-only, or both. The provisioning of IRI and CC are two separate procedures on the MAG-c. If only IRI or CC provisioning is required, perform the applicable procedure for the IRI or CC only. If both IRI and CC are required, you must configure both.

For each subscriber, perform the provisioning as follows:

- Enable 5G IRI under LI_X1 for 5G. See [Configuring FWA LI IRI for 5G RGs](#).
- Enable 4G IRI using SSH CLI (**configure li mobile-gateway target**) for 4G. See [Configuring FWA LI IRI for 4G RGs](#).
- Enable CC for both 4G and 5G RGs using SSH CLI (**configure li target**). See [Configuring FWA LI CC for 4G and 5G RGs](#).



Note: The LI administrator cannot predetermine if the FWA RG is only 4G capable, only 5G capable, or both. For this reason, Nokia recommends configuring both 4G and 5G LI. This guarantees the lawful interface, regardless of the connected access of the subscriber.

UP configuration requirements

The UP requires a minimal set of LI configurations to support MAG-c LI. The mirror destination ID is a key parameter that the MAG-c sends to the UP. You must configure matching mirror destination IDs on the UP and the MAG-c.

See the *7450 ESS, 7750 SR, 7950 XRS, and VSR OAM and Diagnostics Guide* for more information and configuration guidelines.

3.3.2.1 Configuring FWA LI IRI for 5G RGs

Perform the procedure described in this topic to configure the FWA LI IRI solution for 5G RGs.

About this task

For 5G RGs, the LI solution for FWA is based on 3GPP Release 15 TS 33.127 and TS 33.128. This may include RGs that are 5G capable and have the ability to fallback to 4G radio access.

The following requirements apply when configuring FWA LI IRI for 5G RGs:

- Based on TS 33.128, the LI_X1 interface used to provision LI targets requires an associated TLS server profile. The LI_X1 interface only supports IRI over TLS. See [Configuring FWA LI CC for 4G and 5G RGs](#) for more information about CC.
- The LI_X2 interface is the IRI interface and is also TLS based. This requires a TLS client profile configuration.

Perform the following steps to configure FWA LI IRI for 5G RGs:

Procedure

Step 1. Configure the LI targets using the LI_X1 interface with an associated TLS server profile.

Step 2. Configure the LI_X2 interface with an associated TLS client profile for the IRI interface.

3.3.2.2 Configuring FWA LI IRI for 4G RGs

Perform the procedure described in this topic to configure FWA LI IRI for 4G RGs.

About this task

For 4G (LTE) RGs, the LI solution for FWA is based on 3GPP Release 15 TS 33.107 and TS 33.108.

Perform the following steps to configure FWA LI IRI for 4G RGs:

Procedure

Step 1. Use SSH and the following CLI command to configure the IRI destination.

```
configure li mobile-gateway df-peer id df2-addr addr df2-port port
```

Step 2. Associate the FWA LI target type (for example, IMSI) and the ID (for example, IMSI number) with the IRI (DF2) peer.

```
configure li mobile-gateway target type id value peer df-peer-id
```

Step 3. Optional: Use TLS to enable the IRI interface. The IRI interface uses the TPKT protocol based on TS 33.108.

3.3.2.3 Configuring FWA LI CC for 4G and 5G RGs

Perform the procedure described in this topic to configure the FWA LI CC solution for both 4G and 5G RGs.

About this task

Call Content (CC) data-packet mirroring for both 4G and 5G LI uses the same configuration. The following apply for CC:

- CC provisioning for FWA subscribers on the MAG-c is through SSH CLI. You can specify either the International Mobile Subscriber Identity (IMSI) or Mobile Station International Subscriber Directory Number (MSISDN) as the target for both 4G and 5G subscribers.
- The MAG-c instructs the UP to perform LI on the subscriber session via the Sx interface. Each UP sends LI mirrored packets according to the configured mirror destination type: SAP, SDP, or IP-UDP SHIM.

Perform the following steps to configure LI for both 4G and 5G RGs on the MAG-c using SSH CLI:

Procedure

Step 1. Configure the LI target source with the **imsi** or the **msisdn** option to enable CC LI for a specific FWA subscriber.



Note: Use the **imsi** option for both IMSI and SUPI.

Example

```
A:MAG-c# configure li target source id imsi imsi-number
A:MAG-c# configure li target source id msisdn msisdn-number
```

Step 2. Configure additional settings, including the ingress, egress, intercept ID, and session ID.

```
configure li target source
```

3.3.2.4 Additional information for 3GPP IRI messages

Use the following command to customize the 3GPP IRI message for 4G and 5G. For 4G, the command customizes the operator ID within a 4G IRI message, as defined in 3GPP TS 33.107. For 5G, the command customizes the domain ID (DID) field, as defined in ETSI 103 221-2.

```
configure li mobile-gateway operator-id
```

Use the following command to change the 3GPP 4G IRI message network element ID (NEID), as defined in 3GPP TS 33.107.

```
configure li mobile-gateway iri ne-id
```

For 5G, use the following command to change the interception point ID (IPID), as defined in ETSI 103 221-2.

```
configure li mobile-gateway 5g-iri-ip-id
```

The IPID uses TLV type 7 and the encoding for the different values is expressed as a sub-TLV within the TLV as follows.

Table 3: IPID encoding

| | | | | | |
|--------|----------|---|------|--------|-------|
| Type=7 | Length=x | Value | | | |
| | | <table border="1"> <tr> <td>Type</td> <td>Length</td> <td>Value</td> </tr> </table> | Type | Length | Value |
| Type | Length | Value | | | |

The length field is two bytes and represents the length of the TLV.

The value field is two bytes and the type of the sub-TLV in the value field can have the following values:

- Type 0 – reserved
- Type 1 – IPv4 address
- Type 2 – IPv6 address

- Type 3 – string

The value of the sub-TLV in the value field represents an IPv4 address, IPv6 address, or customized string.

Example: IPID encoding for IPv4 address 18.52.86.255 (hexadecimal 12 34 56 FF)

Table 4: IPID encoding example

| | | | | | |
|--------|----------------|---|--------|----------------|-------------------|
| Type=7 | Length=x | Value | | | |
| | | <table border="1"> <tr> <td>Type=1</td> <td>Length=4 bytes</td> <td>Value=12 34 56 FF</td> </tr> </table> | Type=1 | Length=4 bytes | Value=12 34 56 FF |
| Type=1 | Length=4 bytes | Value=12 34 56 FF | | | |

3.3.3 Alternative MAG-c LI solution through the UP

It is possible to provision LI targets on the UP, although it is not recommended for a number of reasons. Users must understand the risks and requirements before considering this option.



Note:

Although Nokia does not recommend it, you can provision the LI target directly on the UP by using the subscriber ID on the UP. However, each time a subscriber logs on to the MAG-c, the UP assigns the subscriber a different subscriber ID. The following methods help the LIG identify the UP where the subscriber is located and the new subscriber ID on the UP.

Using RADIUS accounting messages is one method to help to locate the subscriber and subscriber ID on the UP. Use the **up-info**, **up-subscriber-id**, and **subscriber-id** commands in the following context to configure the MAG-c to include RADIUS attributes in the accounting messages.

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

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.

In addition to using RADIUS, the MAG-c also reports the subscriber and LI events through SNMPv3 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*.

3.3.3.1 Subscriber ID and IP address notifications for LI meditation devices

Review these guidelines and options for enabling MAG-c notifications to the LI mediation gateway about LI and subscriber events.

The MAG-c notifies the LI mediation gateway about the LI and subscriber events. Some key parameters in the notifications include the BNG-UP subscriber ID and the BNG-UP IP address. The LI mediation gateway uses the key parameters 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 example displays a MAG-c log.

Example: 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 example shows VSAs in a MAG-c RADIUS accounting message.

Example: VSAs in a MAG-c RADIUS accounting message

```
Alc-Subsc-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 prepends `_cups_` to the auto-generated subscriber ID; for example, `_cups_549`.

See *7750 SR and VSR BNG CUPS User Plane Function Guide* for more information about the BNG-UP LI target provisioning and LI packet mirroring.

See *7450 ESS, 7750 SR, 7950 XRS, and VSR System Management Guide* for more information about LI access through SNMPv3.

3.4 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.

3.4.1 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.

3.4.2 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:MAG-c>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 (I2-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 following applies:

- 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:MAG-c>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, as follows:

- 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-UP 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](#)

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

3.4.3 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, L2 circuit, MAC address>, for IPoE and 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 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 (interface ID) as defined in RFC 8415 and option 37 (remote ID) as defined in RFC 4649. The enterprise number 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: <l2-circuit, source MAC address> or <l2-circuit, source MAC address, source IP address>
- PPPoE: <l2-circuit, source MAC address, session ID> or <l2-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](#).

3.4.4 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 following context.

```
configure mobile-gateway profile bng entry-point entry subscriber-identification multi-session
```

Related topics

[Subscribers](#)

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

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

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 combination. In this case, two sessions with the same Layer 2 circuit 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. When 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 nonresilient 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 nonresilient 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](#)

3.4.6 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 following context

```
clear mobile-gateway bng
```

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-gateway bng** context.

3.4.6.1 Enabling session lockout

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

Procedure

Step 1. Define a session lockout profile.

```
configure mobile-gateway profile bng session-lockout-profile
```

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.

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


Session lockout is enabled for the sessions that match a BNG EP entry with a referenced session lockout profile.

3.4.7 IPoE

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

The following 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 control plane 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.
 - Use the following command to configure the MAG-c to process an upstream data packet as the trigger.

```
configure mobile-gateway profile bng entry-point entry ipoe allow-data-trigger
```

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](#)

3.4.7.1 Static IPoE sessions

MAG-c supports static IPoE session creation based on the configured trigger parameters.

The method for creating a static session is similar to creating a data-triggered session. The MAG-c evaluates the trigger parameters for the static IPoE session in BNG EP, and authenticates through ADB and AAA.

Static IPoE sessions do not support DHCP and DHCPv6; MAG-c ignores and discards DHCP/DHCPv6 packets from static IPoE sessions.

Use the commands in the following context to configure a static IPoE session.

```
configure mobile-gateway pdn bng static-session
```


3.4.8 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 following 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 PPPoE session ID within the allocation scope, starting with one. However, when randomization is configured, the MAG-c allocates a random unique PPPoE 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 following context.

```
configure mobile-gateway profile bng pppoe-profile
```

- 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 aligning 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 aligning 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 the 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 following context.

```
configure mobile-gateway profile authentication-database entry pppoe
```

- Provide the delay via the Alc-PPPoE-PADO-Delay VSA during RADIUS authentication.

The following figure 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



sw2083

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 false** command in the **config>mobile>profile>adb>entry** context.

```
configure mobile-gateway profile authentication-database entry
```

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 following context.

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

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 following context is enabled.


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

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.

The following figure shows the L2TP LAC network components.

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 L2TP tunnels to the BNG-UP via the PCFP 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.

The following figure shows an example of a LAC-enabled PPPoE session setup flow.

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 L2TP 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 L2TP tunnels provided via authentication is not supported.

To configure a local L2TP group, use the **l2tp-group** command in the following context.

```
configure mobile-gateway profile bng
```

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 following context.

```
configure mobile-gateway profile pfcpc pfcpc-profile
```

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 or the tunnel setup, or both 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 following context or their equivalents in the authentication attributes.

```
configure mobile-gateway profile bng l2tp-group
```

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](#)

[Session keys and anti-spoofing](#)

[Authentication](#)

[Address assignment](#)

3.5 Fixed wireless access sessions

Fixed wireless access (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.

3.5.1 Introduction to FWA sessions and terminology

FWA sessions obtain network connectivity through a mobile RAN, for example 4G (e-)UTRAN or 5G NR, and therefore, they have the following important differences from fixed access sessions:

- FWA sessions perform the initial setup out-of-band and signal directly to the MAG-c, which acts as a 4G SGW-C + PGW-C (combined), or a session management function (SMF), whichever is applicable.
- FWA sessions have tunneled IP connectivity from the client device to the FWA-UP, which acts as a 4G SGW-U + PGW-U (combined), or 5G UPF, whichever is applicable.

FWA sessions are BNG subscriptions, which means they have the same service requirements as fixed access sessions. QoS, service selection, pool management, and charging are all 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.

APN

An Access Point Name (APN) identifies a specific PDN, but it 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:

- Network Identifier (NI): defines the network (PDN).
- 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.

DNN

A Data Network Name (DNN) is the 5G equivalent to an APN.

| | |
|----------------|---|
| GPSI | Similar to SUPI, the Generic Public Subscription Identifier (GPSI) is introduced in 5G to contain public identifiers, such as MSISDN, and it is also ready to support other type of identifiers. For FWA, the GPSI always contains an MSISDN. |
| IMEI | An International Mobile Equipment Identity (IMEI) is 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. |
| IMSI | An International Mobile Subscriber Identity (IMSI) is 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">• Mobile Country Code (MCC): uniquely identifies a country• Mobile Network Code (MNC): uniquely identifies a mobile network within the country• Mobile Subscriber Identification Number (MSIN): a second globally unique number identifying the subscriber |
| MSISDN | For mobile subscriptions, the Mobile Station International Subscriber Directory Number (MSISDN) is 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. |
| NSSAI | Network Slice Selection Assistance Information (NSSAI) is a set of S-NSSAIs that is applicable to a UE. A UE can connect to a single S-NSSAI or to multiple S-NSSAIs. |
| PDN | A Packet Data Network (PDN) is a network the subscriber connects to; for example, the public Internet. |
| PEI | A Permanent Equipment Identifier (PEI) is introduced in 5G that can contain an IMEI and is also ready to support other types of identifiers. For FWA, the PEI always contains an IMEI. |
| RAB | A Radio Access Bearer (RAB) is a radio channel between a UE and the RAN. Multiple RABs per PDN session can exist. |
| RAN | The Radio Access Network (RAN) is antennas providing wireless connectivity. |
| S-NSSAI | Single Network Slice Selection Assistance Information (S-NSSAI) identifies a single network slice, which is a group of mobile functions such as SMF, UPF, and PCF that together provide specific capabilities and characteristics. For example, an FWA slice may contain all FWA UPFs and SMFs, as well as the necessary supporting functions such as CHF or PCF. |
| SUPI | A global unique Subscription Permanent Identifier (SUPI) is introduced in 5G that can contain an IMSI and is also ready to support other type of identifiers. For FWA the SUPI always contains an IMSI. |
| UE | The User Equipment (UE) is an end device used by the subscriber. The UE and RG can be integrated in one device or can be two separate devices. |

Related topics[Residential Gateway models](#)

[Service selection](#)
[Fixed access sessions](#)

3.5.2 Residential Gateway models

For FWA, there are two common models for an 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



sw1252

Related topics

[PAP/CHAP authentication](#)

3.5.3 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 following context.

```
configure mobile-gateway pdn apn fixed-wireless-access
```

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

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

Related topics

[Service selection](#)

[PAP/CHAP authentication](#)

3.5.4 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](#)

3.5.5 FWA-UP selection

Because FWA session setup messages are sent out-of-band, an FWA session is not automatically tied to an FWA-UP as is the case with fixed BNG access.

To select an FWA-UP in a deployment using FWA-UP resiliency, see [Session-to-FSG mapping and load-balancing](#).

To select an FWA-UP in a non-resilient deployment, the following configuration is required:

- Use the following command to define a UP peer list.

```
configure mobile-gateway profile pfc up-peer-list
```

- Use the following command to configure all FWA-UPs as peers in the UP peer list.

```
configure mobile-gateway profile pfc up-peer-list peer
```

The IP address to use as parameter in the **peer** command is the source IP address of the PFCP association of the FWA-UP.

- Use the following command to set the FWA-UP selection to **true** for each peer.

```
configure mobile-gateway profile pfc up-peer-list peer upf-selection
```

- Use the following command to configure the APNs that are supported on the peer.

```
configure mobile-gateway profile pfc up-peer-list peer apn
```

- Use the following command to make a reference to the configured UP peer list.

```
configure mobile-gateway pdn up-peer-list
```

During setup, the FWA session 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.

Additionally, the MAG-c supports a generic policy-based UPF selection. This can include different selection criteria, such as APN and location, where the location can be specified as a list of tracking area identities (TAIs), a list of e-UTRAN cell global identification (ECGI) and NR cell global identification (NCGI) values, or both. See [Configuring policy-based FWA-UP selection](#).

For FWA-UP selection, the MAG-c does not consider a UP for which the PFCP association or path fails, including a headless state, even if the UP is configured under the UP peer list. The MAG-c terminates FWA-UP sessions for which the association or path is down, not including a headless state, and sets a reactivation cause. The affected FWA RGs are expected to reconnect using a new session establishment procedure and chose a new non-failed UP, thereby providing a basic UP resiliency mechanism.

3.5.5.1 Configuring policy-based FWA-UP selection

The MAG-c supports a generic policy-based FWA-UP selection with different selection criteria, such as APN and location. The location can be specified as a list of TAIs, as a list of ECGI and NCGI values, or both.

About this task

The following procedure configures a TAI list, an ECGI and NCGI list, or both as selection criterion for policy-based FWA-UP selection.

Procedure

Step 1. Create a TAI list, an ECGI and NCGI list, or both.

Create the TAI lists by configuring a range of TAIs or individually listed TAIs.

```
configure mobile-gateway profile list tai-list
configure mobile-gateway profile list tai-list tai
```

Create the ECGI and NCGI list by configuring a range of ECGI and NCGI values or individually listed ECGI and NCGI values.

```
configure mobile-gateway profile list cell-id-list
```

Step 2. Define group labels for policy-based FWA-UP selection.

Start a policy editing session to define group labels and then save your changes.

```
configure mobile-gateway profile node-selection begin
configure mobile-gateway profile node-selection group-label
configure mobile-gateway profile node-selection commit
```

Step 3. Configure one or more group labels for FWA-UP peers in the FWA-UP peer list. A maximum of five group labels can be associated with a FWA-UP peer.

```
configure mobile-gateway profile pfc up-peer-list peer group-label
```

Step 4. Configure the target profile for FWA-UP node selection. Target profile entries can be configured with group labels that the MAG-c uses to determine eligible FWA-UP nodes by matching these to the group labels in the FWA-UP peer list that are configured in step 3.

Start a policy editing session and use the **target-profile** command and commands in its context to configure the target profile.

```
configure mobile-gateway profile node-selection begin
configure mobile-gateway profile node-selection target-profile
```

...

Each entry in the target profile has a selection priority based on the numeric value of its associated ID. If the number of eligible FWA-UPs for an evaluated entry is below its configured threshold, the entry is not considered for FWA-UP selection and the next entry is evaluated. If a valid entry is found, a FWA-UP from the eligible list of FWA-UP nodes for that entry is selected to anchor the session. To ensure equal distribution of sessions amongst the list of eligible FWA-UPs for the entry, the MAG-c assigns sessions to FWA-UPs on a least-loaded-first basis, where the FWA-UP with the least load is picked to anchor a new session. When multiple FWA-UPs have the same least load, the MAG-c randomly selects one of those FWA-UPs to anchor the new session.

- Step 5.** Configure the client profile with one or more entries that match on client parameters (for example, APN, TAI list, APN and TAI list, or ECGI and NCGI list) to determine the target profile for the session.

Use the **client-profile** command and the commands in its context to configure the client profile, and then save your changes.

```
configure mobile-gateway profile node-selection client-profile
...
configure mobile-gateway profile node-selection commit
```

Each entry in the client profile has a selection priority processing based on the numeric value of its associated ID. The parameters configured per entry (for example, APN, TAI list, APN and TAI list, or ECGI and NCGI list) are applied collectively. The action for the matching entry sets the target profile for the session. If no matching entry is found, the configured fallback default action sets a default target profile for the session. If no fallback default action is configured under the client profile, the session fails.

- Step 6.** Configure the client profile to apply on the SMF.

```
configure mobile-gateway pdn up-selection pgw-u-client-profile
```

Example: TAI-based FWA-UP selection

The following example shows the configuration for a TAI-based FWA-UP selection.



Note:

Nokia recommends to also configure throttling limits for the `tmnxMobProfNodeSelTrgtSugstRevrt` and `tmnxMobProfNodeSelTrgtSugstRebal` log events when the same FWA-UP is shared across multiple target profiles.

```
A:MAG-c# configure mobile-gateway profile list
tai-list "tai-list-1"
  tai mcc 206 mnc 01 tac 0x000001
  tai mcc 206 mnc 01 tac 0x00000b
exit
tai-list "tai-list-2"
  tai mcc 206 mnc 01 tac-range-start 0x123410 tac-range-end 0x123419
exit

A:MAG-c# configure mobile-gateway profile node-selection
begin
group-label "site_1"
description "Label for UPFs serving TAI 1"
exit
```

```
group-label "site_2"
  description "Label for UPFs serving TAI 2"
  exit
commit
exit

A:MAG-c# configure mobile-gateway profile pfc up-peer-list "up-peer-list1"
peer 2.2.2.2
  group-label "site_1"
peer 3.3.3.3
  group-label "site_1"
peer 4.4.4.4
  group-label "site_2"
peer 5.5.5.5
  group-label "site_2"

A:MAG-c# configure mobile-gateway profile node-selection
begin
client-profile "client_profile1"
  entry 1
    match
      tai-list "tai-list-1"
    exit
    action
      set-target-profile "target_prof_1"
    exit
  exit
  entry 2
    match
      tai-list "tai-list-2"
    exit
    action
      set-target-profile "target_prof_2"
    exit
  exit
  default-action
set-target-profile "default_target_prof"
  exit
exit

target-profile "target_prof_1"
  entry 1
    threshold 2
    match
      group-label "site_1"
    exit
  exit
  entry 2
    threshold 2
    match
      group-label "site_2"
    exit
  exit
exit

target-profile "target_prof_2"
  entry 1
    threshold 2
    match
      group-label "site_2"
    exit
  exit
  entry 2
    threshold 2
```

```

        match
            group-label "site_1"
        exit
    exit
exit
target-profile "default_target_prof"
description "default for non-matching TAIs"
exit
commit
exit
A:MAG-c# configure mobile-gateway pdn 1
up-peer-list "up-peer-list1"
up-selection
pgw-u-client-profile "client_profile1"
exit

```

Example: FWA-UP selection based on cell ID

```

A:MAG-c# configure mobile-gateway profile list
cell-id-list "cell-id-list-1"
plmn mcc 206 mnc 01
eci 0x0000011 end 0x0000014
exit
nci 0x00120001 end 0x001200F9
exit
exit
exit
A:MAG-c# configure mobile-gateway profile node-selection
begin
group-label "CoLo-site_1"
exit
commit
exit
A:MAG-c# configure mobile-gateway profile pfc up-peer-list "up-peer-list1"
peer 2.2.2.2
group-label "CoLo-site_1"
peer 3.3.3.3
group-label "CoLo-site_1"
peer 4.4.4.4
group-label "site_1"
peer 5.5.5.5
group-label "site_1"
peer 6.6.6.6
group-label "site_2"
A:MAG-c# configure mobile-gateway profile node-selection
begin
client-profile "client_profile1"
entry 1
match
cell-id-list "cell-id-list-1"
exit
action
set-target-profile "tp-cell-id-primary-selection"
exit
exit
entry 2
match
tai-list "tai-list-1"
exit

```

```

        action
            set-target-profile "tp-tai-primary-selection"
        exit
    exit
    default-action
set-target-profile "default_target_prof"
    exit
    exit

    target-profile "tp-cell-id-primary-selection"
        entry 1
            threshold 2
            match
                group-label "CoLoSite_1"
            exit
        exit
    exit
exit

target-profile "tp-tai-primary-selection"
    entry 1
        threshold 2
        match
            group-label "site_1"
        exit
    exit
    entry 2
        threshold 2
        match
            group-label "site_2"
        exit
    exit
    exit
target-profile "default-target-prof"
    description "default for non-matching TAIs"
    exit
commit
exit

A:MAG-c# configure mobile-gateway pdn 1
up-peer-list "up-peer-list1"
up-selection
    pgw-u-client-profile "client_profile1"
    exit

```

3.5.5.2 Static FWA-UP selection per IMSI

The MAG-c supports selecting a preferred FWA-UP for an IMSI via an ADB match.

The preferred FWA-UP from the ADB match overrides the FWA-UP resulting from the normal policy-based FWA-UP selection, based on matches in **client-profile** and **target-profile**.

The preferred FWA-UP must be part of the candidate FWA-UP pool available to policy-based FWA-UP selection. Associate the FWA-UP with one or more **group-label** values that match in the selected **target-profile** during the FWA-UP selection process. See [Configuring policy-based FWA-UP selection](#) for information about using **group-label**, **target-profile**, and **client-profile**.

To configure static FWA-UP selection per IMSI, use the following commands:

- Configure the IMSI match attribute in the ADB match configuration.

```
configure mobile-gateway profile authentication-database match
```

- Configure the IMSI match criterion in the ADB entry.

```
configure mobile-gateway profile authentication-database entry match imsi
```

- Configure the preferred FWA-UP for the matched IMSI.

```
configure mobile-gateway profile authentication-database entry fixed-wireless-access prefer-up
```

3.5.5.3 Session rebalancing and revert operations

The MAG-c supports a rebalancing operation to restore an equal distribution of sessions among the eligible set of highest-priority UPs, that is, the set of UPs belonging to the highest-priority entry in the target profile for a TAI list or ECGI and NCGI list. Clearing a configurable percentage of sessions on each of the currently active UPs in the highest-priority entry of the target profile via CLI or via a network element management system, such as NSP, triggers the rebalancing operation. When a new UP becomes active and belongs to the highest-priority entry in the target profile for a TAI list or ECGI and NCGI list, a log event is generated to signal that a rebalance operation should be performed.

When the number of surviving UPs in the highest-priority entry of the target profile falls below the configured threshold, sessions may be anchored on UPs that are not highest-priority UPs. When the number of surviving UPs increases up to the configured threshold, a revert operation can be performed to reanchor the sessions from the lower-priority UPs to the highest-priority UPs. Clearing a configurable percentage of sessions on each of the lower-priority UPs for the target profile via CLI or via a network element management system, such as NSP, can trigger the revert operation. A log event is generated to signal that a revert operation should be performed.

Use the following command to trigger the rebalance or revert operation.

```
clear mobile-gateway pdn bearer-context target-profile
```

Example: Configuration details

```
A:MAG-c>config>mobile>profile>node-selection# info
-----
      group-label "site1"
      exit
      group-label "site2"
      exit
      client-profile "pgw_tai_client_prof"
        entry 1
          match
            tai-list "tai1"
          exit
          action
            set-target-profile "target_prof_tail"
          exit
        exit
      default-action
      exit
    exit
  target-profile "target_prof_tail"
    entry 1
      threshold 3
      match
        group-label "site1"
```

```

        exit
    exit
    entry 2
        threshold 2
        match
            group-label "site2"
        exit
    exit
exit
-----

```

Example: Revert and rebalance triggers

```

A:MAG-c# clear mobile-gateway pdn 1 bearer-context target-profile "target_prof_tail"
action rebalance up-peer 30.0.1.188
A:MAG-c# clear mobile-gateway pdn 1 bearer-context target-profile "target_prof_tail"
action revert up-peer 30.0.1.241

A:MAG-c# clear mobile-gateway pdn 1 bearer-context target-profile "target_prof_tail"
action rebalance up-peer 30.0.1.188 percentage 50
A:MAG-c# clear mobile-gateway pdn 1 bearer-context target-profile "target_prof_tail"
action revert up-peer 30.0.1.241 percentage 50

A:MAG-c# clear mobile-gateway pdn 1 bearer-context session-deletion-rate 100 target-
profile "target_prof_tail" action rebalance up-peer 30.0.1.188 percentage 50 A:MAG-
c# clear mobile-gateway pdn 1 bearer-context session-deletion-rate 100 target-profile
"target_prof_tail" action revert up-peer 30.0.1.241 percentage 50

```

3.5.6 Address signaling methods and deferred allocation

The 3GPP specifications define the following 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 for 4G, or the N1 PDU Session Establishment Accept message for 5G.
- **via DHCP**
The IPv4 address is signaled after the GTP or PDU session setup completes.

The method used depends on the following parameters and configuration, listed in order of precedence (item 1 being highest priority). Each parameter or configuration can be absent, or can have the value NAS or DHCP:

1. the RADIUS Alc-FWA-IPv4-Signaling-Method VSA
2. the configuration of the **ipv4-signaling-method** command in the following context

```
configure mobile-gateway profile authentication-database entry fixed-wireless-access
```

3. 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.
4. the configuration of the **default-ipv4-signaling-method** command in the following context of the selected APN where the session is created

```
configure mobile-gateway pdn apn fixed-wireless-access
```

5. 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: In case of 4G, any allocated SLAAC prefix is also 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 5G, the SLAAC prefix is not signaled to the AMF.

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](#)

3.5.7 QoS

FWA sessions generally support the generic BNG QoS functionality and are additionally subject to specific mobile QoS constructs. In mobile networks, the concept of a bearer defines 4G QoS functionality and a QoS flow defines the 5G QoS.

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](#) and [5G 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.

3.5.7.1 4G QoS attributes

The concept of a bearer defines the 4G FWA QoS functionality in mobile networks. Sessions support default bearers of the non-GBR type, as determined by specific parameters, and only APN-AMBR is supported for mobile rate-limiting.

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 a two-step process.

1. The MAG-c retrieves the subscribed values using the following parameters, in order of precedence:

- a. from RADIUS, the 3GPP-GPRS-Negotiated-QoS-Profile attribute

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 functionality, use the following command.

```
configure mobile-gateway profile bng radius-authentication-profile include-attribute gprs-negotiated-qos-profile
```

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

2. The MAG-c determines the authorized values using the following parameters, in order of precedence:

- a. from the PCF, using the `authSessAmbr` and `authDefQos` information elements during the `Npcf_SMPolicyControl_Create` service operation
The MAG-c signals the subscribed QoS values to the PCF in this operation, and the PCF replies with the final authorized parameters. Session AMBR is mapped directly to APN-AMBR in this case.
- b. if the session does not use PCF, or PCF failed but the session continues in ap-continue mode: from the ADB, the QoS profile configured using the following command

```
configure mobile-gateway profile authentication-database entry fixed-wireless-access
default-qos-profiles authorized-4g
```



Note: The QoS profile must reference a profile configured using the following command.

```
configure mobile-gateway profile qos-profile
```

- c. from the subscribed QoS values retrieved in step 1

The authorized QoS parameters are used for the generic mobile QoS functionality. They are installed on the FWA-UP using PFCP QER constructs and 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:

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

These 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.

3.5.7.2 5G QoS attributes

In 5G mobile networks, the QoS flow is a key concept. The MAG-c interacts with the PCF to create and then signal QoS flows, along with corresponding default-bearer information, to the RAN (gNodeB), UPF, and the UE.

ARP, 5QI, and Session-AMBR for the default QoS flow are determined using a two-step process.

1. The MAG-c retrieves the subscribed values using the following parameters, in order of precedence:
 - a. from RADIUS, the 3GPP-GPRS-Negotiated-QoS-Profile attribute
 - b. from the UDM, the 5G QoS profile and `sessionAmbr` information elements
 - c. from the ADB, the QoS profile configured using the following command

```
configure mobile-gateway profile authentication-database entry fixed-wireless-access
default-qos-profiles subscribed-5g
```



Note: The QoS profile must reference a profile configured using the following command.

```
configure mobile-gateway profile qos-5g-profile
```

2. The MAG-c determines the authorized values using the following parameters, in order of precedence:
 - a. from the PCF, using the authSessAmbr and authDefQos information elements during the Npcf_SMPolicyControl_Create service operation
The MAG-c signals the subscribed QoS values to the PCF in this operation, and the PCF replies with the final authorized parameters.
 - b. if the session does not use PCF, or PCF failed but the session continues in ap-continue mode: from the ADB, the QoS profile configured using the following command

```
configure mobile-gateway profile authentication-database entry fixed-wireless-access
default-qos-profiles authorized-5g
```



Note: The QoS profile must reference a profile configured using the following command.

```
configure mobile-gateway profile qos-5g-profile
```

- c. from the subscribed QoS values retrieved in step 1

After interaction with the PCF, the MAG-c creates QoS flows for each 5QI associated with the session. The QoS flow is identified by a QoS Flow Identifier (QFI), which is also downloaded to the UPF and sent on the wire as an extension header to GTP-U. By default, the MAG-c allocates a dynamic QFI for each QoS flow, but it can also use the 5QI itself as a QFI. The latter can be achieved by enabling the following command.

```
configure mobile-gateway profile qfi-mapping-profile 5qi-as-qfi
```

Use the following commands to apply a QFI mapping profile per PDN or per APN.

```
configure mobile-gateway pdn qfi-mapping-profile
configure mobile-gateway pdn apn qfi-mapping-profile
```



Note: The advantage of using 5QI as a QFI is you can map the data plane traffic directly to a specific 5QI. However, the following limitations apply and for these reasons Nokia does not recommend using 5QI as QFI in a production network:

- It is not allowed to modify the 5QI for a QoS flow mid-session. This includes the default 5QI for the default QoS flow.
- It is not possible to use standard or vendor-specific 5QI values above 64.

When QoS flows are created, the MAG-c acting as the SMF signals the QoS flows and the corresponding parameters to the RAN (gNodeB), UPF, and the UE. Toward the RAN and UE this includes parameters that are typically not applicable to SMF or UPF and are passed through as they are; for example, averaging window.

3.5.7.3 Dynamic QoS based on PCC rules

When a session is provisioned with policy and charging control (PCC) rules, the MAG-c enforces QoS as specified in 3GPP TS 23.501 and TS 23.503. The MAG-c supports the following PCC rule parameters for QoS enforcement:

- precedence versus other PCC rules in case of overlapping packet classifiers

- list of service data flows (SDFs) that act as classifiers to associate packets with PCC rules. SDFs can contain the 5-tuple or the type of service (TOS) field classification, and can be bidirectional or unidirectional.
- 5G quality of service identifier (5QI) and allocation and retention priority (ARP) to determine the end-to-end 5G QoS characteristics
- indicator to bind the PCC rule to the default QoS flow, which is created based on the session level 5QI and ARP values. The 5QI and ARP values are ignored for the QoS flow binding.
- indicator to enable reflective QoS
- gate status to allow or disallow the PCC rule to forward traffic
- maximum bit rate (MBR) and guaranteed bit rate (GBR) values
- QoS characteristics specific to a radio access network (RAN) such as priority level, averaging window, maximum data burst volume, and maximum packet loss rate. These are passed as is to the RAN and act as overrides for the QoS characteristics associated with the 5QI.



Note: QoS characteristics for a standard 5QI are defined in 3GPP TS 23.503 table 5.7.4-1.

Based on the provided 5QI and ARP values and the default QoS flow binding, the MAG-c generates QoS flows as specified in [5G QoS attributes](#). For any GBR QoS flows, it sets the guaranteed flow bit rate (GFBR) and the maximum flow bit rate (MFBR) value to the sum of the GBR and MBR values of all PCC rules tied to that QoS flow.

On top of the QoS flow rates, the MAG-c instructs the FWA-UP to enforce any MBR rate on a per PCC-rule level. This MBR is known as the SDF MBR. The FWA-UP combines the rate enforcements as follows:


- Downlink packets tied to a GBR flow are first subjected to the SDF MBR, and subsequently to the GFBR and MFBR.
- Downlink packets tied to a non-GBR flow are first subjected to the SDF MBR, and subsequently to the session aggregate maximum bit rate (AMBR).
- Uplink packets tied to a GBR flow are only subjected to the SDF MBR because the RAN enforces the GFBR and MFBR.
- Uplink packets tied to a non-GBR flow are first subjected to the SDF MBR, and subsequently to the session AMBR.



Note: The Nokia FWA-UP by default uses policing to enforce all UL rates and the DL PCC rule MBR, and shaping to enforce the DL QoS flow rates (GFBR, MFBR and session AMBR).


The following figures show the rate enforcements.

Figure 9: Downlink QoS enforcement on a FWA-UP



sw1449

Figure 10: Uplink QoS enforcement on a FWA-UP



sw1450

The FWA-UP also performs the following based on the PCC rule classifier:

- QoS flow identifier (QFI) marking
- reflective QoS identifier (RQI) marking
- differentiated services code point (DSCP) marking, if configured on the MAG-c

The MAG-c automatically generates a default any-to-any PCC rule that uses the session level 5QI and ARP in addition to any explicit per-session PCC rules. When the creation of the any-to-any rule is disabled, the MAG-c does not automatically install a default rule and drops any traffic that does not match a specific PCC rule.

Perform the following steps to disable the creation of the any-to-any PCC rule:

1. Use the following command to enable 5QI to be used as the QFI value.

```
configure mobile-gateway profile qfi-mapping-profile 5qi-as-qfi
```

2. Provision a PCC rule with the same 5QI and ARP as the session 5QI and ARP.

Related topics

[DSCP marking](#)

3.5.7.4 DSCP marking

In addition to applying base 4G and 5G QoS, Nokia's FWA solution can also apply DSCP marking based on QCI/5QI and ARP values. The FWA-UP remarks the DSCP values as follows:


- In the uplink direction, the FWA-UP marks the DSCP value of the forwarded IPv4 and IPv6 PDU packets. This is useful for QoS treatment based on QCI/5QI in the data network beyond the FWA-UP.
- In the downlink direction, the FWA-UP marks the DSCP value of both the forwarded IPv4 and IPv6 PDU packets and marks the DSCP value of the outer GTP-u header. This is useful to honor 4G/5G QoS properties in the aggregation network between UP and RAN.



Note: If marking is not applied, the outer GTP-u header DSCP uses the DSCP value of the inner PDU packet by default.

As shown in the following figure, in both directions statically preprovisioned QoS treatment is applied based on QCI/5QI and ARP via the DSCP values in Network Elements (NEs) that are typically not 4G or 5G QoS aware.

Figure 11: 4G/5G QoS aware versus non-QoS aware NEs



Do the following to apply DSCP marking on the UP:

1. Use the following command to configure a QCI policy.

```
configure mobile-gateway profile policy-options qci-policy
```

2. Use the following commands to configure per QCI and ARP combination, the DSCP value in the uplink (**in**) and downlink (**out**) directions and to disable DSCP preserve.



Note: The **no dscp-preserve** command specifies that the MAG-c signals the configured DSCP values to the FWA-UP in PFCP messages. See the *MAG-c CLI Reference Guide* for more information about this command.

```
configure mobile-gateway profile policy-options qci-policy qci
```

```
configure mobile-gateway profile policy-options qci-policy qci dscp
configure mobile-gateway profile policy-options qci-policy qci no dscp-preserve
```

3. Apply the QCI policy in the following context.

```
configure mobile-gateway pdn apn qci-policy
```

3.5.8 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 following context.

```
configure mobile-gateway profile bng radius-authentication-profile
```

- The PAP/CHAP credentials are reflected in RADIUS instead of the **password** configured in the following context.

```
configure mobile-gateway profile bng radius-authentication-profile
```

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](#)

3.5.9 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 following context.

```
configure mobile-gateway profile authentication-database entry fixed-wireless-access
```

context. The hold time can be set to a fixed amount of minutes or be based on the longest remaining address lifetime. In both cases, the hold time value is limited to a maximum of 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 a 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 held 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.

3.5.10 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 as follows:

- GTP for 4G
- N1 NAS/UDM SDM for 5G

The PDN type is checked against the local configuration that is done with the following command.

```
configure mobile-gateway pdn apn pdn-type
```

The PDN type acts as an allow-list. If the signaled PDN type does not match the configured PDN type, but the PDN types are compatible (for example, IPv4 is configured but IPv4v6 is signaled), MAG-c allows the FWA session and performs a downgrade to a compatible PDN type (for example, IPv4). If a downgrade is possible to both IPv4 and IPv6 (for example, both IPv4 and IPv6 are configured and IPv4v6 is signaled), by default MAG-c chooses IPv4, 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.

3.5.11 Mobility and idling

FWA supports mobility and idling procedures, and tracks location.

FWA sessions support the following mobility and idling procedures:

- mobility procedures – S1-based handover (4G), X2-based handover (4G), N2-based handover (5G), and Xn-based handover (5G)



Note: During S1-based and N2-based handover procedures where there are RAN requests for an indirect forwarding tunnel, the requests proceed without the installation of an indirect forwarding tunnel on the FWA-UP.

- idling procedures – S1 release (4G) and AN release (5G)
- reactivation and paging procedures – UE triggered service request (4G+5G) and network triggered service request (4G+5G).



Note: Variants of 4G procedures with an SGW change are not supported.

Most of these procedures are required for interoperability with standard RAN deployments, which assumes these procedures are available by default. Handover procedures provide RAN resiliency when a UE or RG uses them to switch to a new antenna if its current antenna fails.

FWA-UP updates for mobility and idling procedures

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

- For mobility procedures, the F-TEID field in the Forwarding Parameter IE is updated.
- For idling procedures, the forwarding rule is removed and replaced by a buffering rule that instructs the FWA-UP to buffer the packets. The FWA-UP is also instructed to notify the MAG-c when buffering a packet, so the MAG-c can start paging procedures.
- When completing reactivation procedures, the buffering rule is removed and replaced by a forwarding rule with an F-TEID. The FWA-UP sends the buffered packets to the new F-TEID.

3.5.11.1 Location tracking

The MAG-c tracks and learns the location of an FWA session using procedures such as X2/Xn-based handover, UE-triggered service request, location reporting, and tracking area update. For 4G, 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. For 5G the MAG-c does not subscribe to specific location updates and only learns location changes as part of other procedures.

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.

Use the following commands to configure location tracking for RADIUS:

- Use the following command to enable location tracking.

```
configure mobile-gateway profile charging bng-charging radius session update-triggers user-location-change
```

- Use the following command to send the user location in the 3GPP-User-Location-Info attribute.

```
configure mobile-gateway profile charging bng-charging radius session include-attribute user-location-info
```

3.5.12 Headless mode for FWA

Headless mode as described in [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 [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 pfcf pfcf-profile path-restoration-time
```

3.5.13 4G and 5G NSA option 3 sessions

[Figure 12: 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 following figure shows 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 12: 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 following context.

```
configure mobile-gateway pdn s5 interface
```

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 following context.

```
configure mobile-gateway pdn s11 interface
```

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 **interface-realm** parameter of the **gtp-c** command in the following contexts.

```
configure mobile-gateway pdn s11 interface
configure mobile-gateway pdn s5 interface
```

Both realms must be configured the same.

- Configure the EPC node name using the **epc-node** command in the **configure mobile-gateway 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.

The following figure shows an example of a 4G FWA session setup.

The authentication flow of the FWA node configured using the **authentication-flow** command in the following context is used to authenticate the session.

```
configure mobile-gateway pdn apn fixed-wireless-access
```

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.

Related topics

[Address signaling methods and deferred allocation](#)


[Address assignment protocols](#)

[QoS](#)

3.5.14 5G standalone sessions

[Figure 14: Basic 5G FWA network](#) shows the elements involved in a basic 5G FWA network with a MAG-c operating as an SMF and an FWA-UP operating as a UPF. The figure highlights the direct communication between the interface elements and the protocols used, if applicable. The figure does not include the N1 (NAS signaling from and to the UE) and N2 (RAN control plane) interfaces because these interfaces may terminate on both AMF and SMF depending on the type of control plane signaling.

Figure 14: Basic 5G FWA network



The following are high-level descriptions of the role of each function:

- **AMF**

The Access and Mobility Function (AMF) orchestrates the basic connectivity of the UE and how this connectivity evolves during the lifetime of a UE (for example, mobility, idling, and paging).

- **SMF**

The Session Management Function (SMF) maintains 5G sessions. It connects to other functions such as AAA, UDM, PCF, and CHF to retrieve and manage subscription, policy (including QoS), and charging data for an active session. The SMF provides a similar function to a SGW-C + PGW-C in 4G, however, in contrast to 4G it performs all the signaling for session management directly. For example, the SMF constructs all N1 (NAS) session-management messages for the UE directly, as compared to 4G, which uses the MME to proxy relevant data over GTP to the UE.

- **UDM**

The Unified Data Management (UDM) provides subscription data for a user. This subscription data is queried by multiple functions, for example, the AMF retrieves access and mobility subscription data and the SMF retrieves session management subscription data. Additionally, functions such as the SMF register themselves with the UDM to indicate they are an active function for a specific UE or session.

- **NRF**

The Network Repository Function (NRF) keeps a repository of all 5G functions. When a function becomes available, for example the SMF, it registers itself with the NRF. Subsequently, other functions such as the AMF can query the NRF to discover the functions with which they need to communicate.

- **PCF**

The SMF queries the Policy Control Function (PCF) at setup to retrieve policy information applicable to a session. This policy information can be QoS, ACL, or charging related. Some policy information is directly applied by the SMF, but most is passed to other functions to enforce, such as the UPF, gNodeB, UE, or CHF. The PCF can also request updated policies on some triggers, such as location changes or revalidation timeouts and make direct requests for a policy update without requiring a trigger.

- **CHF**

The Charging Function (CHF) is used for both online and offline charging. In both cases CHF manages charging based on Rating Groups (RGs), which can either represent an entire session or a subset of traffic. For offline charging, the CHF requests to signal statistics upon reaching a specific time or volume threshold; however, it does not enforce any limits. For online charging, the CHF additionally enforces quota limits that trigger the SMF and UPF to take immediate action without communicating to the CHF; for example, to stop forwarding for the RG.

- **UPF**

The User Plane Function (UPF) provides connectivity to an IP network, enforcing data plane policies such as QoS and ACLs, reporting statistics, and quotas.

Enabling 5G FWA sessions requires configuration of the following:

- SMF PDUSession services
- AMF client communication services
- 5G UDM function
- FWA sessions

3.5.14.1 SMF PDUSession service server configuration

To enable 5G FWA sessions, do the following:

1. Use the following command to configure an SMF PDUSession service server instance on the MAG-c, to handle service calls from the AMF to the SMF.

```
configure mobile-gateway pdn sba-server-services nsmf-pdusession
```

2. Configure the service instance with the following required elements:

- an interface on which the MAG-c listens for incoming connections from the AMF
- the N3 interface realm that is sent to the FWA UP as the realm (service) where N3 GTP-u tunnels are terminated



Note: When 4G and 5G interworking is required, configure the **interface-realm** option for the S11 interface using the following command.

```
configure mobile-gateway pdn s11 interface gtp-c
```

See [4G and 5G NSA option 3 sessions](#) for more information about configuring the S11 interface.

3. Optionally provision the following for the server instance:

- An HTTP2 profile to influence generic HTTP/2 and SBI parameters; see [HTTP/2, OpenAPI, and Python](#) for more details.

- An API prefix to be part of the service URI; see [URI construction](#) for more information about how URIs are constructed.
- An FQDN to identify the service by other NF peers. For example, the AMF can use this when constructing URIs or when discovering the SMF using DNS.
- Multiple attributes to send to the NRF when registering the service:
 - PLMNS using the following command.

```
configure mobile-gateway pdn sba-server-services nsmf-pdusession allowed-plmns
```

- slices using the following command.

```
configure mobile-gateway pdn sba-server-services nsmf-pdusession allowed-slices
```

See [NF registration](#) for more information about how services and related attributes register with the NRF.

4. Use the following command option to enable the service instance.

```
configure mobile-gateway pdn sba-service-realm server-service service-instance
```

3.5.14.2 AMF client communication service configuration

In addition to the SMF server instance, you must configure the AMF communication client service. This service generates service calls from the SMF to the AMF, to forward N1 or N2 messages to the UE or RAN.

To configure the AMF client service, do the following:

1. Use the following command to configure the AMF service.

```
configure mobile-gateway pdn sba-client-services amf-client namf-comm
```

2. Use the following command to configure an interface for the service instance that the MAG-c uses to communicate with the AMF.

```
configure mobile-gateway pdn sba-client-services amf-client namf-comm interface
```

3. Optionally configure the following for the client service instance:

- An FQDN to identify the service by other NF peers. For example, when generating callback functions the SMF can use the FQDN in the URI, rather than the IP address; see [URI construction](#) for more information about how URIs are constructed.
- An NF ID list to specify which NF instances can be used for the AMF service; see [Enabling discovery based on a local NF ID list](#) for more information. This command is required if the AMF is not dynamically discovered using NRF.
- An HTTP2 profile to influence generic HTTP/2 and SBI parameters; see [HTTP/2, OpenAPI, and Python](#) for more information.
- An AMF profile to handle AMF failure handling; see section [Failure handling](#) for more information.

- An N1 profile configured using the following command.

```
configure mobile-gateway profile n1-profile
```

Use the following commands to configure the N1 profile to influence N1 signaling:

- The N1 profile timer governs retransmit behavior for an N1 PDU SESSION MODIFICATION COMMAND message.

```
configure mobile-gateway profile n1-profile n1-t3591
```

- The N1 profile timer specifies retransmit behavior for an N1 PDU SESSION RELEASE COMMAND message.

```
configure mobile-gateway profile n1-profile n1-t3592
```

- The **cause-code** commands in the following context specify which error codes are signaled to the UE and AMF for several session removal and rejection cases.

```
configure mobile-gateway profile n1-profile
```

4. This service generates service calls from SMF to AMF, to forward N1 or N2 messages to the UE or RAN. Use the following command option to subsequently enable this service instance.

```
configure mobile-gateway pdn sba-service-realm client-service service-instance
```

3.5.14.3 UDM function configuration

In 5G, the MAG-c acting as the SMF directly communicates with the UDM function to retrieve session subscription data, such as subscribed default QoS. The MAG-c subscribes to any subscription data changes to apply these changes. Use the MAG-c SBI IE Search Tool to see a list of supported attributes.

This differs from 4G where a similar function is provided by the HSS, but the PGW does not retrieve this context directly. In 4G, the MME retrieves the subscription data from HSS and signals this via GTP messages to the PGW.

3.5.14.3.1 Enabling the UDM SDM client service

About this task

The UDM SDM client service enables UDM-based subscription retrieval.

Procedure

- Step 1.** Configure a UDM SDM client service.

```
configure mobile-gateway pdn sba-client-services udm-client nudm-sdm
```

Step 2. Configure an interface for the service instance over which the MAG-c communicates with the UDM.

```
configure mobile-gateway pdn sba-client-services udm-client nudm-sdm interface
```

Step 3. Configure the following optional parameters:

- An FQDN to identify the service by other NF peers. For example, when generating callback functions, the MAG-c can use the FQDN in the URI, instead of the IP address. See [URI construction](#) for more information about how URIs are constructed.
- An NF ID list to specify NF instances that can be used for the UDM function. See [Enabling discovery based on a local NF ID list](#) for more information. This configuration is required if the UDM is not dynamically discovered using NRF.
- An HTTP2 profile to influence generic HTTP/2 and SBI parameters. See [HTTP/2, OpenAPI, and Python](#) for more information.
- A UDM SDM profile that contains parameters on how to handle UDM failure handling. See [Failure handling](#) for more information.

Step 4. Enable this service instance.

```
configure mobile-gateway pdn sba-service-realm client-service
```

3.5.14.3.2 Enabling the UDM UECM client service

About this task

The UDM UECM (UE Context Management) service is used to register the MAG-c as an SMF providing service for the UE.

Procedure

Step 1. Configure a UDM UECM client service.

```
configure mobile-gateway pdn sba-client-services udm-client nudm-uecm
```

Step 2. Configure an interface for the service instance over which the MAG-c communicates with the UDM.

```
configure mobile-gateway pdn sba-client-services udm-client nudm-uecm interface
```

Step 3. Configure the optional parameters as described for the UDM client service. See [Enabling the UDM SDM client service](#).

Step 4. Enable this service instance.

```
configure mobile-gateway pdn sba-service-realm client-service
```

3.5.14.3.3 Configuring a 5G QoS profile for the UDM failure handling

About this task

UDM communication supports generic SBI failure handling as specified in [Failure handling](#). When the **ap-continue** action is used, session setup continues without SDM data, an SDM subscription, or a UECM registration. However, a locally configured subscribed 5G QoS profile must be present to provide fallback default QoS values. If this profile is not present, session setup fails anyway.

Procedure

Step 1. Configure a new 5G QoS profile.

```
configure mobile-gateway profile qos-5g-profile
```

Step 2. In the 5G QoS profile, configure the **5qi**, **arp**, **dl-ambr**, and **ul-ambr** parameters as needed.

Step 3. Apply the profile in an ADB entry that is used for the 5G FWA sessions. For more information about matching sessions to ADB entries, see [BNG EP and ADB lookup](#).

```
configure mobile-gateway profile authentication-database entry fixed-wireless-access  
default-qos-profiles subscribed-5g
```




Note: If authorized QoS parameters are also provisioned, for example, via PCF or ADB, those parameters take precedence over the subscribed QoS profile.

3.5.14.4 FWA session setup for 5G

Typically a UDM (see [UDM function configuration](#)) and PCF (see [PCF-based policy management](#)) are configured for a 5G session, in addition to the SMF PDUSession service and the AMF client communication service. The following figure shows the basic 5G FWA session setup when all these components are configured.

Figure 15: 5G FWA session setup



The following describes the message flow and actions to set up a 5G FWA session:

1. An initial registration sequence is initiated between the UE and the AMF. This procedure is orthogonal to the SMF because it does not involve any session-management functionality. Upon completion of the registration, an FWA UE requests creation of a PDU session to forward traffic. To do this, it generates a PDU Session Establishment Request NAS message and signals this to the AMF. Upon receipt of this message, the AMF may obtain additional subscription parameters from the UDM and determines which SMF should handle the session based on local configuration. For example, a specific FWA slice (S-NSSAI) could be used, with which the SMF registers (see [NF registration](#)) and the AMF discovers this SMF by calling the `Nnrf_Discovery` service that has the S-NSSAI configured in the query parameters. Upon discovering the SMF, the AMF creates a session management resource on the SMF by calling the `Nsmf_PDU_Session_CreateSMContextRequest` service operation. In this it includes the PDU Session Establishment Request NAS message as it is, and any additional metadata it can provide, for example the DNN.

2. Upon receiving the `Nsmf_PDUSession_CreateSMContextRequest` the MAG-c uses the signaled DNN to look up the configured APN/DNN in the following context.

```
configure mobile-gateway pdn apn
```



Note: To enable 5G FWA functionality, the APN must be in the enabled state and configured with an FWA node using the following command.

```
configure mobile-gateway pdn apn fixed-wireless-access
```

If a UDM client service is configured (see [UDM function configuration](#)), a `Nudm_SDM_Get` Request service call is executed to obtain subscription parameters such as session AMBR, default 5QI, default S-NSSAI, and so on.

If the UDM is not configured, these parameters must either be provided by the PDU Session Create message, or locally configured on the MAG-c. Subsequently, the authentication flow of the FWA node (configured using the following command) is used to authenticate the session.

```
configure mobile-gateway pdn apn authentication-flow
```

FWA-specific ADB match criteria or RADIUS include attributes can be used. Match criteria use data from the PDU Session create call and the UDM. For example, if an S-NSSAI is not provided in the create call, but a default S-NSSAI is provided in UDM subscription data, this is matched.

3. After the UDM and authentication flow is completed, the MAG-c answers the call from the AMF `Nsmf_PDUSession_CreateSMContext`. This creates the session context in the AMF only; it does not yet include an answer to the UE-originated PDU Session Establishment Request message.
4. Optionally, after the UDM and authentication flow completes, the MAG-c can obtain policy data from a PCF. See [PCF-based policy management](#) for more information about policy management and the PCF service.
5. When all authentication and policy management is completed, the MAG-c selects an FWA-UP and sets up the corresponding PFCP session that carries the full session data, including all BNG-specific parameters. The RAN (gNodeB) 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 back to the MAG-c. See the *7750 SR and VSR BNG CUPS User Plane Function Guide* for information about the BNG CUPS UPF.
6. The MAG-c acknowledges the UE-originated PDU Session Establishment Request message by generating an N1 PDU Session Establishment Accept message. It also generates an N2 SM Information message which is sent to the RAN (gNodeB), including the F-TEID it received from the FWA-UP. Both messages also include any policy QoS rules applicable to either the UE or RAN. It includes these messages in a `Namf_Communication MessageTransfer` service call toward the AMF. The AMF in turn creates an N2 PDU Session Request call to the RAN to create context on the RAN, including the N2 SM Information and N1 PDU Session Establishment Accept message generated by the SMF. The RAN and UE finalize the radio resources setup. The uplink data can now flow.
7. The RAN acknowledges the session context creation using an N2 PDU Session Response message to the AMF. In this message it also includes an N2 SM Information message with its F-TEID for downlink data, which the AMF forwards to the MAG-c using the `smf_PDU_Session_UpdateSMContext` service. The MAG-c forwards the F-TEID to the FWA-UP using a PFCP Session Modification Request message. This changes the downlink rules from buffering to forwarding with the provisioned F-TEID. The FWA-UP acknowledges the PFCP message to the MAG-c, which in turn acknowledges the `smf_PDU_Session_UpdateSMContext` call. Downlink traffic is now possible.

8. If the UDM is enabled, the MAG-c registers itself for subscription updates for this session context using the Nudm_SDM_Subscription service call. Additionally, it also registers itself as the active SMF for this UE session using the Nudm_UECM_Registration service call.
9. The data layer connectivity is now fully ready, although 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 as needed. See [Address signaling methods and deferred allocation](#) for more information.

3.5.14.5 5G slicing

In 5G networks, a slice identifies an end-to-end logical network, from the UE to the network, on top of the shared physical infrastructure. The user can pick the physical or virtual components on a per-slice basis. This is a very flexible and powerful concept that can be used for many purposes. For example, FWA can be expressed as a single slice with dedicated FWA SMF and FWA UPF components.

MAG-c acting as the SMF is not involved in the slice-selection function itself. However, when selecting a peering NF instance via NRF, the MAG-c can use S-NSSAI as a query parameter to discover only peers associated with the selected slice. The MAG-c also passes the S-NSSAI to any peers, such as the UDM, PCF, or CHF, which require the MAG-c S-NSSAI to identify the services to offer. Finally, the SST and SD part of the S-NSSAI can be used as a match parameter in the authentication database, which allows the S-NSSAI to influence a wide set of parameters on itself.

3.5.14.6 5G-4G interworking

Interworking of 5G-4G allows seamless mobility between a 5G and a 4G radio coverage area.

3.5.14.6.1 Overview of interworking between 5G-4G

Interworking between 5G-4G is done using the MAG-c CPF and the 7750 SR UPF call-flow procedures and the N26 interface, which allows seamless mobility between 5G and 4G systems.

Interworking between 5G-4G is performed with support for call-flow procedures using the following:

- a combined SMF and SGW-C/PGW-C for the control plane function (MAG-c)
- a combined UPF and SGW-U/PGW-U for the user plane function (7750 SR)

The 5G-4G interworking is supported with the N26 interface only. N26 is the interface between the AMF and the MME that enables seamless mobility between 5G and 4G systems.

The combined SMF and SGW-C/PGW-C (MAG-c) also supports serving of 4G-only UEs (UEs without 5GS NAS support), while using 5G SBI (N7 and N40) with the PCF and CHF.

To configure the MAG-c for 5G-4G interworking, the following is required:

- To facilitate discovery of the MAG-c for sessions that require 4G interworking support, the PGW FQDN must be configured for use in the NF profile during NF registration, and for update to the NRF. Use the following command to configure the PGW FQDN.

```
configure mobile-gateway pdn nf-profile-attributes pgw-fqdn
```

- 5G SBI interface support is required for UEs without 5GS NAS support (4G LTE). Use the following command to configure the 5G SBI interface support.

```
configure mobile-gateway pdn sba-client-services converged-interface-support pdu-session-id-base
```

- The serving node IP address that the PGW sends to the various interfaces for combined SGW/PGW sessions must be configured to send the MME IP address or the SGW IP address. Use the following command with the corresponding option to configure the MME or SGW.

```
configure mobile-gateway pdn serving-node-for-combo-sessions
```

- Instead of a trigger from a peering node (for example, the PCF) on a call-flow failure because of a handover or TAU in progress, a retry time and delay are required for a local retry. The call procedure retry is supported for the procedures initiated by the SMF/PGW while the handover or TAU is in progress. The retry delay specifies the duration the SMF/PGW must wait for the next retry attempt after the TAU/HO is completed. Use the following command to configure these settings.

```
configure mobile-gateway pdn call-procedure-retry
```

3.5.14.6.2 Bearer and QoS mapping between 5G-4G

In 5G-4G interworking, the bearer and QoS parameter-mapping mechanisms for a 5G-capable UE depend on the access via 4G/LTE or 4G or 5G handover.

If a 5G-capable UE initially attaches via 4G/LTE access, or moves to 4G/LTE access because of a handover from 5G, the following bearer and QoS parameter mappings apply:

- Mapping of QoS parameters and PDU session takes place during PDN connection establishment via the S11 interface.
- The dedicated bearer is currently not supported. As the UE moves between 4G and 5G access, bearer mapping is handled without support for dedicated bearer. For example, the non-GBR and the GBR QoS flow on 5G is rejected.
- The combined SMF and SGW-C/PGW-C maps the EPS QoS parameters to 5GC QoS parameters; for example, QCI-1=5QI-1, QCI-5=5QI-5.
- The 5G UE allocates the PDU session ID to the PDN connection and sends it in the PCO to the PGW-C or SMF. The UE stores the received 5G QoS rules to use when it moves to 5G access.
- 5G QoS parameters (session AMBR, QoS rules (TFT), QoS flow descriptions, and S-NSSAI) corresponding to the PDU connection are sent to the UE in a PCO, while the UE is anchored via 4G access.

If a 5G-capable UE moves to 5G access because of a handover from 4G, the following apply:

- Mapping of the QoS parameters and the bearer ID takes place during the PDU session and GBR flow establishment via the N11 interface.
- The combined SMF and SGW-C/PGW-C maps the 5GC QoS parameters to the EPS QoS parameters; for example, 5QI-1=QCI-1, 5QI-5=QCI-5.
- The AMF manages the bearer ID allocation and the UE/MM context mapping.

3.5.14.7 AMF set support

The AMF set feature provides high availability for the AMF by allowing multiple AMF instances to operate together as a single redundant group. AMF instances within a set serve the same geographic area and share UE context data, enabling any instance in the set to take over a UE session context if the primary AMF fails. This supports service continuity of the mobile core during hardware or software disruptions, including partial failures within the AMF set.

Each UE initially uses a specific AMF instance, while the MAG-c maintains a logical mapping of the entire AMF set. AMF instances in an AMF set share the same set region ID and AMF set ID values. This grouping enables the MAG-c to move sessions to alternative AMF instances based on the health of individual nodes.

The MAG-c subscribes to the NRF for updates about the AMF set, allowing it to maintain visibility into all available nodes. It also monitors HTTP/2 connectivity to AMF peers to determine when to move sessions to another AMF.

When an AMF node fails, the MAG-c does not immediately move all affected UEs. Instead, it selects an alternative active AMF for a UE during the next service request or signaling interaction for that UE. The MAG-c uses a round-robin algorithm to select among the remaining active AMF instances.

In all cases, the MAG-c updates its signaling path to the new AMF to maintain session continuity during both unplanned failures and planned maintenance windows. This capability allows operators to scale their networks and perform system updates while maintaining active sessions.

For maintenance use cases, session transition within an AMF set uses the following mechanisms:

- offloading (non-directed) – This is typically used during maintenance; the active AMF triggers a 307/308 Redirect message to move sessions to a peer AMF before going out of service.
- rebalancing (directed) – This uses the same 307/308 Redirect mechanism to move sessions to another AMF within the set for load distribution.
- 503 + UNSPECIFIED handling – If a session receives a 503 response with cause UNSPECIFIED, the MAG-c selects another AMF from the set using a round-robin algorithm. Unlike the 307/308 Redirect, no specific target AMF is indicated.

3.5.15 Signaling IPv4 MTU to UE

The MAG-c communicates a configurable maximum transmit unit (MTU) for IPv4 packets to a requesting user equipment (UE) during session establishment. The MTU value is conveyed in the protocol configuration option (PCO) or extended protocol configuration option (ePCO) information element (IE) via the GTPv2 (S11 interface) or HTTP/2 (N11 interface). The MAG-c only sends the MTU value if the UE explicitly requests it. Configuring the MTU to an optimal value avoids fragmentation of the data the UE transmits.

Use the following command to configure the IPv4 MTU value for transmission in the PCO or ePCO IE.

```
configure mobile-gateway profile authentication-database entry fixed-wireless-access ipv4-link-mtu
```

3.5.16 Load control, overload control, and congestion mitigation

In the MAG-c system, the load and overload control mechanisms are distinct but complementary concepts that help manage the flow of traffic for subscriber sessions. Load control uses preventive approaches to help avoid overload conditions. For sessions that are experiencing overload conditions, overload control enables the MAG-c node to gracefully reduce traffic for the existing sessions. In addition, the MAG-c provides support for node-level congestion mitigation.

3.5.16.1 Load control configuration guidelines

Load control is an overload prevention concept. It enables the originator node (for example, SGW-C + PGW-C or SMF) to send its load information to a consumer node (for example, the MME or the AMF). During the session setup, the consumer node uses the load information it receives from multiple originator nodes to select nodes for load balancing the session across all the nodes, according to their effective loads.

Use the following command to enable GTP-C load control.

```
configure mobile-gateway pdn load-control enable-gtp-load-control
```



Note: The **enable-gtp-load-control** command enables load control using default configurations. Use the commands in the **load-control** context to change the default configurations, as required.

To support load control with 4G or 5G NSA, the originating node sends the Load Control Information (LCI) IE in the GTP-C messages to the consumer node. The MAG-c can send GTP-C load control information in every *n*th GTP-C message, or on every *x* percentage load change.

The load metric reported in the LCI represents the current load level of the MAG-c node as a percentage, within the range of 0 to 100, where 0 means no load and 100 means the maximum or 100% load (no further load is desirable).

Use the following command to configure the GTP-C load control settings based on the percentage of load change.

```
configure mobile-gateway pdn load-control load-change-trigger
```

Use the following command to configure the GTP-C load control settings to send the load control information in every *n*th message.

```
configure mobile-gateway pdn load-control nth-message
```

To support indirect load reporting to the consumer nodes, the producer node updates the NRF with its local load information. The load reported by the producer (originating) node to the NRF is made available to the consumer node during node selection. The consumer node can load balance the signaling load by selecting a producer node that has a lower load compared to a similar node with a higher load. This method is used for load control for SBI interfaces only.

Use the following command to enable SBI load reporting.

```
configure mobile-gateway pdn load-control enable-sbi-load-reporting
```

3.5.16.2 Overload control configuration guidelines

Overload control is an overload mitigation concept. It enables an originator node (SGW-C + PGW-C, SMF, MME, or AMF) that is overloaded, or at risk of becoming overloaded, to gracefully reduce the incoming signaling. The originator node is in an overload state when it operates beyond its signaling capacity, resulting in diminished performance and impacting the handling of incoming and outgoing traffic.

Overload control configuration and examples

The MAG-c overload control mechanism supports traffic reduction in overload conditions by instructing the consumer node (MME, AMF) to reduce the volume of traffic it is sending for existing sessions.

Use the following command to enable GTP-C overload control. This automatically enables default threshold and traffic-reduction values for the different overload levels.

```
configure mobile-gateway pdn overload-control enable-gtp-overload-control
```

To support overload control with 4G or 5G NSA, the overloaded originator node sends an Overload Control Information (OCI) message to inform the consumer node about the overload condition. The OCI information includes the traffic-reduction value that the consumer node applies to throttle signaling traffic toward the originator node. Overload levels are identified as minor, major and critical. The thresholds that define when the node reaches a minor, major, or critical overload level, can be configured on the MAG-c.



Note: The MAG-c comes out of an overload condition when the overload is 2% below the configured (or default) minor value.

Use the commands shown in the following example to modify the default thresholds for the different overload levels.

Example: Modify threshold default values for overload levels

```
A:MAG-c# configure mobile-gateway pdn overload-control thresholds minor 85
A:MAG-c# configure mobile-gateway pdn overload-control thresholds major 90
A:MAG-c# configure mobile-gateway pdn overload-control thresholds critical 95
```

Use the following command to modify the default time period during which the reported overload remains in effect.

```
configure mobile-gateway pdn overload-control validity-time
```

Use the commands shown in the following example to modify the default traffic-reduction values that are signaled to the MME node for the overload levels (within the supported ranges).

Example: Modify traffic reduction default values for overload levels

```
A:MAG-c# configure mobile-gateway overload-control overload-reduction minor 40
A:MAG-c# configure mobile-gateway overload-control overload-reduction major 70
A:MAG-c# configure mobile-gateway overload-control overload-reduction critical 85
```

You can also enable and disable overload throttling at the local level. Use the following command to enable local overload throttling. This automatically enables local overload throttling values for the different overload levels.

```
configure mobile-gateway pdn overload-control local-throttling local-overload-throttling
```

Use the commands shown the following example to modify the percent of the traffic to be throttled when the MAG-c node reaches the minor, major, or critical overload level (within the supported ranges).

Example: Modify local throttling default values for overload levels

```
A:MAG-c# configure mobile-gateway pdn overload-control local-throttling minor throttling-percent 0
A:MAG-c# configure mobile-gateway pdn overload-control local-throttling major throttling-percent 50
A:MAG-c# configure mobile-gateway pdn overload-control local-throttling critical throttling-percent 80
```

The following table describes the overload reporting to the MME and the local throttling applied for the associated overload condition levels, based on the settings used in the preceding overload-reduction and local-throttling examples.

Table 6: Example overload reporting and local throttling

| Overload condition | Overload reporting to the MME | Local throttling |
|--------------------|--|---|
| Minor | The SGW-C + PGW-C reports the minor overload condition to the MME, informing it to reduce traffic by the configured minor overload value of 40%. | Zero (0) percent of the traffic is throttled. The SGW-C + PGW-C does not throttle any GTP-C traffic. At this overload level the MME is permitted to reduce traffic with its own priority procedures. |
| Major | The SGW-C + PGW-C reports the major overload condition to the MME, informing it to reduce traffic by the major overload of 70% (configured, as shown in the preceding example) | 50% of the incoming Create Session Requests are throttled. Any Network initiated modifications that could trigger an Update Bearer Request are also throttled at the same percentage. |
| Critical | The SGW-C + PGW-C reports the critical overload condition to the MME, informing it to reduce traffic by the critical overload of 85% (as configured in the preceding example) | 80% of the traffic is throttled, as configured in the preceding example. 80% of the received GTP-C Create Session Requests, Bearer Resource command, or Modify Bearer command call flows are throttled. Any Network initiated modifications that could trigger an Update Bearer Request are also throttled at the same percentage. |

As a consumer, the SGW-C + PGW-C also receives overload information from the MME. When it receives a GTP-C OCI IE from the MME in overload, the SGW-C + PGW-C throttles signaling traffic toward the MME. Based on the traffic reduction received in the OCI IE from the MME, the SGW-C + PGW-C acts on any network-initiated call flows to the MME, based on the priority shown in the following table.

Table 7: Message throttling priority

| Throttling priority | Network-initiated messages |
|----------------------|----------------------------|
| Low (dropped last) | Release PDN connection |
| High (dropped first) | Update PDN connection |

Throttling restrictions for overload conditions

Throttling is not applied to the following overload conditions:

- major overload
 - Wireless Priority Sessions (WPS)
 - handovers, except those based on CSReq which fall into throttling at major overload
- critical overload
 - WPS sessions
 - Delete PDN session
 - handovers, except those based on CSReq which fall into throttling at critical overload

Viewing load and overload control information

Use the following command to view overload throttling information.

```
show mobile-gateway pdn load-overload-control summary
```

3.5.16.3 Congestion mitigation

The MAG-c supports node-level N11 congestion mitigation based on 3GPP TS29.500, 29.502. SBI congestion mitigation is triggered on a consumer node based on the HTTP status codes the node receives from the originator (producer) node.

When acting as an SMF, the MAG-c as the producer node on an N11 interface supports sending the HTTP 503 code during overload conditions. You can configure the overload threshold for congestion mitigation using CLI commands.

During overload, the PDU Session Establishment Create Request is rejected, and the HTTP 503 error code is generated toward the consumer node. The MAG-c responds to the AMF with an HTTP 503 message, including the error attribute ProblemDetails set to "NF_CONGESTION". If the **sbi-retry-after-timer** command is also configured, the value is included in the HTTP 503 response.

When no longer in overload, the MAG-c stops rejecting any new PDU Session Establishment Create Requests.

Use the following command to enable SBI overload control.

```
configure mobile-gateway pdn overload-control enable-sbi-overload-control
```

Use the following commands to modify the default values of the SBI overload threshold and retry timer.

```
configure mobile-gateway pdn overload-control sbi-overload-threshold  
configure mobile-gateway pdn overload-control sbi-retry-after-timer
```

3.5.17 Wireless Priority Service

Users can configure Wireless Priority Service (WPS) on the MAG-c to enable subscribers who provide national security and emergency services to make priority calls (priority sessions) using public networks during network congestion conditions. The MAG-c can be configured to provide network support for end-to-end session-priority treatment in the core and radio networks.

3.5.17.1 WPS overview

Government-authorized personnel, emergency management officials, and other authorized users who provide disaster response and management rely on the ability to communicate during congestion conditions. These subscribers are expected to receive priority treatment, in support of mission critical multimedia communications.

There are two methods to qualify a high-priority session:

- **subscription-based WPS**

For subscription-based WPS, the Message-Priority value received in the GTP message header on the S11 interface during the initial session setup for 4G/5G NSA would indicate the session is high priority. In this case, the signaling messages toward other peers (UDM, CHF, and PCF) also carry a high-priority indication in the SBI message headers, so that the signaling messages on the receiving node are processed at high priority as well. High-priority indication over the SBI interface is achieved using the Message-Priority value in the HTTP/2 header.

The subscription-based priority service can be enabled on the MAG-c by mapping the Message-Priority value received in the GTP signaling message on S11 for 4G/5G NSA to indicate it is a high-priority session.

- **invocation-based WPS**

For invocation-based WPS, the authorized ARP value (received from the PCF) maps the session as a high-priority session. The invocation-based priority service is enabled by mapping the authorized ARP value (for example, received from PCF) for 4G/5G NSA or 5G SA to indicate it is a high-priority session.

In both cases, the authorized ARP value determines the final priority level of the session. The priority level based on authorized ARP value overwrites the priority level assigned based on the Message-Priority value received in the initial GTP messages on the S11 interface.

3.5.17.2 WPS configuration guidelines

The MAG-c uses configurable priority-mapping profiles to map high-priority sessions based on the Message-Priority value received as follows:

- in GTP messages for 4G and 5G NSA sessions
- in SBI messages for 5G SA sessions
- on the authorized ARP for 4G and 5G NSA and 5G SA sessions

The user defines the priority level for the priority-mapping profile by assigning a priority value. The following example shows priority-level configurations for priority-mapping profiles. Based on this example, the priority sessions are processed as follows:

- An ARP value of 1 is processed as a high priority session and all other values of ARP are processed at regular priority.
- All Message-Priority values received in GTP and SBI messages are processed as high-priority sessions, indicated by the wildcard (*).

Example: Priority mapping profile configuration

```
A:MAG-c>config>mobile>profile# info
-----
    priority-mapping "profile1"
      arp 1
        priority 2
      exit
      gtp-rx-message-priority *
        priority 2
      exit
      sbi-rx-message-priority *
        priority 2
      exit
    exit
-----
```



Note: See the *MAG-c Release Notes* for information about authorized ARP, GTP Message-Priority, and SBI Message-Priority values supported on MAG-c.

High-priority sessions can be treated with additional priority among themselves. You must define a priority profile for the differentiated treatments among the high-priority sessions, where each high-priority session is identified by one of the priority values (1 to 4).



Note: See the *MAG-c Release Notes* for information about the number of priority levels that MAG-c supports.

Use the following commands to configure a priority profile and assign a priority identifier.

```
configure mobile-gateway priority-profile
configure mobile-gateway priority-profile priority
```

The following example shows a priority profile configuration for exceptional high-priority sessions, such as emergency first responders. An egress policy is assigned for high-priority sessions that map to a priority level based on the value assigned to the **priority** command. The following information applies, based on the configuration shown in the example:

- The priority identifier value (2 in the example) is assigned under the priority-mapping profile, based on the priority criteria for each option.
- A Message-Priority value 0 is included in GTP messages sent over the s11 interface to the MME, and the GTP messages are marked with DSCP value 56.

- A Message-Priority value 0 is included in the SBI messages sent to peers, for example, the UDM, CHF, or PCF.
- The priority level is `critical` while processing PFCP messages marked with DSCP value of 56.



WARNING: Use the `critical` PFCP priority level sparingly because it bypasses any overload protection mechanisms. Contact your Nokia customer service representative for information about the recommended configuration when using the `critical` option.

Example: Priority profile configuration for exceptional high-priority WPS sessions

```
A:MAG-c>config>mobile>profile# info
-----
      priority-profile "profile1"
        priority 2
          egress
            gtp-tx-message-priority 0
            sbi-tx-message-priority 0
          gtp-dscp 56
          pfc-priority-level critical
          pfc-dscp 56
          exit
        exit
-----
```

3.5.18 Radio UP security profiles

The MAG-c supports signaling of UP security profiles to the RAN based on the `UpSecurity` IE in the UDM subscription data.

The MAG-c supports signaling of UP security profiles to the RAN based on UDM subscription data, as specified in 3GPP TS 33.501 clause 6.6. When the UDM subscription data includes the `UpSecurity` IE, the MAG-c enables the UP security profile feature. The UP security profile consists of the following parameters:

- **UP confidentiality**
Possible values are:
 - required
 - preferred
 - not needed

For more information, see 3GPP TS 29.571 clause 5.4.3.5.

- **UP integrity**
Possible values are:
 - required
 - preferred
 - not needed

For more information, see 3GPP TS 29.571 clause 5.4.3.4.

When the UP integrity parameter is set to **required** or **preferred**, the MAG-c learns the UE maximum integrity protection data rate from N1 signaling.

For 5G communication, the MAG-c signals the UP security profile and the learned UE maximum integrity protection data rate, if applicable, over N2 to the gNB.

For 4G communication, the behavior depends on whether the user plane integrity protection (UPIP) feature is supported by the client and the RAN. Support for the UPIP feature is signaled during session setup for both 4G and 5G to support 4G-5G interworking procedures. Support for the IPIP feature only affects the UP integrity parameter; the SMF does not signal the UP confidentiality parameter for 4G communication.

The MAG-c behavior for 4G communication is as follows:

- **UPIP feature is supported**
The MAG-c signals the UP integrity value along with the UE maximum integrity protection data rate, if applicable, to the MME in a GTP Create Session Response. The MME then delivers these values to the eNB. 4G-5G interworking is allowed.
- **UPIP feature is not supported and the UP integrity is preferred or not needed**
The MAG-c sets up 4G communication without signaling the UP integrity value or the UE maximum integrity protection data rate. 4G-5G interworking is allowed. During a 4G to 5G handover, the MAG-c sends the UP security profile and the UE maximum integrity protection data rate, if applicable, to the 5G gNB.
- **UPIP feature is not supported and UP integrity is required**
The MAG-c rejects a 4G session setup. 4G-5G interworking is not allowed.



Note: If the MAG-c does not interact with the UDM, or the UDM subscription data does not include the UpSecurity IE, the MAG-c sets up sessions without any UP security profile or UE maximum integrity protection data rate signaling. 4G session creation and 4G-5G interworking are not restricted.

3.5.19 Aggregate statistics collection

The MAG-c supports the automatic collection of the following aggregate statistics for specified scopes:

- number of sessions currently connected within the specified scope
- number of forwarded and dropped packets and octets in the uplink and downlink direction

The MAG-c calculates the latter statistics based on statistics collected for individual sessions. Each time a PFCP Usage Report is received, the MAG-c adds the reported statistics to the scope in which the session is currently active.



Note: Because of this mechanism, these counters often lag actual consumed data traffic that has not yet been reported to the MAG-c.

The MAG-c supports the following scopes:

- per QCI or 5QI for each APN in a UP, enabled using the following command

```
configure mobile-gateway pdn aggregate-statistics-collection per-up per-apn per-qos-identifier
```

- per radio access type (RAT) for each APN in a UP, enabled using the following command

```
configure mobile-gateway pdn aggregate-statistics-collection per-up per-apn per-radio-access-type
```

- per 5QI or QCI for each slice in a UP, enabled using the following command

```
configure mobile-gateway pdn aggregate-statistics-collection per-up per-slice per-qos-identifier
```

- per APN for each UP, enabled automatically when the scope per RAT or per 5QI or QCI for each APN in a UP is enabled
- per slice for each UP, enabled automatically when the scope per 5QI or QCI for each slice in a UP is enabled
- per UP, enabled automatically when one of the above scopes is enabled



Note: The following commands do not enable a scope for statistics collection, they merely enable a context containing configuration commands.

```
configure mobile-gateway pdn aggregate-statistics-collection per-up
configure mobile-gateway pdn aggregate-statistics-collection per-up per-apn
configure mobile-gateway pdn aggregate-statistics-collection per-up per-slice
```

The following commands do enable a scope for statistics collection, and as a consequence, also enable the parent scopes.

```
configure mobile-gateway pdn aggregate-statistics-collection per-up per-apn per-qos-identifier
configure mobile-gateway pdn aggregate-statistics-collection per-up per-apn per-radio-access-type
configure mobile-gateway pdn aggregate-statistics-collection per-up per-slice per-qos-identifier
```

For example, the **configure mobile-gateway pdn aggregate-statistics-collection per-up per-apn per-qos-identifier** command enables three scopes for statistics collection:

- per UP
- per APN for each UP
- per QCI or 5QI for each APN in a UP

Use the following command to show the statistics:

- **classic CLI**

```
show mobile-gateway pdn statistics aggregates
```

- **MD-CLI**

```
state mobile-gateway pdn statistics aggregates
```

See [YANG state](#) for more information about accessing state information via the MD-CLI interface.

3.5.20 DNS-based content filtering

The MAG-c uses the ESA-AA to support enrichment of the DNS requests with a per-subscriber policy ID. This enables a DNS-based per-subscriber management scheme that differentiates DNS answers based

on the policy ID of the subscriber. Users can leverage this feature to offer a DNS-based approach for web-filtering solutions.

The MAG-c receives the per-subscriber policy ID from the PCF in the Nokia-specific PFCP Content-Filtering-Policy-Id IE, which has a type value of 32840 and a fixed length of 4 bytes.

To include this IE in the PFCP message, use the following command with **content-filtering-policy-id** as the field name.

```
configure mobile-gateway profile pfcpc pfcpc-translate-profile field
```

3.5.21 Enabling multiple sessions per APN or DNN for each subscriber

The MAG-c supports setup of multiple PDN sessions per APN and multiple PDU sessions per DNN, for each subscriber in the 4G or 5G network.

Multiple PDN and PDU sessions can be established from each UE, independently from the PDN type (IPv4, IPv6, or IPv4v6). The following methods are used to identify unique PDN and PDU sessions:

- For 4G and 5G NSA, the PDN sessions are identified by a unique pair of EPS bearer ID (EBI) and APN identifier.
- For 5G SA, the PDU sessions are identified by a unique pair of PDU session ID (SEID) and DNN identifier.

Use the following command to enable MAG-c support for multiple PDN sessions per APN per UE, and multiple PDU sessions per DNN per UE.

```
configure mobile-gateway pdn multiple-sessions-per-apn-per-ue
```

Use the following commands to display information about the subscriber and EBI for 4G sessions.

```
show mobile-gateway pdn pdn-context imsi [eps-bearer-id]
show mobile-gateway bng session imsi [eps-bearer-id]
```

Use the following commands to display information about the PDU SEID and SUPI for 5G sessions.

```
show mobile-gateway bng session imsi [n1-pdu-session-id]
show mobile-gateway pdn pdn-context supi [detail]
show mobile-gateway pdn pdu-session supi [n1-pdu-session-id]
show mobile-gateway pdn pdu-session supi detail [dnn] [n1-pdu-session-id]
```

3.6 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 8: Address assignment protocols per session type

| Address assignment protocol | IPoE session | PPPoE session |
|---|--------------|---------------|
| DHCP | ✓ | |
| IPCP | | ✓ |
| DHCPv6 NA | ✓ | ✓ |
| DHCPv6 PD | ✓ | ✓ |
| SLAAC (see ICMPv6 Router Advertisements and SLAAC) | ✓ | ✓ |

3.6.1 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 following context.

```
configure mobile-gateway pdn apn
```

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 following context

```
configure mobile-gateway profile bng
```

- 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.
 - Specific 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](#)

3.6.2 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.

3.6.3 ICMPv6 Router Advertisements and SLAAC

The MAG-c periodically generates ICMPv6 Router Advertisements (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 following context.

```
configure mobile-gateway profile authentication-database entry interface
```

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 following context to configure the RA profile locally.

```
configure mobile-gateway profile bng
```

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 a DHCPv6 IA-NA or an 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.



Note: Nokia recommends that preferred lifetime is at least double or more than the configured maximum advertisement interval. This avoids the expiration of the preferred lifetime on the client side because of the loss of a single ICMPv6 RA message.

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).

3.6.4 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.

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

- an allocated IA-NA address, an 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 following context.

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

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 following context.

```
configure mobile-gateway pdn apn
```

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 IPv6CP 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 VSA
- bulk options derived from a locally configured DHCP profile using the **dhcpv6-profile** command in the following context

```
configure mobile-gateway profile bng
```

- 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 and SLAAC](#)

3.7 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
<logout and login again>
//
state mobile-gateway bng queries
```

The first command configures the MD-CLI engine, after which the user needs to logout from the CLI and login again, the second command switches between the classic CLI and the 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

```
A:MAG-c>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

```
A:MAG-c>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
  }
}
}

```

4 Service Based Interfaces


The 5G core network is a Service Based Architecture (SBA), which consists of multiple Network Functions (NFs). Each interface in the Service Based Interfaces (SBI) consists of one or more services that are used for the communication between the functions.

4.1 Introduction to Service Based Architecture

The 5G core network is an SBA, which consists of multiple Network Functions (NFs) communicating using a common service-based protocol. This protocol is based on [OpenAPI v3](#) using JSON signaled over HTTP/2. Multiple instances of Network Functions can exist (for example, multiple SMFs) and each instance is identified by a Network Function Instance ID (NFInstanceID), which is modeled as a Unique Universal Identifier (UUID).

Communication between functions is based on services offered by each function. Each service consists of a service producer (server) and a service consumer (client). Each service is typically oriented at executing a very specific function. An interface can consist of multiple services, as shown in [Figure 16: Relationship between service and interface](#) for the N11 interface. The N11 interface consists of at least the Nsmf_PDUSession service where the SMF acts as server, and the Namf_Communication service where the AMF acts as server. This concept is called SBI, where each interface (such as N11) consists of one or more services. Services can be reused over interfaces; for example, the Nudm_SDM service is used both on the N8 interface by AMF and N10 interface by SMF.

Figure 16: Relationship between service and interface




sw1401

In general, any consumer can consume a service, subject to access control. This allows for easy extensibility and re-use. Examples of this are the Nnrf_NFDiscovery and Nnrf_NFManagement services, both of which are served by the NRF. These services are used by other NFs to discover other NFs and register themselves for discovery respectively. For this reason, a 5G core architecture is often presented in an alternative form where the focus is not on interfaces such as shown in [Figure 14: Basic 5G FWA network](#) but on the provided services. An example of such a representation is shown [Figure 17: Basic 5G](#)

[network with service representation](#). This representation also makes it clear that the SBA is focused on the control plane functions.

Figure 17: Basic 5G network with service representation



sw1402

4.2 NF peer discovery

When acting as a client, the MAG-c performs the following steps to establish a peer to host the NF service:

1. Discover the list of NF instances identified by UUIDs and select one NF instance:

- The NF instance with highest priority is usually selected.




Note: A lower priority value means a higher priority.

- When multiple NF instances have the same priority, the UPF does hash-based selection per session.
- In some cases, a specific UUID may already be signaled and picked from the list of discovered instances, ignoring the priority. For example, the AMF signals in the Nsmf_PDUSession operations the AMF UUID to use for the related Namf_Communication service.

2. Determine a list of IP addresses associated with the selected NF instance and select one IP address to use.

For the discovery of NF instances, the MAG-c supports both NRF-based discovery using the Nnrf_NFDiscovery service (see [NRF-based discovery](#) for more information) and a locally configured NF ID list (see [Enabling discovery based on a local NF ID list](#) for more information). When configured, NRF-based discovery has precedence and the NF ID list is only used if the NRF lookup fails or does not return any NF instances. The following figure shows the NF instance discovery.

Figure 18: NF-instance discovery flowchart



sw1403

4.2.1 NRF-based discovery

To enable NRF-based discovery, perform the procedure [Enabling NRF-based discovery](#).

When NRF discovery is provisioned, it is automatically enabled for AMF, CHF, PCF, and UDM. Use the following command to suppress discovery for any of these functions.

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-disc suppress-discovery
```

The following query options are included when discovering PCF, CHF, or UDM:

- **S-NSSAI**

S-NSSAI is included by default for PCF and UDM discovery, however, you can use one of the following commands to disable it.

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-disc exclude-snsais
configure mobile-gateway pdn apn pcf-selection exclude-snsais
```

- **SUPI**

SUPI is included by default, however, you can disable it for each function using the following commands:

- for the UDM

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-disc udm-selection-
options exclude-supi
```

- for the CHF

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-disc chf-selection-
options exclude-supi
```

- for the PCF

```
configure mobile-gateway pdn apn pcf-selection exclude-supi
```

- **preferred-locality**

This option is not included by default, however, you can enable it and set a global locality for UDM, PCF, and CHF discovery using the following commands:

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-disc use-locality
configure mobile-gateway pdn nf-profile-attributes locality
```

To enable this option and set a locality separately for each function, use the following commands:

- for the UDM

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-disc udm-selection-
options preferred-locality
```

- for the CHF

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-disc chf-selection-
options preferred-locality
```

- for the PCF

```
configure mobile-gateway pdn apn pcf-selection preferred-locality
```

When both global and per-function locality are present, the per-function locality takes precedence over the global locality.

- **GPSI**

GPSI is included by default for PCF and CHF discovery. You can disable it using the following commands:

- for the CHF

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-disc chf-selection-
options exclude-gpsi
```

- for the PCF

```
configure mobile-gateway pdn apn pcf-selection exclude-gpsi
```

- **DNN**

DNN is included by default for the PCF.

- **Serving scope**

This option is included by default for the PCF.

The NRF keeps a cache of discovered NF instances per query. This cache expires based on the validityPeriod IE returned from the NRF. Use the following command to override this value.

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-disc nrf-cache-interval
```

When a valid cache entry is discovered, it is used instead of contacting the NRF. When an NRF cannot be reached or returns an error, it is temporarily blocked for discovery.

4.2.2 Enabling NRF-based discovery

Enable NRF-based discovery using the Nnrf_NFDiscovery service for the discovery of NF instances.

About this task

Configure and enable a client service instance to enable NRF-based discovery.

Procedure

Step 1. Configure a client service instance.

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-disc
```

Step 2. Within the NNRF client service, configure the following:

a. Specify the interface that MAG-c uses to communicate with the NRF.

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-disc interface
```

b. Specify that NF instances can be used for the discovery service; see [Enabling discovery based on a local NF ID list](#) for more information.

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-disc nf-id-list
```

Step 3. Provision an optional HTTP/2 profile to change generic HTTP/2 and SBI functionality, such as HTTP/2 level retries, or enable Python. See [HTTP/2, OpenAPI, and Python](#) for more information.

Step 4. Enable the client service instance.

```
configure mobile-gateway pdn sba-service-realm client-service
```

4.2.3 Enabling discovery based on a local NF ID list

About this task

Use this task to configure an NF ID list for local discovery.

Procedure

- Step 1.** Use the commands in the following context to configure an NF ID list with multiple NF profile ID entries, each identifying an NF instance and its profile.

```
configure mobile-gateway profile list nf-id-list nf-prof-id
```

- Step 2.** Configure the following per NF instance:

- Configure a mandatory UUID to globally identify the NF instance.

```
configure mobile-gateway profile list nf-id-list nf-prof-id uuid
```

- Configure one or both of the following commands:
 - Configure an FQDN to retrieve a list of IP addresses using a DNS client (see [Enabling DNS-based IP lookup](#) for more information).

```
configure mobile-gateway profile list nf-id-list nf-prof-id fqdn
```

- Configure a prioritized address to directly determine a list of IP addresses.

```
configure mobile-gateway profile list nf-id-list nf-prof-id prioritized-address-list
```

- Configure an NF priority to determine the priority of this NF Instance compared to other NF instances. A lower value has a higher priority.

```
configure mobile-gateway profile list nf-id-list nf-prof-id nf-priority
```

- Configure a custom API prefix to insert in the URI for the service calls, as specified in TS 29.501.

```
configure mobile-gateway profile list nf-id-list nf-prof-id api-prefix
```

- Activate this NF instance for selection. For this, the UUID and either the FQDN or prioritized address list must be configured.

```
configure mobile-gateway profile list nf-id-list nf-prof-id enable
```

- Step 3.** To use the NF ID list for NF discovery, configure the **nf-id-list** command for any of the following services:



Note: If enabled for the service, the NRF-based discovery takes precedence over the NF ID list configuration. However, the NF ID list may still be used as a fallback, or for cases where the NRF returns a UUID but not an FQDN or IP address. See [NF peer discovery](#) and [NRF-based discovery](#) for more information.

- Nnrf_NFDiscovery service

```
configure mobile-gateway pdn sba-client-services nrf-client nnrf-disc nf-id-list
```



Note: This service does not support NRF-based discovery.

- Nnrf_NFManagement service

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-nfm nf-id-list
```



Note: This service does not support NRF-based discovery.

- Namf_Communication service

```
configure mobile-gateway pdn sba-client-services amf-client namf-comm nf-id-list
```

- Nudm_SDM service

```
configure mobile-gateway pdn sba-client-services udm-client nudm-sdm nf-id-list
```

- Nudm_UECM service

```
configure mobile-gateway pdn sba-client-services udm-client nudm-uecm nf-id-list
```

- Npcf_SMPolicyControl service

```
configure mobile-gateway pdn apn pcf-selection nf-id-list
configure mobile-gateway pdn sba-client-services pcf-client npcf-smpolicycontrol
nf-id-list
```



Note: If both are present, the APN-level configuration has priority over the SBA client-level configuration.

- Nchf_ConvergedCharging service

```
configure mobile-gateway pdn apn charging nchf chf-selection nf-id-list
```

-  **Note:** See [CHF peer discovery](#) for more information about CHF peer selection priority.

4.2.4 Enabling DNS-based IP lookup

About this task

Configure and apply a DNS profile to enable DNS-based IP lookup.

Procedure

Step 1. Configure a DNS profile.

```
configure mobile-gateway profile dns-profile
```

The following parameters can be configured for the DNS profile:

- **IP DSCP** and **IP TTL**

These parameters determine the DSCP and TTL values in the IP headers of outgoing DNS request packets.

- **Message retransmit**

This parameter determines how often and how quickly an unanswered DNS request is retransmitted.

- **Primary, secondary, and tertiary DNS**

These parameters determine the DNS servers to attempt. At least a primary DNS must be configured.

Step 2. Apply the configured DNS profile.

```
configure mobile-gateway pdn dns-client dns-profile
```

Step 3. Configure an interface to send DNS requests.

```
configure mobile-gateway pdn dns-client dns-interface
```

4.3 URI construction

When acting as a client, the MAG-c constructs SBI URIs as specified in clause 4.4 in 3GPP TS 29.501. The following table lists the URI scheme.

Table 9: URI scheme

| apiRoot | | | apiName | apiVersion | apiSpecific Resource |
|---------|---|----------------------------|--------------|-------------|----------------------|
| Scheme | Authority | Deployment-specific prefix | | | |
| http:// | IP endpoint or FQDN, and an optional port | API prefix | Service name | API version | |

The IP Endpoint, FQDN, API prefix, Service name, and API version are either found in the NF profiles or they are hard coded. The NF profiles are discovered either from the NRF or from the local NF ID list; see [NRF-based discovery](#) and [Enabling discovery based on a local NF ID list](#) for more information.

4.4 NF registration

The MAG-c can register itself as an SMF with the NRF using the Nnrf_NFManagement service.

To configure an Nnrf_NFManagement client service instance, do the following:

1. Use the following command to configure the client service.

```
configure mobile-gateway pdn sba-client-services nrf-client nnrf-nfm
```

2. Within the client service, configure the following:

- Use the following command to configure the interface the MAG-c uses to communicate with the NRF.

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-nfm interface
```

- Use the following command to configure the NF instances to use for the registration service; see [Enabling discovery based on a local NF ID list](#) for more information.

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-nfm nf-id-list
```

- Use the following command to configure an optional HTTP/2 profile to change generic HTTP/2 and SBI functionality such as HTTP/2 level retries, or to enable Python. See [HTTP/2, OpenAPI, and Python](#) for more information.

```
configure mobile-gateway pdn sba-client-services nrf-client nrf-nfm http2-profile
```

3. Use the following command option to enable the client service.

```
configure mobile-gateway pdn sba-service-realm client-service service-instance
```

The MAG-c automatically registers itself as an SMF, if the Nsmf_PDUSession service is enabled using the following command.

```
configure mobile-gateway pdn no shutdown
```

4. For registration to succeed, use the following command to provision at least one slice.

```
configure mobile-gateway pdn slices slice-list
```

5. To verify the registration, use the following command to check if the NF Status is registered.

```
tools dump mobile-gateway nf-profile
```

The MAG-c updates its registration when relevant configuration changes are applied; for example, adding or removing an APN, changing a slice list, or changing FQDNs.

Use the commands in the following context to configure many of the NF profile attributes that are sent to the NRF when the MAG-c is registered as the SMF.

```
configure mobile-gateway pdn nf-profile-attributes
```

If an attribute is dependent on a configuration that is not present, the attribute is not included.

The following table describes the NF profile attributes that are sent to the NRF during the MAG-c registration as the SMF, including the commands used to configure the attributes. When no command context is specified, the **nf-profile-attributes** context is used.

The MAG-c can include the smfInfo or the smfInfoList attribute when registering. The smfInfoList is used when multiple sets of smfInfo attributes are discovered and they cannot be merged. For example, if a **tailist** is used with multiple priorities configured, a separate smfInfo object is created for each priority, and each smfInfo object only includes the TAIs or TAI ranges configured with that priority.

Table 10: Attributes sent during NRF registration

| NFProfile attribute | Source |
|--|---|
| nfInstanceId | Based on the configured nf-instance-id command |
| nfType | Always SMF |
| nfStatus | NF status is as follows: <ul style="list-style-type: none"> UNDISCOVERABLE, if the following command is configured <pre>configure mobile-gateway pdn suspend</pre> REGISTERED, if not |
| plmnList | Based on the configuration of the following command <pre>configure mobile-gateway pdn home-plmn-list</pre> |
| sNssais | Based on the superset of all configured slice-lists, as determined by the configuration of the following command <pre>configure mobile-gateway pdn slices slice-list configure mobile-gateway pdn apn slices slice-list</pre> |
| nsiList | Based on the configuration of the following command <pre>configure mobile-gateway pdn slices slice-instance-list</pre> |
| capacity nfServices/capacity | Based on the configured capacity command |
| priority nfServices/priority | Based on the configured priority command |
| smfInfo/priority smfInfoList/smfInfo/priority | Based on the priority configured in the linked tailist . If no such list is linked or if the list does not specify a priority, this IE is not included. If multiple priorities are configured in the list, an smfInfo object is created for each different priority. |
| recoveryTime | Included and automatically generated after a service restart |
| locality | Based on the configured locality command |

| NFProfile attribute | Source |
|--|--|
| fqdn | Based on the configured fqdn command |
| nfServices | List of services provided by the SMF, each entry of which is created based on all the SBA server service instances configured in the following context <pre>configure mobile-gateway pdn sba-server-services</pre> |
| ipv4Addresses ipv6Addresses nfServices/ipEndPoints | IP addresses configured under the interface referenced in the following command <pre>configure mobile-gateway pdn sba-server-services nsmf-pdusession interface</pre> |
| nfServices/serviceInstanceId | Name of the SBA service instance configured using the following command <pre>configure mobile-gateway pdn sba-server-services nsmf-pdusession</pre> |
| nfServices/fqdn | FQDN configured using the following command <pre>configure mobile-gateway pdn sba-server-services nsmf-pdusession</pre> If not provisioned, the FQDN configuration is used in the following context <pre>configure mobile-gateway pdn sba-server-services</pre> |
| nfServices/serviceName | Always nsmf-pdusession |
| nfServices/version/apiVersionInUri | Always v1 |
| nfServices/version/apiFullVersion | Always 1.0.2 |
| nfServices/scheme | Always http |
| nfServices/nfServiceStatus | NF service status is as follows: <ul style="list-style-type: none"> UNDISCOVERABLE, if the following command is configured <pre>configure mobile-gateway pdn suspend</pre> REGISTERED, if not (same as nfStatus) |

| NFProfile attribute | Source |
|--|--|
| nfServices/allowedPlmns | <p>Based on the configuration using the following command</p> <pre>configure mobile-gateway pdn sba-server-services nsmf-pdusession allowed-plmns</pre> |
| nfServices/allowedNssais | <p>Based on the configuration using the following command</p> <pre>configure mobile-gateway pdn sba-server-services nsmf-pdusession allowed-slices</pre> |
| smfInfo/sNssaiSmfInfoList smfInfoList/smfInfo/sNssaiSmfInfoList | <p>List of SMF-specific information for all slices available on this SMF, based on the configuration of the following command</p> <pre>configure mobile-gateway pdn slices slice-list</pre> <p>Or, based on the following command for all APNs that are in an enabled state, where the signaled list is the superset of all configured slice lists</p> <pre>configure mobile-gateway pdn apn slices slice-list</pre> |
| smfInfo/sNssaiSmfInfoList/sNssai smfInfoList/smfInfo/sNssaiSmfInfoList/sNssai | S-NSSAI identifying this slice in the list |
| smfInfo/sNssaiSmfInfoList/dnnSmfInfoList smfInfoList/smfInfo/sNssaiSmfInfoList/dnnSmfInfoList | <p>List of DNNs applicable to this slice, depending on the presence of a DNN-specific slice list configured in the following context</p> <pre>configure mobile-gateway pdn apn slice-list</pre> <p>Includes the DNN if the slice is part of the existing DNN-specific slice list</p> <p>Includes the DNN if no DNN-specific slice exists, and the slice is part of the slice list at the PDN level configured in the following context</p> <pre>configure mobile-gateway pdn slice-list configure mobile-gateway pdn apn slices slice-list</pre> |
| smfInfo/sNssaiSmfInfoList/dnnSmfInfoList/dnn smfInfoList/smfInfo/sNssaiSmfInfoList/dnnSmfInfoList/dnn | <p>The DNN configured using the following command</p> <pre>configure mobile-gateway pdn apn</pre> |


| NFProfile attribute | Source |
|--|--|
| | APN name signaled as-is, with APN and DNN synonymous in this context |
| smfInfo/taiList smfInfo/taiRangeList smfInfoList/smfInfo/taiList smfInfoList/smfInfo/taiRangeList | Based on the configuration of the tai-list command |
| smfInfo/pgwFqdn smfInfoList/smfInfo/pgwFqdn | Based on the configuration of the pgw-fqdn command |

When an NRF cannot be reached or returns an error, it is temporarily blocked for registration as described in [Failure handling](#). The registration proceeds with a different NRF if a backup NF instance or NF peer IP is configured.

4.5 HTTP/2, OpenAPI, and Python

SBI communication uniformly uses an HTTP/2 stack as shown in [Figure 19: HTTP/2 stack](#). Messages are sent over this HTTP/2 stack [OpenAPI v3](#). Most of these messages use JSON encoding, but a few exceptions exist. For example, N1 NAS and N2 AP messages are signaled directly as binary format using the application/vnd.3gpp.5gnas and application/vnd.3gpp.ngap content types respectively.

Figure 19: HTTP/2 stack



sw1411

Use the commands in the following context to change the common HTTP/2 and SBI behavior.

```
configure mobile-gateway profile http2
```



Note: The following profile can be applied in multiple specific service instances. The commands in this topic are in the context of this HTTP/2 profile, unless stated otherwise.

Use the following command to configure a Python policy to apply Python scripts to any JSON-encoded messages. See [Python support](#) for more information.

```
configure mobile-gateway profile http2 python-policy
```

The MAG-c maintains multiple HTTP/2 connections to a peer, but does not use these equivalently. This is to ensure that if one connection fails, for example, because of stream ID exhaustion, a backup connection is available. When not enough connections are available, the MAG-c tries to periodically set up a new connection as determined by the following configuration, until enough connections are available.

```
configure mobile-gateway profile http2 connection-timer
```

While a connection is established, the MAG-c periodically sends an HTTP/2 PING message as determined by the configuration of the **ping-period** command.

HTTP/2 messages are sent out with the configured **IP dscp** command value in the IP header, which defaults to 56.

When acting as an SBI client, the MAG-c creates a new stream for each SBI message sent, subject to the maximum number of outstanding streams. When HTTP/2 messages are unacknowledged, the MAG-c retransmits these messages periodically based on the following command.

```
configure mobile-gateway profile http2 fail-retry-interval
```

The default interval is applied (default 5s) until a maximum number of attempts (default 0) has been reached based on the configuration of the following command.

```
configure mobile-gateway profile http2 fail-num-retries
```

Additionally, an outstanding message is also subject to the configuration of the following command (default 20s).

```
configure mobile-gateway profile http2 request-timeout
```

After the request timeout, the message is canceled. The reason for the request timeout, in addition to the **fail-retry-interval** and the **fail-num-retries** command configurations, is that the latter only takes effect from the moment the message is sent on a stream. The configured **request-timeout** value is started from the moment the client service wants to send the message, including any time the message is queued if the maximum number of outstanding streams is reached.



Note: The client services can override a request timeout. For example, the transaction-timer configuration in the following context acts as an override to this configuration.

```
configure mobile-gateway profile namf-profile transaction-timer
```

4.6 Failure handling

The MAG-c supports flexible and consistent failure handling over multiple SBI applications such as UDM, AMF, PCF, and CHF. Use the following commands to configure failure handling for these applications:

- For timeout of SBI operations, use the **fh-session-no-response** command in the applicable context.

- To modify timeout actions for SBI operations, use the **retry-count** and **transaction-timer** commands in the applicable context.
- For specific HTTP and application error codes, use the **http-status-code** command in the applicable context.

Use these commands to configure the following services in the contexts indicated:

- **Namf_Communication service**

```
configure mobile-gateway profile amf-profile
```

- **get operation of the Nudm_SubscriberDataManagement service**

```
configure mobile-gateway profile udm-sdm-profile get-sm-subscription-data
```

- **subscribe operation of the Nudm_SubscriberDataManagement service**

```
configure mobile-gateway profile udm-sdm-profile create-subscription
```

- **registration operation of the Nudm_UEContextManagement service**

```
configure mobile-gateway profile udm-uecm-profile register-smf
```

- **Npcf_SMPolicyControl service**

```
configure mobile-gateway profile pcf-profile
```

Specify per-operation values in the **sbi-srvc-operation** context of this profile.

Applications support two or more of the following actions when there is a failure:

- **terminate**
Terminate the session.
- **attempt-continue**
Continue the triggering procedure and terminate the session if the failure cannot be handled. For example, a failure on Namf_Communication terminates the session when an update of F-TEID programming upon mobility fails. However, the session continues a failing QoS update by rejecting the triggering procedure.
- **ap-continue**
Take fallback actions to continue the session, known as assume positive (AP) mode. In rare cases the session may still terminate, usually when the fallback action is not properly configured. The fallback option depends on the service or the operation:
 - **Npcf_SMPolicyControl**
The SMPolicy association is removed and all communication with the PCF stops for the remainder of the session lifetime. The session can keep working with the latest PCF rules supplied, or fall back to local rules only. See [PCF failure handling](#) for more information.
 - **Get operation of the Nudm_SubscriberDataManagement**
The session setup continues with local configuration. A default QoS (5QI/ARP/session-AMBR) must be provisioned separately from the UDM. See [Configuring a 5G QoS profile for the UDM failure handling](#) for more information.
 - **Subscribe operation of Nudm_SubscriberDataManagement**
The session continues without UDM subscription.

- **Registration operation of the Nudm_UEContextManagement service**
The session continues without UECM registration.
- **Nchf_ConvergedCharging**
Charging is paused and removed from the FWA-UP. See [CHF failure handling](#) for more information.

The following alternative actions are also available. These behave the same as the main variants. However, after the initial failure, the operation is first retried for a maximum of two alternative NF peers, if any are available:



Note: For Npcf_SMPolicyControl, the retry options only apply to the create operation.

- **retry-and-terminate**
- **retry-and-attempt-continue**
- **retry-and-ap-continue**

See [NF peer discovery](#) for more information about how to discover multiple NF peers.

4.6.1 NRF failure handling

NRF services use a different mechanism for failure handling, because the NRF services are usually used system-wide instead of per session. Instead of per-operation failure handling, a failed NRF operation puts the NRF peer in a blocklist. Use the following command to configure the duration of the block for the applicable service:

- **Nnrf_NFManagement service**

```
configure mobile-gateway pdn sba-client-services nrf-client nnrf-nfm blocklist-duration
```

For the Nnrf_NFManagement service the behavior is granularly configurable using the following commands:

- The **heartbeat-nrf-blocklist** command configures how often to retry and wait for a heartbeat NfUpdate operation before blocklisting.
- The **nfm-nrf-blocklist** command configures how often to retry and wait for any other operation before blocklisting. This includes non-heartbeat NfUpdate operations, for example when updated SMF parameters have to be sent to the NRF.

The following specific status codes immediately block a peer, without retrying the Nnrf_NFManagement operation:

- For the NfRegister operation, any HTTP status code in the range 400 to 499, and HTTP status code 501 (Not Implemented)
- For the NfUpdate operation, including the heartbeat, HTTP status code 400 (Bad Request) and 501 (Not Implemented).

For the NfUpdate operation, including heartbeat, if an HTTP status code 404 (Not Found) is received, the MAG-c locally removes the registration and trigger a new NfRegister operation toward the same NRF peer.

If the MAG-c is registered with an NRF peer that is blocklisted, the MAG-c triggers a new NfRegister operation to the next peer, taking into account any configured NF Instance or IP priority.

- **Nnrf_NFDiscovery service**

```
configure mobile-gateway pdn sba-client-services nrf-client nnrf-disc blacklist-duration
```

For the Nnrf_NFDiscovery service, the peer is blacklisted immediately if a single operation is rejected or times out and the discovery procedures use other configured NRF peers.

When an HTTP 429 (Too Many Requests) status code is received, the blacklist duration is set to the value of the **retry-after** command header.

4.6.2 Operational state and CLI commands

The MAG-c supports retrieving service statistics per NF type, including failure statistics. Use the following command or state data to retrieve the statistics:

- **show** command

```
show mobile-gateway pdn service-stats
```

- operational state data

```
state mobile-gateway pdn sba-client-services
```

See [YANG state](#) for more information about the operational state data.

When the **ap-continue** failure handling mode is enabled, it is not always clear that the MAG-c performs failure handling, because the triggering procedures (for example, an attach) are successfully completed. Use the following operational state data and **show** commands to detect or debug this situation for the CHF, PCF, UDM SDM, and UDM UECM NF types:

- Use the following command or state data to display the total number of sessions and the number of sessions in **ap-continue** failure handling mode per NF type:

- **show** command

```
show mobile-gateway pdn apn fh-stats
```

- operational state data

```
state mobile-gateway pdn apn statistics failure-handling
```

- A log event can be generated when the percentage of sessions in failure handling for an NF type exceeds a configurable upper threshold (default 10%). A second event is generated when the percentage subsequently falls below a configurable lower threshold (default 5%). Use the following command to configure the thresholds.

```
configure mobile-gateway pdn failure-handling-alarm-threshold
```

- Use the following command or state data to display the available per-NF-type failure handling information for a session:

- **show** command

```
show mobile-gateway bng session
```

- operational state data

```
state mobile-gateway bng queries sessions session result
```

- Use the **fh-active-chf**, **fh-active-pcf**, **fh-active-udm-sdm**, and **fh-active-udm-uecm** filters to display only sessions in failure handling. The filters can be used in the following ways:

- as filter for the following **show** command

```
show mobile-gateway bng session
```

- as filter in the session state query configuration in the following context

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

When the **attempt-continue** failure handling mode is enabled for PCF, only sessions with initial SMPolicy creation failures are counted toward the **ap-continue** failure handling statistics. Sessions with subsequent SMPolicy update failures are not included in the **ap-continue** failure handling statistics.

4.7 SCP model D communication

In 5G Service Based Architecture (SBA), the model D method supports indirect communication via a service communication proxy (SCP) that mediates between the NF consumer and producer.

Implementation overview

In the model D method of indirect communication, the SCP is deployed between the NF consumer and the NF producer. The SCP performs the discovery of the NF producer via the NRF. After discovering and selecting the NF producer, further communication between the consumer and the producer only takes place via the SCP. There is never direct communication between the NF consumer and the NF producer while using the model D method.

When the SCP is deployed between the SMF and the UDM or the PCF:

- The SMF acts as the NF consumer.
- The UDM or the PCF acts as the NF producer.

The SMF adds the NF discovery parameters in the HTTP header and the SCP executes the actual NF discovery with the NRF. This is also known as the delegated-discovery method. Inclusion of discovery parameters in the HTTP headers follows the same rules and configuration as described in [NRF-based discovery](#). Some configurations require an NRF client context, but even with this context, NRF is not used for peer discovery in model D communication.

4.7.1 SCP client configuration

The SCP client configuration defines the communication model. Use the commands in the following context to configure the SCP client settings:

```
configure mobile-gateway pdn sba-client-services scp-client
```

4.7.2 SCP failure handling

The SMF can use the communication model D with an option to fallback to model A, or assume-positive response in the event of a failure from the SCP. The SCP failure handling is triggered if the SCP rejects the request, or it does not respond. In the case of a failure between the SMF and the SCP, the client-service configuration defines the communication model and the optional fallback mechanism.

During both a failure and timeout event from an SCP peer on a request, the SMF does not retry that SCP. If multiple SCP peers are available, the SMF makes only one more attempt to send the same request to another available SCP. This is the default behavior and the user cannot modify it through configuration.

Use the following commands to configure the mode and fallback options with the UDM (SDM and UECM).

```
configure mobile-gateway pdn sba-client-services udm-client nudm-sdm mode
configure mobile-gateway pdn sba-client-services udm-client nudm-uecm mode
```

Use the following command to configure the mode and fallback options with the PCF.

```
configure mobile-gateway pdn sba-client-services pcf-client npcfsmpolicycontrol mode
```

4.7.3 SCP load distribution

The MAG-c supports configuration of active load balancing across multiple SCP peers to enable SMF communication via multiple SCP nodes in parallel.

The SMF can communicate via multiple SCP nodes in parallel. Active load balancing is supported across multiple SCP peers such that the load is evenly distributed across all the peers. In the case of a failure of one SCP peer, the impacted sessions are redistributed to the remaining SCP peers.

Load balancing is implemented by default, based on the number of SCP addresses configured for the NF ID list profile that is referenced in the SCP client configuration. The load is distributed amongst the configured SCPs per service.



Note: Consult Nokia technical support for the number of SCP peers supported on the SMF.

5 Address assignment

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

5.1 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
- external DHCPv4 and DHCPv6 server 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.

5.2 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.

5.2.1 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
```



Note: The term micro-net in the documentation, state output, or show commands refers to the full prefix when using dedicated mode.

- **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
    -----
```

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 **primary-dns** and **secondary-dns** commands in the following contexts:

```
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](#)

5.2.2 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 following context.

```
configure mobile-gateway pdn local-address-assignment network-realm pool ipv6 pd delegated-prefix
```

- 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 following context.

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

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 following context, is adapted as follows.

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

- 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.

5.2.3 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, IPv6 NA) or the **exclude-prefixes** command (SLAAC, IPv6 PD) in the following contexts:

- ```
configure mobile-gateway pdn local-address-assignment network-realm pool ipv4 prefix
```
- ```
configure mobile-gateway pdn local-address-assignment 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 IPv4, 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:

- ```
configure mobile-gateway pdn local-address-assignment network-realm pool ipv4 prefix
```
- ```
configure mobile-gateway pdn local-address-assignment 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 from RADIUS, HSS, or UDM](#)

5.3 AAA-based address assignment from RADIUS, HSS, or UDM

AAA services such as RADIUS, HSS, or UDM 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 specific situations including 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 following context.

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

- In the absence of any pool during authentication, a pool can be provisioned per APN using the **unmanaged** command in the following context.

```
configure mobile-gateway pdn apn address-assignment-defaults
```

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 having a non-dedicated pool for AAA-based address assignment and a separate non-dedicated pool for local address assignment.

5.4 AAA framed routes

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. Framed routes are supported for all address assignment types, and not restricted to AAA-based address allocation.

5.5 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 following context.

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

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 following context.

```
configure mobile-gateway pdn apn
```

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](#)

5.6 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.

5.7 Address assignment for data-triggered IPoE sessions

Data-triggered IPoE sessions support the following address assignment methods:

- AAA-based static address assignment
- static address assignment from ADB
- using the source IP address of the data trigger

The addresses and prefixes assigned during data-triggered IPoE session setup are also used in DHCP/DHCPv6 address assignment when MAG-c receives DHCP/DHCPv6 packets from data-triggered IPoE sessions.

Use the commands in the following context to configure data-triggered source IP handling in address assignment.

```
configure mobile-gateway profile authentication-database entry address-assignment initial-ip-handling data-trigger
```

Use the following command to configure source IP options for data-triggered IPoE sessions. The **ignore** option for this command forces address assignment from the ADB entry or AAA. The **unmanaged-default** option populates an unmanaged address or prefix based on a source IP address.

```
configure mobile-gateway profile authentication-database entry address-assignment initial-ip-handling data-trigger source-ip
```

Configure the following command together with the **unmanaged-default** option, to define the source IPv6 address modeling in the IPoE session.

```
configure mobile-gateway profile authentication-database entry address-assignment initial-ip-handling data-trigger ipv6-address-assignment-type
```

5.8 External DHCPv4 and DHCPv6 server address assignment

An external DHCPv4/v6 server can assign a session address via MAG-c acting as a DHCPv4/v6 relay agent, which relays DHCP messages between the client and the external server. During authentication, only the ADB (and not RADIUS) can return DHCP relay configuration, including an external server address, a tracking pool, a DHCPv4 giaddr, and a DHCPv6 link address. The MAG-c uses that ADB configuration to relay the DHCP messages.



Note: Relaying a DHCP server initiated exchange (such as DHCPv4 Force-Renew or DHCPv6 Reconfigure) is not supported.

5.8.1 Local ODSA pool tracking

In case of DHCP relay, the external server assigns the address and the MAG-c tracks the assigned address, subnet, and default-gw to make sure there is no conflict in the assignments.

To track the DHCP relay address assignments, specify a local ODSA tracking pool under the DHCP relay configuration in the ADB. The subnets and prefixes configured in the tracking pool must be same as in external server. The tracking pool is not used to allocate micro-nets; it is a dedicated pool that can only be used to track the DHCP relay assignments.


Use the following command with the **tracking dhcp-relay** keywords to create a tracking pool. Use the commands in this context to configure the tracking pool.

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

5.8.2 DHCPv4 relay

The following figure shows the DHCPv4 call flow.

Figure 20: DHCPv4 call flow



The DHCPv4 server sends its responses to the giaddr. Therefore, the giaddr must be configured on a MAG-c IP interface that is reachable by the DHCPv4 server. Typically, this is a /32 address on a loopback interface.

If the giaddr is also used for pool or subnet selection by the server, the MAG-c must advertise a /32 route for the giaddr to make sure that the DHCP response from the server reaches the MAG-c, and not the BNG-UP.

Example: ADB configuration for DHCPv4 relay

```

A:MAG-c>config>mobile>profile>adb>entry# info
-----
      address-assignment
      dhcp-relay
      ipv4
      gi-address 21.21.0.1
      pool "p2"
      router "Base"
      server 172.16.20.100
      exit
    exit
  exit
-----
  
```

5.8.3 DHCPv6 relay

The MAG-c as a DHCPv6 relay server supports the following cases:

- The DHCPv6 client starts the call flow with a Router Solicit (RS) message.
- The DHCPv6 client does not send the RS message.


- The DHCPv6 client is behind a lightweight DHCPv6 relay agent (LDRA).

Call flow with RS message

The DHCPv6 client first sends an RS message. When the M bit in the Router Advertisement (RA) message is set, the client starts the DHCPv6 message exchange.

The following figure shows the DHCPv6 call flow for a client sending the RS.

Figure 21: DHCPv6 call flow with RS message




Call flow without RS message

The DHCPv6 client starts the DHCPv6 message exchange without first sending an RS message.

The following figure shows the DHCPv6 call flow when a client does not send the RS.

Figure 22: DHCPv6 call flow without RS message



Call flow behind an LDRA

The DHCPv6 client is behind an LDRA. The LDRA sends a relay-forward message to the MAG-c and the MAG-c replies with a relay-reply message.

Configuration specifics

The configuration of an RA profile in the ADB customizes the content of the RA message. Use the following command to configure an RA profile.

```
configure mobile-gateway profile authentication-database entry ra-profile
```

Use the following command to configure the source address of the relay-forward message to the server. The DHCPv6 server sends its response to the source address of the received request message.

```
configure mobile-gateway profile authentication-database entry address-assignment dhcp-relay
  ipv6 source-ip-address
```

Use the following command to configure a link address that the server can use to select a pool and prefix.

```
configure mobile-gateway profile authentication-database entry address-assignment dhcp-relay
  ipv6 link-address
```

Example: DHCPv6 relay configuration

```
A:MAG-c>config>mobile>profile>adb>entry>address-assignment# info
-----
                dhcp-relay
                ipv6
```

```

link-address 2001:9999::1
pool "p2"
router "Base"
source-ip-address 2001:beef::1
server 2001:beef::100
exit
exit
-----

```

5.8.4 DHCP and DHCPv6 for data-triggered IPoE sessions

For data-triggered IPoE sessions, MAG-c supports address assignment for IPv4, IPv6 NA, and IPv6 PD address stacks.

The following packets can trigger an address-assignment process with lease-state creation and signaled-address association with the IP stack created by a data-trigger.

- DHCP discover and request
- DHCPv6 solicit, request, renew, and rebind

DHCP and DHCPv6 release messages that are received on a data-triggered IP stack, trigger removal of the entire IPoE session.

Related topics

[Address assignment for data-triggered IPoE sessions](#)

5.8.5 Dual-stack relay

If the ADB returns both a DHCPv4 and a DHCPv6 relay configuration, the MAG-c treats this as a dual-stack session.

The MAG-c behavior when relaying dual-stack sessions is as follows:

- When the MAG-c receives a DHCPv4 Discover or a DHCPv6 Solicit message for a new session, it starts a 60s timer to guard the session setup time. The MAG-c relays the messages and waits until the message exchange with the external server for both stacks is finished (for example, until the reception of both DHCPv4 ACK and DHCPv6 Reply messages from the server) before starting the final PFCP session modification message exchange.

If the 60s timer expires before the message exchange for both stacks is finished, the session setup fails.

- When the MAG-c receives a DHCPv4 or DHCPv6 Release message, it sends both the DHCPv4 and DHCPv6 Release messages to the server.
- If the MAG-c initiates a session deletion, it sends both the DHCPv4 and DHCPv6 Release messages to the server.



Caution: The preceding behavior causes failure of a single-stack session setup when the ADB returns both a DHCPv4 and a DHCPv6 configuration. Therefore, for single-stack sessions, configure only a single-stack relay in the ADB.

Because the MAG-c removes both stacks at the same time, the DHCP server must use the same lease time for both DHCPv4 and DHCPv6 to avoid a stale lease on the client side.

5.8.6 DHCP options

DHCPv4

Use the following commands to configure a DHCPv4 profile for the specified direction:

- **from the MAG-c toward the DHCP server**

```
configure mobile-gateway profile authentication-database entry address-assignment dhcp-relay
  ipv4 options request-dhcp-profile
```

This profile specifies the DHCP options sent to the server, including relay agent options like circuit ID and remote ID.

- **from the MAG-c toward the DHCP client**

```
configure mobile-gateway profile authentication-database entry address-assignment dhcp-relay
  ipv4 options response-dhcp-profile
```

This profile specifies the DHCPv6 options sent to the client. Relay agent options do not apply in this direction.

DHCPv6

Use the following command to configure a DHCPv6 profile that specifies the DHCPv6 options for the relay-forward message.

```
configure mobile-gateway profile authentication-database entry address-assignment dhcp-relay
  ipv6 options request-dhcp6-profile
```

Because a relay agent cannot modify the encapsulated message (see RFC 8415), the MAG-c inserts the options at the top level of the generated relay-forward message, and not in the encapsulated message.

This profile can include relay agent options, such as interface ID and remote ID.

5.8.7 DHCP relay and BNG-UP redundancy



Note: The content of this chapter only applies if the following conditions are both met:

- BNG-UP redundancy is enabled.
- The server relies on the MAG-c insert option for pool/subnet and client identification.

If the external server uses a DHCP option inserted by the access node, not configuring any **request-dhcp-profile** or **request-dhcp6-profile** makes sure that the option inserted by the access node is not changed during BNG-UP switchover.

Access ports and VLANs on multiple BNG-UPs can be part of the same UP group in case of BNG-UP redundancy. Consequently, the same set of sessions may come from different ports, VLANs, or BNG-UPs depending on the active BNG-UP. Therefore, when the DHCP server assigns an address or prefix to a session with BNG-UP redundancy enabled and while relying on an option inserted by the MAG-c to select pool/subnet, it must take the UP group into consideration to avoid sharing subnets or prefixes between different UP groups.

Use one of the following options to prevent sharing subnets or prefixes between different UP groups:

- If the DHCP server uses the giaddr or link address to select a pool or subnet, configure one unique giaddr or link address per UP group on the MAG-c and use one pool or subnet per UP group on the DHCP server.
- Use the following commands to configure the UP group ID as circuit ID, remote ID, or interface ID in the DHCPv4 or DHCPv6 profile for sessions with BNG-UP redundancy enabled. The UP group does not change during BNG-UP switchover.

```
configure mobile-gateway profile bng dhcp-profile options relay-agent-options circuit-id
configure mobile-gateway profile bng dhcp-profile options relay-agent-options remote-id
configure mobile-gateway profile bng dhcpv6-profile options relay-agent-options remote-id
configure mobile-gateway profile bng dhcpv6-profile options relay-agent-options interface-id
```

- If the preceding options are not feasible, provision a Python script on the MAG-c that converts the UP group VSO into an option and format that the server supports.

For the DHCP renew to work, the client identification as seen by the DHCP server must be the same during the BNG-UP switchover; for example, in case of a BNG-UP switchover to a different BNG-UP, port, or VLAN, the client identification must be the same. When the DHCP server uses the circuit ID, remote ID, or interface ID as the client identification, use one of the following methods to obtain the same client identification:

- If the port ID is different on the BNG-UPs for a specific UP group, configure the same Layer 2 access ID alias on all BNG-UPs in the UP group, or use a Python script to change the circuit ID, remote ID, or interface ID to be the same.
- If the VLAN tags are different on different BNG-UPs, use a Python script to change the circuit ID, remote ID, or interface ID to be the same.

6 Authentication

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

6.1 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](#)

6.2 BNG entry point

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 sx-n4 signaling ibcp bng-entry-point
```



Note: Each Sx-N4 reference point can reference only one BNG EP; that is, BNG EP e1 and e2 must not both be referenced in the following context.

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

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

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

To configure the content of the BNG EP in the BNG profile, use the following command.

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

Example

The following example shows an EP configuration in the BNG profile.

```
A:BNG-CPF>config>mobile>profile>bng# info
-----
  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
exit
-----
```

6.3 Authentication database

Each authentication database (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 following context.

```
configure mobile-gateway profile bng
```

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 following context.

```
configure mobile-gateway profile authentication-database entry match
```

Configure the password using the **action local-auth** command in the following context.

```
configure mobile-gateway profile authentication-database entry
```

The MAG-c uses the session creation configuration parameters of all ADBs. The authentication flow contains an ordered list of ADBs. If ADBx comes before ADBy in the ordered list of ADBs, the values of the parameters in ADBy have priority over the values of the parameters in ADBx. 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 configures **bng-charging-profile discard**, the MAG-c creates the session without **bng-charging-profile**.

Related topics

[BNG EP and ADB lookup](#)

6.4 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](#)

6.5 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**), data packet field (for example, **source-ip-prefix**), 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 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 **string-mask** 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 **string-mask** 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

- **I2-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.

6.6 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 11: 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 ¹ | ADB, RADIUS |
| sub-profile ¹ | ADB, RADIUS |
| group-interface-template ¹ | ADB, RADIUS |
| sap-template ¹ | ADB, RADIUS |

¹ 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.

6.7 RADIUS authentication profile

RADIUS authentication is performed when the **action** parameter in the best-matched ADB entry is set to **radius**. The RADIUS authentication profile defines the RADIUS authentication behavior. To define the profile, use the **radius-authentication-profile** command in the following context.

```
configure mobile-gateway profile bng
```

RADIUS authentication is triggered by the ADB lookup; consequently, it is possible to have multiple rounds of RADIUS authentication during the authentication flow lookup. If the same attributes are returned in the Access-Accept message during multiple authentication rounds, 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 **configure mobile-gateway profile** context. Reference this RADIUS group in the following context.

```
configure mobile-gateway profile bng radius-authentication-profile
```

- **user-name-format**

Use the following command to define the username format for the RADIUS server.

```
configure mobile-gateway profile bng radius-authentication-profile user-name-format
```

Use the following command option to send the source IP address of the data-trigger packet.

```
configure mobile-gateway profile bng radius-authentication-profile user-name-format ipoe  
format data-trigger-source-ip
```

- **password**

Use the following command to define the password of the RADIUS user.

```
configure mobile-gateway profile bng radius-authentication-profile password
```

- **include-attribute**

The **include-attribute** command defines the RADIUS attributes to be included in an Access-Request message. Use the commands in the following context to define the attributes to include.

```
configure mobile-gateway profile bng radius-authentication-profile include-attribute
```

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

6.8 RADIUS CoA and DM

A RADIUS Change of Authorization (CoA) message or a Disconnect Message (DM) asks for changes in the session or subscriber object.

To enable support for RADIUS CoA and DM messages, use the **interface** command in the following context.

```
configure mobile-gateway pdn bng radius-coa
```

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 *MAG-c RADIUS Attributes and IU Triggers* list defines the message attributes that can be used to identify one or multiple sessions as target object. Filter on the value **True** 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 *MAG-c RADIUS Attributes and IU Triggers*.

If the MAG-c applies all requested changes successfully to all targeted objects, the MAG-c sends a CoA-ACK message back to the RADIUS server. 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 with an ERROR-CAUSE code 506, 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 following context.

```
configure mobile-gateway pdn bng radius-coa
```

Use the following command to manually send a CoA request that can be used to test or modify the attributes of the existing session.

```
tools perform mobile-gateway profile bng radius coa
```

6.9 Example configuration

The example configurations in this section are 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

This setup uses an authentication flow with three ADBs for which the following are 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.

Example: ADB configuration with three ADBs

```

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 entry point (EP).

Example: BNG EP configuration

```

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 a reference to the BNG EP configured in the following context.

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

Example: BNG EP reference configuration

```
A:BNG-CPF>config>mobile>pdn>sx-n4# info
```

```

-----
pfcf-association-list "pfcpassoc1"
interface
  pfcf "system"
  ibcp "system"
exit
signaling
  pfcf
    profile "pfcfpro-1"
  exit
  ibcp
    bng-entry-point "e1"
    triggers ipoe-dhcp
  exit
exit

```

6.10 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 23: 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 RADIUS Alc-Portal-Url VSA, the value of the RADIUS attribute is used for HTTP redirection. The URL in the RADIUS Alc-Portal-Url VSA 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 MAG-c 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`

6.10.1 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.

```
configure mobile-gateway pdn bng wpp interface
```

Step 2. Define a WPP portal.

```
configure mobile-gateway profile bng wpp portal
```

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

```
configure mobile-gateway profile bng wpp portal-group
```

Step 4. Define a RADIUS authentication profile.

```
configure mobile-gateway profile bng radius-authentication-profile
```

Step 5. Configure the WPP context in an ADB entry.
Use the following commands in the following context.

```
configure mobile-gateway profile authentication-database entry wpp
```

- 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.

```
configure mobile-gateway profile authentication-database entry http-redirect url
```

Step 7. Configure the after-auth profiles in an ADB entry.
Define the after-auth SLA and subscriber profiles.

```
configure mobile-gateway profile authentication-database entry subscriber-mgmt
```

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  
-----
```

7 PCF-based policy management

The Policy Control Function (PCF) is used to retrieve policy operations related to traffic forwarding, QoS, and charging. This is a pure policy-management function and does not include any authentication functionality, in contrast for example to RADIUS, where authentication and policy management (authorization) are often combined. The PCF uses the SMPolicyControl service. See [Service Based Interfaces](#) for more information.

To enable PCF communication, perform the procedure [Enabling PCF communication](#).

When a PCF client service is configured, FWA sessions automatically attempt to contact the PCF. Use the following command to configure a dynamic policy in the ADB and bypass PCF for specific sessions.

```
configure mobile-gateway profile authentication-database entry dynamic-policy
```

You can also configure per DNN parameters using commands in the following context:

```
configure mobile-gateway pdn apn pcf-selection
```

Options to configure include the following:

- **nf-id-list** and **npcf-profile** commands to override their equivalents in the client service instance
- **dnn-format** command to modify how the DNN is sent to the PCF
- **exclude** commands to exclude sending specific IEs to the PCF

PCF interaction is always performed after the authentication flow procedure, including any RADIUS authentication, is completed. The PCF can provide multiple policy decisions, such as an SLA profile, subscriber profile, subscriber filters, and QoS overrides. If any of these values are also provided during the authentication flow, the PCF value overrides these.

Nokia supports vendor-specific attributes in the SMPolicyDecision context, using the vendorSpecific-6527 container. The following attributes are supported:

- subscriberProfile – string indicating the subscriber profile
- slaProfile – string indicating the SLA profile
- sapTemplate – string indicating the SAP template
- groupInterfaceTemplate – string indicating the group interface template
- intermediateDestinationID – string indicating an intermediate destination ID, typically used to apply VPORT QoS
- QoSOverride – string representing specific QoS overrides, using the same format as the RADIUS Al-Subscriber-QoS-Override attribute which can be present up to 18 times
- A subscriber filter – container that can contain the following sub-objects:
 - ingressIpv4 – string representing an IPv4 filter that is applied on ingress
 - egressIpv4 – string representing an IPv4 filter that is applied on egress
 - ingressIpv6 – string representing an IPv6 filter that is applied on ingress
 - egressIpv6 – string representing an IPv6 filter that is applied on egress

The MAG-c can retrieve an updated policy from the PCF using the Npcf_SMPolicyControl Update service call. The SMF only does this if the PCF sets triggers when providing a policy decision. The following trigger codes are supported:

- RES_MO_RE – a PDU Session Modification Request has been received from the UE to add or remove a policy and charging control (PCC) rule
- UE_IP_CH – an IP address is allocated or freed using DHCP/DHCPv6
- DEF_QOS_CH – default QoS parameters are changed via another interface; for example, because of a UDM subscription update
- SE_AMBR_CH – the Session AMBR is changed via another interface; for example, because of a UDM subscription update
- SAREA_CH – a location change was detected that changes the Serving Area
- SCNN_CH – a location change was detected that changes the Serving CN Node
- RE_TIMEOUT – the PCF provided a revalidation timer that triggered the update
- REF_QOS_IND_CH – support for reflective QoS is changed in a PDU Modification Request message
- SCELL_CH – a location change was detected that changed the Serving Cell
- CM_SES_FAIL – a credit control failure was detected; see [CHF failure handling](#) for more information

Alternatively, the PCF can directly update the policy on the MAG-c using the SmPolicyUpdateNotification callback service call. This call is executed against an explicit URI provided by the MAG-c during setup, which contains all required session context.

The Npcf_SMPolicyControl service creates a policy association for a session that can be updated or removed at any time. Each policy association has its own identifier assigned by the PCF, which may be different from MAG-c session-specific identifiers such as SUPI. The MAG-c supports the following operations of the Npcf_SMPolicyControl service, as shown in [Figure 24: PCF Policy Control Operations](#):

- The Create operation is used to create the policy control association. In the request, the MAG-c sends various metadata such as SUPI, GPSI, PEI, DNN, S-NSSAI, and subscribed QoS. The MAG-c also includes a notification callback URI for the UpdateNotify operation. The PCF responds to this request with an association ID identifying the policy association, and the initial policy data.

The Create operation is part of the initial session setup; see [FWA session setup for 5G](#) for more information.


- The Update operation that is used to inform the PCF that one or more signaled triggers have been activated. The MAG-c signals the policy association ID, applicable update trigger, and any metadata to the PCF. The PCF responds to the update with the changed policy data, if any.
- The UpdateNotify operation is used by the PCF when it requires to update a policy decision immediately. In this case it sends the request to the URI provided during the creation of the policy association, including the changed policy data. The MAG-c applies the changed policy and sends a response.



Note: This operation does not use the association ID; the MAG-c provides any local identifier it requires as part of the callback URI.

- The Delete operation is used by the MAG-c to stop policy enforcement and remove the policy association. This operation is executed during FWA session deletion.

Figure 24: PCF Policy Control Operations



sw1404

7.1 Enabling PCF communication

About this task

The PCF is used to retrieve policy operations related to traffic forwarding, QoS, and charging. Configure and enable a PCF SMPolicyControl client service to enable PCF communication.

Procedure

Step 1. Configure a PCF SMPolicyControl client service.

```
configure mobile-gateway pdn sba-client-services pcf-client npcf-smpolicycontrol
```

Step 2. Configure an interface for the MAG-c to use to communicate with the PCF.

```
configure mobile-gateway pdn sba-client-services pcf-client npcf-smpolicycontrol  
interface
```

Step 3. Provide the following optional configuration.

- An FQDN to identify the service by other NF peers. For example, when generating callback functions, the MAG-c can use the FQDN in the URI, instead of the IP address; see [URI construction](#) for more information about how URIs are constructed.
- An NF ID list to specify NF instances that can be used for the PCF function; see [Enabling discovery based on a local NF ID list](#) for more information. This command is required if the PCF is not dynamically discovered using NRF.
- An HTTP2 profile to influence generic HTTP/2 and SBI parameters; see [HTTP/2, OpenAPI, and Python](#) for more information.
- A PCF profile that contains parameters to handle PCF failure handling as configured using the following command. See [Failure handling](#) for more information.

```
configure mobile-gateway profile npcfc
```

This also allows configuration of feature support signaling to the PCF using the **supported-features** commands.

Step 4. Enable the client service.

```
configure mobile-gateway pdn sba-service-realm client-service
```

7.2 PCC rule retrieval

The PCF can install PCC rules using the following methods:

- **dynamic PCC rules**
The PCF fully defines all parameters of the PCC rules, including the packet classifiers. There is no static preconfigured aspect involved in these rules.
- **predefined PCC rules**
The PCF only defines a PCC rule ID, but no parameters. The PCC rule is assumed to be predefined on the FWA gateway.


For both methods, the PCF can provide a list of PCC rules to enforce. The PCF can add, modify, or remove PCC rules, unsolicited or triggered by the SMF, in any policy association update message. The MAG-c re-evaluates the set of PCC rules and passes through any necessary changes to the other functions, such as the RAN or the UPF.

The UE can send a PDU Session Modification Request message with an explicit request to create a PCC rule toward the MAG-c. The MAG-c never accepts or installs autonomously a PCC rule that is requested by a UE. The following exchange takes place:

- The MAG-c reflects the request to the PCF using the RES_MO_RE trigger.
- The PCF can accept the PCC rule creation and reflects it to the MAG-c as a regular PCF-provisioned PCC rule.
- The MAG-c sends a PDU Session Modification Response message confirming the new PCC rule toward the UE.

The following figure shows the message exchange between the involved entities.

Figure 25: PCC rule request initiated by the UE



For dynamic PCC rules, the PCF signals the full PCC rule, including the following parameters:

- precedence
- SDF templates for the 5-tuple-based or TOS field-based filtering
- 5QI, ARP, RQI, and bind-to-default QoS flow indicators
- MBR and GBR values
- 5QI QoS characteristics overrides to pass to the RAN
- gate status to control forwarding of the PCC rule

Dynamic PCC rules support the time-of-day (TOD) activation and deactivation by passing condition data in the PCC rule. The MAG-c automatically enables or disables the PCC rule based on the TOD information.

Related topics

[Dynamic QoS based on PCC rules](#)

7.3 Predefined PCC rules

The MAG-c supports dynamic and predefined PCC rules enforcement on the SMF/GW-C and UPF/GW-U.

The MAG-c supports two methods for policy enforcement on the SMF/GW-C and UPF/GW-U:

- **dynamic PCC rules**

The PCF provisions dynamic PCC rules by sending the complete rule description to the SMF/GW-C.

- **predefined PCC rules**

The PCF references predefined PCC rules configured on the SMF/GW-C, to activate or deactivate session-level or SDF-level policies.

Configure predefined PCC rules as follows:

- Use the following command to configure individual policy rules.

```
configure mobile-gateway profile policy-options policy-rule
```

- Use the commands in the following context to configure a group of rules as a collection of policy rules.

```
configure mobile-gateway profile policy-options policy-rule-base
```

- Use the commands in the following context to configure policy rule units (PRUs).

```
configure mobile-gateway profile policy-options policy-rule-unit
```

To activate and deactivate predefined policy rules for a PDN session, the PCF references the corresponding policy rule names and policy rule bases configured on the gateway. Use the commands in the following context to configure static predefined PCC rules under an APN.

```
configure mobile-gateway pdn apn static-predefined-pcc
```

Static predefined PCC rules are applied to new PDU or PDN sessions that have no PCF control. A PDU or PDN session may be without PCF control, for example, if the PCF is not deployed or the PCF is unreachable and a fallback to the local policy occurs for new sessions.

7.3.1 Evaluation of predefined PCC rules during session setup

During session setup, the MAG-c evaluates rules referenced by the PCF (policy rule bases, policy rule names, or both), and applies them based on the rule and priority.

The MAG-c evaluates predefined PCC rules that the PCF sends during the session setup as follows:

- If the PCF sends a reference to a policy rule base during the session setup, the MAG-c evaluates all the policy rules configured under the policy rule base for the final set of policy rules to apply. If multiple policy rules match, the MAG-c applies the one with the highest priority (lowest precedence value) for that session.
- If the PCF sends a reference to one or more policy rule names during the session setup, the MAG-c evaluates the policy rule against other policy rules sent by the PCF. If multiple policy rules match, the MAG-c applies the one with the highest priority (lowest precedence value) for that session.
- If the PCF sends a reference to policy rule base and a policy rule name during the session setup, the MAG-c evaluates the policy rules configured under the policy rule base and the policy rules referenced by the PCF. If multiple policy rules match, the MAG-c applies the one with the highest priority (lowest precedence value) for that session.

See [Predefined PCC rules configuration example](#) and [Use cases for predefined PCC rules](#) for more information about the session setup behavior when the PCF activates and deactivates predefined PCC policy rules.

7.3.2 Evaluation of predefined PCC rules during session update

During session update, the MAG-c evaluates policy rules referenced by the PCF, together with rules referenced during the previous session setup or update. Rules with the highest priority are applied for the session.

If the PCF sends a reference to policy rule base and policy rule name during a session update, the MAG-c evaluates the policy rules configured under the policy rule base and the policy rules referenced by the PCF, together with the policy rules configured under the policy rule base and the policy rules referenced previously during the session setup or update. If multiple policy rules match, the MAG-c applies the one with highest priority (lowest precedence value) for that session.



Caution:

Nokia recommends deleting the policy rule base and policy rule name during the update by the PCF, in case the intention is to not apply them on existing sessions. Configured precedence for policy rules applies as long as the PCF does not remove the policy-rule references. Relying

on the configured precedence to implicitly remove the policy rules could have unintended consequences.

If the PCF is not available, static policies predefined on the SMF can be used as a fallback option.

See [Predefined PCC rules configuration example](#) and [Use cases for predefined PCC rules](#) for information about the session update behavior when the PCF activates and deactivates predefined PCC policy rules.

7.3.3 Predefined PCC rules configuration example

This topic provides an example of policy options configuration.

The following example shows the configuration of policy options, including policy rules, policy rule units (PRUs), and action rule units (ARUs). The predefined PCC rules are applied at the APN level, including the policy rules, PRUs, and ARUs.

Example: PCC rules configuration

```
A:MAG-c>config>mobile>profile>policy-options# info
  policy-rule-unit "pru-any-normal"
    flow-description 1
    exit
    qos
      down-link gbr 5000 mbr 5000
      up-link gbr 5000 mbr 5000
    exit
  exit
  policy-rule-unit "pru-any-degraded"
    flow-description 1
    exit
    qos
      down-link gbr 250 mbr 250
      up-link gbr 250 mbr 250
    exit
  exit
  policy-rule-unit "pru-dns-normal"
    flow-description 1
    match
      remote-port eq 53
    exit
    qos
      down-link gbr 5000 mbr 5000
      up-link gbr 5000 mbr 5000
    exit
  exit
  policy-rule-unit "pru-dns-degraded"
    flow-description 1
    match
      remote-port eq 53
    exit
    qos
      down-link gbr 250 mbr 250
      up-link gbr 250 mbr 250
    exit
  exit
  action-rule-unit aru-edr-only
    app-profile-name app-prof-edr-only
  exit
  action-rule-unit aru-edr-redirect
```

```

    app-profile-name app-prof-edr-redirect
  exit
  action-rule-unit aru-edr-aa-rate-limit
    app-profile-name app-prof-edr-aa-rate-limit
  exit
  policy-rule "pr-normal-edr-only" policy-rule-unit "pru-any-normal" \
    action-rule-unit "aru-edr-only" qci * arp * precedence 65000
  policy-rule "pr-normal-edr-redirect" policy-rule-unit "pru-any-normal" \
    action-rule-unit "aru-edr-redirect" qci * arp * precedence 1300
  policy-rule "pr-degraded-edr-only" policy-rule-unit "pru-any-degraded" \
    action-rule-unit "aru-edr-only" qci * arp * precedence 1200
  policy-rule "pr-degraded-edr-redirect" policy-rule-unit "pru-any-degraded" \
    action-rule-unit "aru-edr-redirect" qci * arp * precedence 1100
  policy-rule "pr-normal-5-tuple" policy-rule-unit "pru-dns-normal" \
    qci * arp * precedence 200
  policy-rule "pr-degraded-5-tuple" policy-rule-unit "pru-dns-degraded" \
    qci * arp * precedence 100
  policy-rule-base "prb-all-cases"
    policy-rule "pr-normal-edr-only"
  exit

```

7.3.4 Use cases for predefined PCC rules

Review these use cases to understand the behavior when the PCF references and applies PCC rules during session setup and update.

This topic describes use cases for predefined PCC rules.

Update the application profile to enable redirect

The following table lists session setup and update behavior when the PCF applies policy rules to update the application profile (app-profile) to enable redirect.

Table 12: Update the app-profile to enable redirect

| PCC rules applied | Session setup | Session update |
|---|--|--|
| Using policy rules only | When the PCF sends the policy rule "pr-normal-edr-only": <ul style="list-style-type: none"> Policy-rule "pr-normal-edr-only" is activated App-profile "app-prof-edr-only" is activated | When the PCF sends the policy rule "pr-normal-edr-redirect": <ul style="list-style-type: none"> Policy rule "pr-normal-edr-redirect" is activated App-profile "app-prof-edr-redirect" is activated |
| Using the configured policy rule base and policy rule | When the PCF sends the policy rule base "prb-all-cases": <ul style="list-style-type: none"> Policy-rule "pr-normal-edr-only" is activated App-profile "app-prof-edr-only" is activated | When the PCF sends the rule name "pr-normal-edr-redirect": <ul style="list-style-type: none"> Policy rule "pr-normal-edr-redirect" is activated App-profile: "app-prof-edr-redirect" is activated |

Update policy to throttle all traffic

The following table lists the session setup and update behavior when the PCF applies policy rules to throttle all traffic.

Table 13: Update policy to throttle all traffic

| PCC rules applied | Session setup | Session update |
|--|---|--|
| Using policy rules only | When the PCF sends the policy rule “pr-normal-edr-only”: <ul style="list-style-type: none"> Policy-rule “pr-normal-edr-only” is activated App-profile “app-prof-edr-only” is activated | When the PCF sends the policy rule “pr-degraded-edr-only”: <ul style="list-style-type: none"> Policy rule “pr-degraded-edr-only” is activated App-profile “app-prof-edr-only” remains unchanged |
| Using policy-rule-base and policy-rule | When the PCF sends the policy rule base “prb-all-cases”: <ul style="list-style-type: none"> Policy-rule “pr-normal-edr-only” is activated App-profile “app-prof-edr-only” is activated | When the PCF sends the rule name “pr-degraded-edr-only”: <ul style="list-style-type: none"> Policy rule “pr-degraded-edr-only” is activated App-profile: “app-prof-edr-only” remains unchanged |

Update policy to throttle selective traffic

The following table lists the behavior when the PCF applies policy rules to throttle selective traffic, for example, all traffic except DNS traffic.

Table 14: Update policy to throttle selective traffic

| PCC rules applied | Session setup | Session update |
|--|--|--|
| Using policy rules only | When the PCF sends the policy rules “pr-normal-edr-only” and “pr-normal-5-tuple”: <ul style="list-style-type: none"> Policy-rule “pr-normal-edr-only” and “pr-normal-5-tuple” are activated App-profile “app-prof-edr-only” is activated | When the PCF sends the policy rule “pr-degraded-edr-only”: <ul style="list-style-type: none"> Policy rule “pr-degraded-edr-only” is activated and “pr-normal-5-tuple” remains unchanged App-profile “app-prof-edr-only” remains unchanged |
| Using policy-rule-base and policy-rule | When the PCF sends the policy rule base “prb-all-cases” and policy-rule “pr-normal-5-tuple”: <ul style="list-style-type: none"> Policy-rule “pr-normal-edr-only” and “pr-normal-5-tuple” are activated App-profile “app-prof-edr-only” is activated | When the PCF sends rule-name “pr-degraded-edr-only”: <ul style="list-style-type: none"> Policy rule “pr-degraded-edr-only” is activated and “pr-normal-5-tuple” remains unchanged App-profile “app-prof-edr-only” remains unchanged |

7.4 PCF failure handling

The MAG-c supports failure handling for PCF as defined in [Failure handling](#). During failure handling using the **ap-continue** option, the MAG-c performs the following actions for the ongoing session:

- If the failure is mid-session, the MAG-c internally terminates the existing SM policy association and sends no further triggers to the PCF, even if the communication with the PCF recovers. The MAG-c does not send any SM Policy Delete procedures to the PCF, even if the communication with the PCF recovers before the session terminates.
- Use the following command to configure the PCC rules to keep in the assume positive (AP) mode.

```
configure mobile-gateway profile pcf-profile ap-action
```

The following options are available for this command:

- **keep-pcf-rules**
The MAG-c keeps the latest rules received from the PCF. If the PCF failure occurs during session setup, there are no PCC rules to keep and this option behaves the same as the **keep-local-rules-only** option.
- **keep-local-rules-only**
The MAG-c removes all PCC rules signaled by the PCF and only keeps locally-defined rules for the session, that is, the internal any-to-any rule described in [Dynamic QoS based on PCC rules](#) and the static predefined PCC rules described in [Predefined PCC rules](#).
- An optional time or volume limit or both can be specified during AP mode. These limits do not account for any time or volume before the failure. The MAG-c terminates the sessions as soon as a limit is reached. Use the following commands to configure the limits.

```
configure mobile-gateway profile pcf-profile fh-time-limit  
configure mobile-gateway profile pcf-profile fh-volume-limit
```

8 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.

8.1 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 following context.

```
configure mobile-gateway profile charging
```

To use a BNG charging profile for a specific set of sessions, use the **bng-charging-profile** command in the following context.

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

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 to 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](#)

8.2 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. The usage report only includes statistics 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, statistics collected since the previous report are lost. However, the MAG-c statistics remain correct and increase monotonically. Similarly, when BNG-UP resiliency is used, the MAG-c can add counters of both BNG-UPs to calculate the 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 26: Statistics collection using the pull model



- **push model**

The BNG-UP sends an unsolicited PFCP Usage Report IE in a PFCP Session Report Request message. Upon receiving this message, the MAG-c records and stores the statistics. The push message does not trigger an update (for example, Radius Accounting Interim Update or CHF Nchf_ConvergedCharging_Update). However, the statistics update may trigger any reporting limit, threshold, or quota; for example, a [CHF offline charging volume limit](#) or [CHF online charging quota](#).

The unsolicited usage report can either be periodic, or threshold- or quota-based. Both push models can be enabled simultaneously as follows:

- **periodic push model**

The BNG-UP sends a report within a fixed interval. If another report was generated (for example, pull-based or threshold-based), the interval is reset.

- **threshold- or quota-based push model**

The MAG-c instructs the BNG-UP to send a report when the amount of collected statistics reaches a specific threshold or quota. For online charging use cases, the MAG-c can additionally instruct the BNG-UP to take immediate action when a specific quota is reached; for example, stop forwarding. See [Online Charging](#) for more information.




Note: Because a URR can serve multiple applications, the threshold or quota values installed on the FWA-UP do not always match a specific application limit, quota, or threshold value. For example, a volume threshold of 1 GB may be installed by online charging, while a volume limit of 300 MB is installed by offline charging. Initially, the MAG-c installs a 300 MB threshold on the BNG-UP, matching the offline charging limit. However,

after the third report 900 MB is consumed, the MAG-c installs a 100 MB threshold on the BNG-UP to match the remainder (1 GB - 3 × 300 MB) of the online charging threshold.

The following figure shows the statistics collection using the push model.

Figure 27: Statistics collection using the push model



To enable the periodic 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 disabling 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 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 potentially 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** is enabled (set to **true**) 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 is 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](#)

8.3 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 following context.

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

The threshold values can also be provided via RADIUS. For more information, see *MAG-c RADIUS Attributes and IU Triggers*.

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](#)

8.4 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.

Use the **radius-group** command in the following context to reference the RADIUS group in the BNG charging profile.

```
configure mobile-gateway profile charging bng-charging radius
```

Use the **radius-group** command in the following context to define and configure the RADIUS group parameters.

```
configure mobile-gateway profile
```

- **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 following context for the session accounting parameters.

```
configure mobile-gateway profile charging bng-charging radius
```

For more information about the accounting attributes, their content, the associated include-attribute configuration, and the messages they can appear in, see the *MAG-c RADIUS Attributes and IU Triggers* and the *MAG-c CLI Reference Guide*.

8.4.1 Enabling RADIUS accounting

Procedure

Step 1. Define a BNG charging profile.

```
configure mobile-gateway profile charging bng-charging
```

Step 2. Configure RADIUS accounting for the BNG charging profile

```
configure mobile-gateway profile charging bng-charging radius
```

Step 3. Reference the BNG charging profile.

```
configure mobile-gateway profile authentication-database entry charging bng-charging-profile
```

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](#)

8.4.2 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, a new server is 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 Accounting Request 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 VSA.



Note: Nokia recommends provisioning 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 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](#)

8.4.3 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 following context.

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

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 following context.

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

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 following context.

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

To enable buffering of one Accounting Start message per session, use the **start** command in the following context.

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

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 following context.

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

- 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](#)

8.5 CHF charging

The guidelines in this chapter provide the basic requirements and options for configuring the CHF function common interface Nchf_ConvergedCharging SBI service.

In a 5G core network, the Charging Function (CHF) provides both online and offline charging functionality using a common interface Nchf_ConvergedCharging SBI service. The Nchf_ConvergedCharging SBI service can offer online-only charging, offline-only charging, or simultaneous online and offline charging. See [Service Based Interfaces](#) for more information about SBI service communication.

8.5.1 Enabling CHF charging

About this task

This task describes the steps to enable a basic Nchf_ConvergedCharging SBI service.

To enable CHF communication, perform the following steps:

Procedure

Step 1. Configure a CHF ConvergedCharging client service.

```
configure mobile-gateway pdn sba-client-services chf-client nchf-convergedcharging
```

Step 2. Configure an interface for the MAG-c to use to communicate with the CHF server.

```
configure mobile-gateway pdn sba-client-services chf-client nchf-convergedcharging
interface
```

Step 3. Use commands in the following context to optionally configure the functionality indicated:

```
configure mobile-gateway pdn sba-client-services chf-client nchf-convergedcharging
```

- Configure an FQDN to identify the service by other NF peers. For example, when generating callback functions, the MAG-c uses the FQDN in the URI, instead of the IP address; see [URI construction](#) for more information.
- Configure an HTTP2 profile to influence generic HTTP/2 and SBI parameters; see [HTTP/2, OpenAPI, and Python](#) for more information.
- Configure an API prefix that indicates the API version used in construction of any URIs. This only overrides the URI values and does not change the MAG-c behavior. See [URI construction](#) for more information.

8.5.2 CHF peer discovery

CHF peer discovery is implemented using a list of mechanisms in priority order.

The CHF extends the peer discovery (defined in [NF peer discovery](#)) with two additional sources for peer discovery. The CHF peer discovery is implemented in the following priority order:

1. The PCF returns a primary and secondary CHF address in the primaryChfAddress and secondaryChfAddress attribute fields of the ChargingInformation data type when creating an SM policy. See [PCF-based policy management](#) for more information.
2. A charging profile NF ID list is used, if configured using the following command.

```
configure mobile-gateway profile charging bng-charging chf chf-selection nf-id-list
```

This command has the following purposes:

- If the PCF signals a primary or secondary address, this list picks up the NF instance ID associated with that address. If this list is not configured or the peer is not found, the default UUID 11111111-0000-0000-0000-111111111111 is used. The MAG-c uses the UUID, for example, for statistics collection.



Note: If the default UUID is used by multiple CHF peers, per-peer statistics are not available.

- If the PCF does not provide a primary or secondary address, this list provisions a specific set of CHF peers associated with this profile, which bypasses any NRF discovery.
3. Use the following command to configure the NF ID list. Discovery based on the NRF or local NF ID list is used as defined in [NF peer discovery](#).

```
configure mobile-gateway pdn apn charging nchf chf-selection nf-id-list
```

8.5.2.1 Selecting predefined CHF peers based on charging characteristics

About this task

A common use case for CHF peer selection is to select a set of predefined CHF peers based on the charging characteristics (CC) values returned by the UDM for FWA sessions.

To configure this on the MAG-c, do the following:

Procedure

- Step 1.** For each CC value, create a BNG charging profile with an NF ID list that corresponds to the CC value.

```
configure mobile-gateway pdn apn charging nchf chf-selection nf-id-list
```

- Step 2.** Create an ADB and configure it to match the CC value configured in step 1; see [Authentication database](#) for more information.

- Step 3.** Under the ADB created in step 2, add an entry for each CC value you want to map, and under that entry apply the corresponding BNG charging profile.

```
configure mobile-gateway profile authentication-database entry charging bng-charging-profile
```

- Step 4.** Apply the ADB as one of the ADBs in the FWA authentication flow. See [Selected and real APN](#) for more information about how to configure the authentication flow for FWA sessions.

8.5.3 Charging sessions and rating groups

The steps in [Starting a charging session toward the CHF](#) describe how to establish a charging session between a MAG-c and CHF. However, any charging itself is performed at the level of a Rating Group (RG).

The following applies to charging RGs:

- RGs can be provisioned locally (see [Session charging using a session RG](#)).
- Each RG defines exactly which traffic is subject to charging and which units (time or volume) are monitored.
- Each RG has a unique RG ID to identify it within a session.
- Each RG is subject to either offline or online charging. The PCF decides this by setting the online or offline charging IE. If the PCF does not signal the IE or the PCF interaction is disabled, the MAG-c follows the configuration of the **default-charging-method** command under the BNG charging profile. It is possible to simultaneously enable offline and online charging for the same RG.


A charging session uses the following operations of the Nchf_ConvergedCharging service:

1. The Nchf_ConvergedCharging_Create operation starts the charging session. The MAG-c initiates this operation with an HTTP/2 POST message that includes the following data:
 - session metadata such as IMSI, MSISDN, APN, ULI, RAT, and S-NSSAI
 - a list of RGs subject to charging and, if online charging is enabled, includes the requested units (by default set to zero)

- a notification callback URI for the Notify operation
2. When the charging session starts successfully, the CHF responds with a 201 Created status code that includes the following data:
 - a CHF session URI (in the location header) that is used to identify the charging session on the CHF
 - session-level trigger conditions for the Update operation
 - the same list of RGs with RG-level trigger conditions for the Update operation and if online charging applies, also includes the granted quota units for that RG
 3. The Nchf_ConvergedCharging_Update operation is used for any updates. The Update operation is performed when an installed trigger condition is fulfilled or explicitly triggered via a Notify operation. See [Charging update triggers](#) for more information. The MAG-c initiates this operation with an HTTP/2 POST message toward the CHF session URI to update the current charging session. The message includes the following data:
 - trigger conditions that are fulfilled
 - session meta data
 - list of RGs that are subject to the fulfilled trigger conditions and, depending on the trigger type, signals for used units only or also requests for new updated quota units
 4. When successful, the CHF responds with a 200 OK status code that includes the following:
 - list of updated trigger conditions
 - list of RGs with new granted quota units, if any
 5. The CHF uses the Nchf_ConvergedCharging_Notify operation to indicate the charging session changed and the MAG-c should either abort or reauthorize charging. The CHF initiates this operation with an HTTP/2 POST message toward the MAG-c notification callback URI. The message indicates whether charging should be terminated or reauthorized.
 6. When the Notify operation is successful, the MAG-c responds with a 204 No Content status code. The MAG-c subsequently triggers the Update procedure. See [Charging update triggers](#) for more information.
 7. The Nchf_ConvergedCharging_Release operation gracefully terminates the charging session. The MAG-c initiates this operation with an HTTP/2 POST message toward the CHF session URI. The message includes the final used units for all RGs.
 8. The CHF confirms with a 204 No Content status code that the charging session is removed.

The following figure shows a high-level flow for the Create, Update, and Release operations. See [Charging update triggers](#) for information and an example of the Notify operation.

Figure 28: Charging Session operations



8.5.3.1 Starting a charging session toward the CHF

About this task

To start a charging session toward the CHF, perform the following steps:

Procedure

Step 1. Configure a BNG charging profile.

```
configure mobile-gateway profile charging bng
```

Ensure the BNG charging profile is used as defined in [BNG charging profiles](#).

Step 2. Use the commands in the following context to configure the optional settings as indicated:

```
configure mobile-gateway profile charging bng-charging chf
```

- Configure the **default-charging-method** command to determine which charging method is signaled when creating an SM policy association toward the PCF using online and offline IEs. See [PCF-based policy management](#) for more information.
- Configure the **dnn-format** command to determine which DNN is sent to the CHF. By default the selected DNN is used.
- Configure the **chf-profile** command to define any service communication behavior, including failure handling.
- Configure the **nf-id-list** command to select a specific set of peers for the CHF profile that takes precedence over NRF-based discovery. See [CHF peer discovery](#) for more information.

Step 3. Use the **no shutdown** command to enable CHF charging for the session.

8.5.4 Session charging using a session RG

The MAG-c supports charging for all traffic of a single session by creating a session RG. Use the following command to enable session charging by allocating a session RG ID.

```
configure mobile-gateway profile charging bng-charging chf session-charging rating-group
```

The session RG uses a single session URR to collect statistics from the FWA-UP, as described in [Statistics collection from the BNG-UP](#). This URR is shared with other applications that require session-level statistics reporting or usage monitoring, for example, PCF failure handling volume limits.

The session RG supports both [online charging](#) and [offline charging](#).



Note: Session RGs do not support RG-level triggers; RG-level triggers are ignored in favor of CHF session-level triggers. See [Charging update triggers](#) for more information about charging triggers.

8.5.5 Custom charging using a static RG

The MAG-c supports charging for a custom subset of traffic by using a static RG. The static RG is linked to a custom statistics group, which the FWA-UP uses to determine the traffic to be counted for this RG. Typically, the custom statistics group specifies a subset of QoS-related counters. Use the following command to configure a static rating group and the associated custom statistics group.

```
configure mobile-gateway profile charging bng-charging chf rating-group
```

Static RGs support only offline charging and do not support triggers set at the RG level. These RG-level triggers are ignored in favor of session-level triggers. The URR used by a static RG is always linked to the URR of the session RG and uses the reporting triggers of the session RG on the FWA-UP.

8.5.6 Offline Charging

When offline charging is enabled for an RG, the MAG-c reports consumed units toward the CHF. Because the MAG-c uses the Nchf_ConvergedCharging SBI service, offline and online charging can be combined. Offline charging is implicitly enabled when online charging is enabled.

Consumed units are always signaled to the CHF when the RG is reported in the Nchf_ConvergedCharging_Update operation, including reports triggered by online charging. Additionally, periodic or volume limit-based reports for offline purposes only can be enabled.

To enable periodic reporting of used statistics to the CHF, the CHF must provide the TIME_LIMIT trigger with a corresponding time limit value for the RG. If this trigger is not set, the MAG-c does not generate periodic offline reports, even if there are periodic usage reports from the FWA-UP to the MAG-c. Upon sending a periodic report, the MAG-c fetches the most recent statistics from the FWA-UP, if the existing statistics are not current. When enabled, these reports are always generated on the specified period, independent of any other reports generated in the reporting period.

To enable volume limit-based reporting of used statistics, the CHF must provide the VOLUME_LIMIT trigger with a corresponding volume limit value for the session RG. This volume limit is always relative to

the last report generated for this RG. That is, when the MAG-c generates any other report, for example periodic, the current volume limit is restarted.

Example: Combination of time and volume limit

The following table shows a time sequence for an RG configured with a combination of a time and volume limit.

Table 15: Time sequence for an RG with a time and volume limit

| Time | Action |
|------------|--|
| 0 h 00 min | The RG is created with a time limit of 1 h and volume limit of 1 GB. |
| 0 h 45 min | The FWA-UP notifies the MAG-c that 1 GB of data has been consumed. The MAG-c reports this to the CHF based on the volume limit. The total consumed data at this point is 1 GB. |
| 1 h 00 min | The periodic time limit expires and the MAG-c requests the consumed data from the FWA-UP. The FWA-UP replies that an additional 300 MB of data has been consumed. The MAG-c reports this to the CHF. The total consumed data at this point is 1.3 GB. |
| 1 h 33 min | The FWA-UP notifies the MAG-c that 1 GB of additional data has been consumed. The MAG-c reports this to the CHF based on the volume limit. The total consumed data at this point is 2.3 GB. |
| 2 h 00 min | The periodic time limit expires and the MAG-c requests the consumed data from the FWA-UP. The FWA-UP replies that an additional 700 MB of data has been consumed. The MAG-c reports this to the CHF. The total consumed data at this point is 3 GB. |
| 3 h 00 min | The periodic time limit expires and the MAG-c requests the consumed data from the FWA-UP. The FWA-UP replies that an additional 150 MB of data has been consumed. The MAG-c reports this to the CHF. The total consumed data at this point is 3.15 GB. |

8.5.7 Online Charging

When online charging is enabled for an RG, the MAG-c requests initial quota units from the CHF. The CHF can reply in two ways:

- indicate that quota management is not applicable
- send initial quota units

When quota management is applicable, the CHF sends a combination of the following quota or thresholds:

- **time or volume quota, or both**
- **optional time or volume threshold or both**
This threshold is a relative threshold for remaining quota. For example, if the CHF signals a volume quota of 10 GB and a volume threshold of 1 GB, the absolute threshold is hit after 9 GB.
- **optional final unit indication**
This indication can have the following values:

- TERMINATE – stop forwarding for the RG
- REDIRECT – redirect traffic when the quota are exhausted

The MAG-c monitors the consumed units of the RG based on usage reports from the FWA-UP and takes the following actions:

- When a threshold is hit, the MAG-c initiates the Nchf_ConvergedCharging_Update operation with the QUOTA_THRESHOLD trigger set. The MAG-c reports the consumed units and requests new units for this RG.
- When a quota is hit and a final unit indication is provisioned, the MAG-c initiates the Nchf_ConvergedCharging_Update operation with the FINAL trigger set. The MAG-c reports any consumed units but does not request new units for this RG. Subsequently, the forwarding rules are updated based on the installed final unit indication:
 - **final unit indication is set to TERMINATE**
The MAG-c instructs the FWA-UP to block all traffic for this RG.
 - **final unit indication is set to REDIRECT**
The MAG-c applies redirect actions specific to that RG. If no RG-specific redirect actions are defined, the MAG-c falls back to the forwarding rules for TERMINATE.
For more information about session RGs, see [Session charging using a session RG](#).
 - **any other final unit indication (for example, RESTRICT)**
The MAG-c falls back to the forwarding rules for TERMINATE.
- When a quota is hit and no final unit indication is provisioned, the MAG-c initiates the Nchf_ConvergedCharging_Update operation with the QUOTA_EXHAUSTED trigger set. The MAG-c reports the consumed units and requests new units for this RG. The MAG-c does not change any traffic-forwarding behavior in this case.

The MAG-c performs the following to expedite quota or threshold monitoring:

- The MAG-c enables quota and threshold monitoring on the FWA-UP, using the URR tied to the RG. This ensures that the FWA-UP generates a usage report soon after reaching the installed quota or threshold, upon which the MAG-c takes any action as described in the preceding content.
- If the final unit indication action equals TERMINATE, the MAG-c instructs the FWA-UP to immediately stop forwarding traffic when this quota is reached.

For more information, see [Statistics collection from the BNG-UP](#)

When quota management is not applicable, the CHF sends the QUOTA_MANAGEMENT_NOT_APPLICABLE result code. The MAG-c takes the following actions:

- Signal the CREDIT_CONTROL_NOT_APPLICABLE error in the creditManagementStatus field toward the PCF.
- Apply offline-only charging to the RG. Contrary to most error result codes, this does not block traffic for the RG.

In case of converged charging, the Charging Data Update messages triggered by offline charging (for example, TIME_LIMIT or VOLUME_LIMIT triggers) do not request new quota units for online charging.

8.5.8 Charging update triggers

The MAG-c supports the CHF to install trigger conditions upon which the Nchf_ConvergedCharging_Update operation is called to request an update. A trigger is categorized as either immediate or deferred, and a default category for each trigger type is defined in 3GPP TS 32.255, table 5.2.1.4.1. As indicated in this table, the CHF can use the TriggerCategory IE to override the category for some triggers. When a deferred trigger condition is fulfilled, the MAG-c does not immediately perform the update operation, but stores the used units for each RG and restarts any ongoing measurements. The used units are included with any trigger that results in an immediate update. The MAG-c can exceptionally generate an update operation for a deferred trigger if the number of pending deferred reports becomes too large.

The following preinstalled trigger types are supported on the MAG-c :

- **USER_LOCATION_CHANGE**

This trigger type is triggered by any change to the session ULI value.

- **TIME_LIMIT**

This trigger type must be signaled with a corresponding **timeLimit** value. This triggers a periodic reporting of used units of RGs for offline charging purposes. If an RG is also subject to online charging, a new quota is not requested and the existing quota stays valid.

- **VOLUME_LIMIT**

This trigger type must be signaled with a corresponding **volumeLimit** value. This triggers reporting of consumed units for the specified RG for offline charging purposes. If the RG is also subject to online charging, a new quota is not requested, and the existing quota stays valid.

- **QUOTA_THRESHOLD**

This triggers reporting of consumed units and requesting new quota when the threshold is consumed. For more information, see [Online Charging](#).

When a quota and quota threshold value are specified, this trigger type is automatically enabled for an RG and cannot be disabled.

- **QUOTA_EXHAUSTED**

This trigger type must be signaled with a corresponding quota value. It triggers reporting of consumed units and requesting new quota when the quota is consumed and no final unit action is specified. Additionally, the MAG-c enforces any installed out-of-quota actions. For more information, see [Online Charging](#).

When a quota value is specified, this trigger type is automatically enabled for an RG and cannot be disabled.

- **FINAL**

This trigger type is always enabled. It is used for the final report of the charging session in the charging release operation. When a quota value is specified, it also triggers reporting of consumed units when the quota is consumed and a final unit action is specified. Additionally, the MAG-c enforces any installed out-of-quota actions. For more information, see [Online Charging](#).

- **VALIDITY_TIME**

This trigger type is automatically enabled when the CHF signals a **validityTime** value; a quota value must be specified. It triggers reporting of consumed units and requesting new quota when the validity time expires. This allows a CHF to periodically assign new quota even if previous quota were not consumed, or if quota were consumed and no new quota were granted.

- **MANAGEMENT_INTERVENTION**

This trigger type cannot be disabled and is always categorized as immediate. It is used whenever the MAG-c has to force an update, for example, when using the following command.

```
clear mobile-gateway pdn charging
```

- **RAT_CHANGE**

This trigger type is triggered by any change to the session RAT type value.

- **FORCED_REAUTHORISATION**

This trigger type cannot be disabled and is always categorized as immediate. It is used whenever the CHF instructs the MAG-c to force an update.

Triggers can be set on the session level, the RG level, or both. At both levels, the triggers act as separate sets and do not hierarchically override each other. In the following example, both trigger levels are combined:

- Configuration:

- Set a TIME_LIMIT trigger with a 10 minutes deferred reporting on the RG level.
- Set a TIME_LIMIT trigger with 1 hour immediate reporting at the session level.

- Result:

- The MAG-c generates a usage report for the RG every 10 minutes without sending it.
- After 1 hour, the MAG-c sends an immediate report with both the current and deferred statistics.




WARNING: The MAG-c ignores unsupported triggers. To prevent a change in behavior when these triggers are supported in a future release, Nokia recommends not enabling any unsupported triggers.

The CHF can explicitly trigger an update via the Nchf_ConvergedCharging_Notify operation. In this operation, the CHF indicates one of the following actions:

- terminate the session, which includes performing the Nchf_ConvergedCharging_Release operation
- request an update using the Nchf_ConvergedCharging_Update procedure, using the FORCED_REAUTHORISATION trigger.

The following figure illustrates the call flow for a reauthorization of all RGs.

Figure 29: CHF-triggered update



sw1470

8.5.9 CHF failure handling

The CHF supports failure handling as defined in [Failure handling](#). During failure handling using the **ap-continue** option, the MAG-c performs the following actions to handle ongoing charging:

- The MAG-c stops any ongoing online charging. Quotas are no longer enforced and thresholds are no longer monitored.
- The MAG-c periodically tries to re-establish a charging session toward the CHF, as configured by the **session-restore-retry-timer** command. If this fails, the failure-handling behavior continues. If this succeeds, the failure-handling behavior stops and the CHF resumes normal charging. For offline charging, any collected statistics up to this point are cleared. For online charging, new quotas and thresholds are monitored starting from this point onward, and any units consumed before or during failure handling are not accounted for.
- The MAG-c can enable a grace period after which the session is removed if a CHF connection cannot be reestablished by the time the period expires. Use the **fh-session-continue-timer** command to enable this functionality.
- If PCF functionality is enabled and the PCF has enabled the CM_SES_FAIL trigger, the MAG-c triggers the Npcf_SMPolicyControl Update operation for any CHF HTTP 4xx error with a cause code specified in the following table. The MAG-c maps the CHF cause code to a value included in the PCF creditManagementStatus code IE. No PCF report is generated for any other cause codes or for 3xx or 5xx HTTP errors.

Table 16: CHF cause code to PCF creditManagementStatus code mappings

| CHF cause code | PCF creditManagementStatus code |
|---------------------------|---------------------------------|
| CHARGING_FAILED | RATING_FAILED |
| CHARGING_NOT_APPLICABLE | CREDIT_CONTROL_NOT_APPLICABLE |
| USER_UNKNOWN | USER_UNKNOWN |
| END_USER_REQUEST_DENIED | END_USER_SERVICE_DENIED |
| END_USER_REQUEST_REJECTED | END_USER_SERVICE_DENIED |

Configure the commands for the preceding actions in the following context.

```
configure mobile-gateway profile charging ccfh-profile
```

When the CHF session is successful, but the CHF returns an RG-specific result code that does not equal SUCCESS to indicate an RG-level error, the MAG-c performs the following actions:

- For any result code except QUOTA_MANAGEMENT_NOT_APPLICABLE, traffic for that RG is blocked, as if [online charging](#) is enabled with a TERMINATE action and quota are exhausted.
- For the QUOTA_MANAGEMENT_NOT_APPLICABLE result code, traffic is still forwarded but only offline charging is applied, even if the local configuration or PCF indicated that online charging should be performed.
- If PCF functionality is enabled and the PCF has enabled the CM_SES_FAIL trigger, the MAG-c triggers the Npcf_SMPolicyControl Update operation for the result codes specified in the following table. The

MAG-c maps the RG result code to a value included in the PCF creditManagementStatus code IE. No PCF report is generated for any other result code.

Table 17: CHF RG result code to PCF creditManagementStatus code mappings

| CHF result code | PCF creditManagementStatus code |
|---------------------------------|---------------------------------|
| RATING_FAILED | RATING_FAILED |
| QUOTA_MANAGEMENT_NOT_APPLICABLE | CREDIT_CONTROL_NOT_APPLICABLE |
| USER_UNKNOWN | USER_UNKNOWN |
| END_USER_SERVICE_DENIED | END_USER_SERVICE_DENIED |
| END_USER_SERVICE_REJECTED | END_USER_SERVICE_DENIED |

Fallback to offline CHF charging

The MAG-c enhances the Session Management Function (SMF) resilience using a fallback mechanism to an offline CHF server, when the converged CHF providing the online charging is unavailable or returns specific error responses. This fallback functionality ensures service continuity and accurate charging-record generation.

The MAG-c can be configured to fall back to the offline CHF server based on configured failure-handling conditions, such as HTTP error codes or transport failures (for example, no response or TCP connection failure). See [Failure handling](#) for information about configuring HTTP error codes and failure-handling actions, including retry behavior.

Converged CHF servers are discovered through the NRF or configured locally. Users can configure a list of offline CHF servers on the MAG-c. The list can include multiple servers for offline fallback.

If a converged CHF server fails, the TCP connection is down, or an HTTP error code is received, the MAG-c behavior depends on the configured failure-handling action:

- **retry-and-fallback-continue**
The MAG-c first retries alternative configured converged CHF peers before falling back to the offline CHF server.
- **fallback-continue**
The MAG-c retries the configured fallback CHF peers specified in the associated **nf-id-list** and does not retry alternative converged CHF peers.

Failure handling with fallback to an offline CHF server applies to all the call-flow stages, including Create, Update, and Release. For unacknowledged messages from the converged CHF server during fallback to the offline CHF server, the MAG-c sets the QuotaManagementIndicator (QMI) parameter to CONVERTED_OFFLINE. The QMI value CONVERTED_OFFLINE indicates the transition from a converged CHF server to an offline CHF server.

Fallback to offline CHF configuration

In the following example, **chf-profile 1** is configured for a converged CHF server and **chf-profile 2** is configured for fallback to an offline CHF server. The **ccfh-profile 1** and **ccfh-profile 2** configurations define failure-handling behavior for the converged and offline CHF servers, respectively.

The **fh-volume-limit** and **fh-session-continue-timer** commands configured under **ccfh-profile 2** apply when the session enters the failure-handling state and transitions to the offline CHF server. This occurs

when charging requests to the converged CHF server fail because of no response, TCP connection failure, or configured HTTP error handling that results in fallback to the offline CHF server.

If both the converged CHF peers and the fallback CHF peers do not respond, the TCP connection is down, or an HTTP error code is received, and the failure-handling action is configured as **ap-continue** or **retry-and-ap-continue**, the session enters failure-handling state. In this case, the MAG-c uses the **fh-session-continue-timer** and **fh-volume-limit** configured under **ccfh-profile 2** to control the session duration and volume while in failure-handling.

The session terminates when the configured volume limit is reached or when the session-continue timer expires. While the session is in the fallback state, if the MAG-c later receives a message from the converged CHF server, such as a reauthorization message, the MAG-c terminates the session.

Example: Configuring fallback to offline CHF

```
charging
  ccfh-profile 1
    cc-failure-handling retry-and-ap-continue
    fh-volume-limit 10000
    fh-session-continue-timer 10
    session-restore-retry-timer 5
    nchf-http-code 410 fallback-continue
    nchf-http-code 500 retry-and-fallback-continue
    nchf-http-code 503 retry-and-fallback-continue
    nchf-http-code 504 retry-and-fallback-continue
  exit

  ccfh-profile 2
    cc-failure-handling retry-and-ap-continue
    fallback-volume-limit 50000000
    fallback-session-continue-timer 30
    fh-volume-limit 10000
    fh-session-continue-timer 5
  exit

  chf-profile 1
    description "Converged Charging Profile"
    ccfh-profile 1
    sbi-srv-operation
      transaction-timer 1500
    exit
  exit

  chf-profile 2
    description "Fallback CHF Charging Profile"
    ccfh-profile 2
    sbi-srv-operation
      transaction-timer 1500
    exit
  exit

  bng-charging "bng-charging-prof-1"
    chf
      chf-profile 1
      chf-selection
        fh-fallback
          chf-profile 2
          nf-id-list "nf-id-list-chf-fallback"
      default-charging-method online
      dnn-format real-ni-only
      subscriber-identifier-format identitytype
      session-charging
```

```
        rating-group 74
        exit
        no shutdown
    exit
exit
```

If failure handling is configured with the **retry-and-fallback-continue** option for the HTTP status codes received from the converged CHF server, the MAG-c first attempts alternative CHF peers. If the retry attempts fail, the MAG-c transitions to the offline CHF server.

If failure handling is configured with the **fallback-continue** option for the HTTP status codes, the MAG-c skips online retry attempts and immediately transitions to the offline CHF servers.

If failure handling is configured for a specific HTTP status code using the **nchf-http-code** command, that configuration takes precedence over the **cc-failure-handling** configuration. If the HTTP status code is not explicitly configured, the MAG-c refers to the **cc-failure-handling** configuration.

9 Residential NAT

Get information about NAT on BNG CUPS, learn about the functional split between MAG-c and BNG-UP, about configuring NAT, logging, and monitoring NAT resources.

9.1 NAT terminology and references

| | |
|-----------------------------------|--|
| private side or inside NAT | 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. |
| public side or outside NAT | 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. |
| NAT pool | 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. |
| NAT flow | <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"> • source IP address • source port • translated IP address • translated port • destination port • destination IP address • protocol <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> |
| NAT and NAT44 | 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.

9.2 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 following figure shows a residential NAT example with two residences.

Figure 30: 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.

9.3 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*.

9.4 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 preconfigured prefix (supernet). ODSA supports exclusion of IP address ranges from the configured prefix. In NAT this means that an IP address range that is excluded in the configuration is not used in the NAT pools.

An excluded IP address range must align with the ODSA micro-net boundary, however, the CLI does not enforce configuration of an excluded IP address range that covers the ODSA micro-net. If you configure a range that only partially covers the ODSA micro-net, the entire ODSA micro-net is excluded from the NAT pool. The result is more IP addresses than are configured for exclusion are taken from NAT pool.

The following example shows a misaligned IP address range that is excluded.

Example: Misaligned configuration for excluded IP address range

```
A:MAG-c>config>mobile>pdn# info
-----
      local-address-assignment
        network-realm vrf-1
          pool pool-1
            ipv4
              prefix 198.51.100.0/24
                exclude-addresses 198.51.100.100 192.51.100.120
                micro-net-length 28
            -----
```

In this example, ODSA creates micro-nets, starting from 198.51.100.0/28, 198.51.100.16/28, and so on, to 198.51.100.240/28. The configured excluded range partially covers two such micro-nets: 198.51.100.96/28 and 198.51.100.112/28. In this case, both of these micro-nets are fully excluded from the NAT pool. The micro-net 198.51.100.96/28 has four IP addresses (.96, .97, .98, and .99) that should not be excluded. Similarly, 198.51.100.112/28 has eight IP addresses (.121 to .128) that should not be excluded. This means that twelve additional IP addresses are excluded from the NAT pool that should not be excluded.

The IP address exclusion feature is not supported for a prefix that is configured as dedicated (**configure mobile-gateway pdn local-address-assignment network-realm pool dedicated**).

See [ODSA and local address assignment](#) for more information about NAT prefix allocation.

9.5 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.

Use the **cp-nat-profile** command in the following context to configure a CP NAT profile.

```
configure mobile-gateway profile bng
```

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 subscribers 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.

9.6 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 RADIUS Alc-Static-Port-Forward VSA is used for allocation or deletion of static port forwards. For more information about the VSA, see the *MAG-c RADIUS Attributes and IU Triggers*.

9.7 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.

9.7.1 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 mode
```

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

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

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".

9.7.2 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
      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
      exit
-----
```

Example: BNG-UP NAT policy configuration

```
A:BNG-CP>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.

9.8 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 *MAG-c RADIUS Attributes and IU Triggers*.

The following table describes the relation between the accounting message type, the NAT events, and the content of the Alc-Nat-Port-Range VSA.

Table 18: 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 deallocated 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 |

9.8.1 RADIUS-based logging

Residential NAT on MAG-c uses RADIUS-based logging integrated with subscriber accounting. The logging uses the same RADIUS infrastructure for NAT as for subscriber management.

NAT logging on MAG-c uses the following relevant VSAs:

- 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 attribute with a one second resolution is updated by the MAG-c to reflect the time when the Interim-Update message is generated on the MAG-c.
- The RADIUS Alc-ISA-Event-Timestamp VSA 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 the integrated subscriber management and RADIUS logging attributes relevant to NAT.

Table 19: Subscriber management and NAT integrated RADIUS accounting and logging

| Accounting message type | Accounting actions (session based) | Comments |
|-------------------------|--|---|
| START | <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"> • the outside IP address and initial port for the subscriber's first NAT-enabled session • the outside IP address, the initial PB, and the extended PBs for any additional NAT-enabled sessions of the subscriber. <p>The NAT-related information is carried in the following RADIUS attribute: Alc-Nat-Port-Range (26.6527.121)</p> <p>This attribute includes the outside IP address, PBs, outside realm, and NAT pool.</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. |
| Regular Interim-Update | <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: 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> | — |

| Accounting message type | Accounting actions (session based) | Comments |
|--------------------------|--|----------|
| | <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>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> <ul style="list-style-type: none"> • NAT-MAP (20) • NAT-FREE (19) <p>This attribute includes the reason for this triggered Interim-Update message, which is an extended PB is allocated (MAP) or deallocated (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> | — |

| Accounting message type | Accounting actions (session based) | Comments |
|-------------------------|---|---|
| STOP | <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 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>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 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 NAT logging

The following example describes relevant NAT-related attributes for NAT logging.

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
```

3. 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
```

4. 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
```

5. 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
```

9.8.1.1 Enabling RADIUS logging on MAG-c

Procedure

Use the **nat-port-range** and the **acct-triggered-reason** commands in the following context to enable session accounting with NAT-related information.

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

The **nat-port-range** command enables sending the RADIUS Alc-Nat-Port-Range and Alc-ISA-Event-Timestamp VSAs for the session accounting.

The **acct-triggered-reason** command together with the triggered Interim-Update message conveys information about the event itself.

9.8.1.2 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 31: Alc-ISA-Event-Timestamp triggered Interim-Update message



sw1366

9.8.1.3 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.

9.8.1.4 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.

9.9 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.

9.10 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.

Configure a local address assignment pool with an outside NAT prefix.

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

Example

```
A:MAG-c>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. Configure a CP NAT profile.

```
configure mobile-gateway profile bng cp-nat-profile
```

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

```
A:MAG-c>configure# mobile-gateway profile bng
-----
cp-nat-profile "profile-1"
  nat-pool "pool-1"
    laa-pool network-realm "realm-1" name "laa-pool-1"
  up-nat-policy "up-pol-1"
  mode napt
    port-reservation ports 2000
  exit
```

- 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.

Reference the CP NAT profile in the ADB.

```
configure mobile-gateway profile authentication-database entry cp-nat-profile
```

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. Explicitly enable the Alc-Nat-Port-Range VSA.

```
configure mobile-gateway profile charging bng-charging radius session include-attribute nat-port-range
```

Example

```
A:MAG-c>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*.

9.11 Monitoring NAT resources

Monitor NAT resources using **show** CLI commands.



Note: This section describes NAT resource monitoring on the MAG-c. For information about monitoring NAT resources on the BNG-UP, see the *7450 ESS*, *7750 SR*, and *VSR Multiservice ISA and ESA Guide*.

9.11.1 Terminology overview

The following terminology is used in the context of monitoring NAT pool resources. [Configuration of NAT pools in the MAG-c](#) shows a sample NAT pool configuration.

- **ODSA**

The on-demand subnet allocation (ODSA) subsystem is used for IP address assignment. Use commands in the following context to configure this address assignment.

```
configure mobile-gateway pdn local-address-assignment
```

- **ODSA pool**

The ODSA pool (pool - 1 in the example configuration) is used for the following applications:

- subscriber session address assignment – addresses for subscriber sessions
- NAT pool addressing – public IP addresses in NAT

An ODSA pool can be shared between the two applications.

- **ODSA pool prefix**

The ODSA pool prefix (198.51.100.0/24 in the example configuration) is used for the allocation of micro-nets.

- **ODSA pool micro-net**

An ODSA pool micro-net (198.51.100.224/28 in the example configuration) is an address range in the ODSA pool that is assigned to the BNG-UP.

- **NAT member micro-net**

A NAT member micro-net is an address range in the ODSA pool micro-net. It is allocated to each NAT member on the BNG-UP.

Typically, a NAT member corresponds to an Extended Services Appliance-virtual machine (ESA-VM). In an active/active ESA-VM inter-chassis redundancy configuration in the BNG-UP, however, each ESA-VM is segmented into multiple logical ESA-VMs. In this case, a single ESA-VM houses multiple NAT members.

- **NAT subscriber**

A NAT subscriber defines sessions that share the same outside IP address and port blocks (PBs).

- **initial PB**


The initial PB is the first PB assigned to each NAT subscriber.

- **shared PB space**

The shared PB space on an outside IP address is allocated for PB expansion per NAT subscriber. Each NAT subscriber may get multiple PB allocations from the shared space.

Example: Configuration of NAT pools in the MAG-c

Figure 32: NAT pools in the MAG-c



In the preceding figure, the MAG-c is configured as follows.

```
A:MAG-c>configure# mobile-gateway pdn
-----
local-address-assignment
  network-realm "vrf A"
  pool "pool-1"
    ipv4
      prefix 198.51.100.0/24
      micro-net-length 28
```

9.11.2 NAT pool resources in ODSA

Use the following command to display a high-level overview of the ODSA pool.

```
show mobile-gateway pdn local-address-assignment pool
```

The IPv4 address usage and micro-net usage in the command output show the percentage of pool resources used by both the subscriber session address assignment and the NAT pool addressing applications. For precise resource tracking within the NAT context, Nokia recommends assigning an ODSA pool exclusively to the NAT application.

While each prefix in the pool can be used by both applications, each micro-net can be used by only one application. Use the following command to display the aggregated usage of the prefixes in the ODSA pool (with no differentiation between the applications).

```
show mobile-gateway pdn local-address-assignment pool prefixes
```

Use the following command for high-level information about each prefix with differentiation between applications.

```
show mobile-gateway pdn local-address-assignment pool prefix
```

When assigned to subscriber sessions, the IP address enters the **hold** state. Therefore, IP addresses in the **hold** state are irrelevant to NAT.

An IP address in the **hold** state that is deallocated from the subscriber cannot be immediately reused because of a default timer of one minute.

Use the following command to display an overview of all micro-nets for a specific prefix, their corresponding associations with the BNG-UPs, and the IP address usage within each micro-net.

```
show mobile-gateway pdn local-address-assignment pool prefix micronets
```

Example: High-level overview of the ODSA pool

```
A:MAG-c>show>mobile-gateway>pdn>laa# pool gateway 1 network-realm "601" name "601-pool-0"
=====
Local Address Assignment Pool
=====
PDN gateway           : 1
Network realm        : 601
Pool                  : 601-pool-0
-----
Description           : (Not Specified)
Tracking              : none
Dedicated             : no
Hold time             : 00h00m00s
Minimum number of free micro-nets : 1
  Rising threshold    : 2
IPv4
  Default gateway     : first-address
  Subnet allocation    : (Not Specified)
  Primary DNS         : (Not Specified)
  Secondary DNS       : (Not Specified)
  Micro-net length    : 24
  Number of prefixes  : 1
  Number of excludes  : 0
  Number of micro-nets : 8
  IP address usage    : 11%
  Micro-net usage     : 13%
IPv6
  Primary DNS         : (Not Specified)
  Secondary DNS       : (Not Specified)
  NA
    Micro-net length  : (Not Specified)
    Number of prefixes : 0
    Number of excludes : 0
    Number of micro-nets : 0
    IP address usage  : 0%
    Micro-net usage   : 0%
  PD
    Delegated prefix  : length 56
    Micro-net length  : (Not Specified)
    Number of prefixes : 0
    Number of excludes : 0
    Number of micro-nets : 0
    IP address usage  : 0%
```

```

Micro-net usage      : 0%
SLAAC
Micro-net length    : (Not Specified)
Number of prefixes  : 0
Number of excludes  : 0
Number of micro-nets : 0
IP address usage    : 0%
Micro-net usage     : 0%

```

Example: Aggregated usage of the prefixes in the ODSA pool

The following example output shows 11% IP address usage within the prefix.

```

A:MAG-c>show>mobile-gateway>pdn>laa pool gateway 1 network-realm "601" name "601-pool-0"
prefixes

=====
Local Address Assignment Pool Prefixes
=====
Prefix                                     Address  Oper  IP Addr Drain
Family                                     State   Usage
-----
130.0.0.0/21                             ipv4     up    11%    no
-----
No. of prefixes: 1

```

Example: High-level information for each prefix with differentiation between applications

```

A:MAG-c>show>mobile-gateway>pdn>laa# pool gateway 1 network-realm "601" name "601-pool-0"
prefix 130.0.0.0/21

=====
Local Address Assignment Prefix
=====
PDN gateway                               : 1
Network realm                             : 601
Pool                                       : 601-pool-0
Prefix                                    : 130.0.0.0/21
-----
Address family                             : ipv4
Micro-net length                           : 24
Operational state                          : up
Drain                                       : no
Micro-nets                                 : 8
Micro-nets in use (%)                      : 1 (13%)
NAT
  Micro-nets in use (%)                    : 1 (13%)
  IP addresses                             : 256
  IP addresses in use (%)                   : 218 (85%)
Address assignment
  Micro-nets in use (%)                    : 0 (0%)
  IP addresses                             : 1792
  IP addresses reserved                    : 21
  IP addresses in use (%)                  : 0 (0%)
  IP addresses active (%)                  : 0 (0%)
  IP addresses holdoff (%)                 : 0 (0%)
  IP addresses excluded (%)                : 0 (0%)

```

Example: Overview of all micro-nets for a specific prefix

```

A:MAG-c>show>mobile-gateway>pdn>laa pool gateway 1 network-realm "601" name "601-pool-0"
prefix 130.0.0.0/21 micro-nets

```

```

=====
Local Address Assignment Micro-nets for Prefix 130.0.0.0/21
=====
Micro-net                UPF/FSG                IP Addresses In Use
                        Active      Hold Usage
-----
130.0.0.0/24            NAT 10.0.1.3          218          85%
130.0.1.0/24            N/A                  0            0    0%
130.0.2.0/24            N/A                  0            0    0%
130.0.3.0/24            N/A                  0            0    0%
130.0.4.0/24            N/A                  0            0    0%
130.0.5.0/24            N/A                  0            0    0%
130.0.6.0/24            N/A                  0            0    0%
130.0.7.0/24            N/A                  0            0    0%
-----
No. of micro-nets: 8

```

9.11.3 Number of subscribers and extended PBs

Use the following command to display the number of allocated micro-nets, the number of users within each micro-net, and the extended PB usage for each micro-net under a specific prefix.

```
show mobile-gateway pdn local-address-assignment pool prefix micro-nets nat
```

Use the following command to display the overall subscriber usage per micro-net and the association between the CP NAT profile and the micro-net.

```
show mobile-gateway pdn local-address-assignment pool prefix micro-net
```

Use the following command to display overall subscriber usage and extended PB usage per NAT member micro-net on the BNG-UP.

```
show mobile-gateway pdn local-address-assignment pool prefix micro-net nat-members
```

Use the following command to display the extended PB usage per IP address under a specific prefix.

```
show mobile-gateway pdn local-address-assignment pool prefix micro-nets nat extended-port-blocks
```

Example: Number of subscribers and extended PBs

```

A:MAG-c>show>mobile-gateway>pdn>laa# pool gateway 1 network-realm "601" name "601-pool-0"
prefix 130.0.0.0/21 micro-nets nat
=====
Local Address Assignment NAT Micro-nets for Prefix 130.0.0.0/21
=====
Micro-net                UPF/FSG                Subscribers (%) Ext Port Blocks (%)
-----
130.0.0.0/24            NAT 10.0.1.3          138853 (84%)   483479 (33%)
-----
No. of subscribers :      138853
No. of NAT micronets:      1

```

Example: Subscriber usage for a specific micro-net

```
A:MAG-c>show>mobile-gateway>pdn>laa# pool gateway 1 network-realm "601" name "601-pool-0"
prefix 130.0.0.0/21 micro-net 130.0.0.0/24

=====
Local Address Assignment Micro-net
=====
PDN gateway                : 1
Network realm              : 601
Pool                       : 601-pool-0
Prefix                     : 130.0.0.0/21
Address family             : ipv4
Micro-net                  : 130.0.0.0/24
-----
UPF/FSG                    : NAT 10.0.1.3
NAT
  CP NAT profile           : Dut-C_pol_poll1_for_601-pool-0
  NAT pool                 : 601-pool-0
  IP addresses             : 256
  IP addresses in use (%)  : 218 (85%)
  Maximum subscribers     : 165120
  Active subscribers (%)   : 138853 (84%)
  Maximum ext port blocks  : 1486336
  Active ext port blocks (%) : 377434 (25%)
```

Example: Subscriber usage and extended PB usage per NAT member micro-net

```
A:MAG-c>show>mobile-gateway>pdn>laa# pool gateway 1 network-realm "601" name "601-pool-0"
prefix 130.0.0.0/21 micro-net 130.0.0.0/24 nat-members

=====
Local Address Assignment Micro-net nat-members
=====
PDN gateway                : 1
Network realm              : 601
Pool                       : 601-pool-0
Prefix                     : 130.0.0.0/21
Address family             : ipv4
Micro-net                  : 130.0.0.0/24
UPF/FSG                    : NAT 10.0.1.3
-----
Member Address Range      IP Addresses  Subscribers  Ext Port Blocks
-----
1 130.0.0.171 - 130.0.0.255    85           46284        86504
2 130.0.0.0 - 130.0.0.85       86           46284        86808
3 130.0.0.86 - 130.0.0.170    85           46285        80243
-----
```

Example: Extended PB usage per IP address under a specific prefix

```
A:MAG-c>show>mobile-gateway>pdn>laa# pool gateway 1 network-realm "601" name "601-pool-0"
prefix 130.0.0.0/21 micro-nets nat extended-port-blocks

=====
Local Address Assignment NAT Micro-nets for Prefix 130.0.0.0/21
=====
Micro-net      UPF/FSG
  IP Address                Subscribers (%) Ext Port Blocks (%)
-----
130.0.0.0/24    NAT 10.0.1.3
```

| | | |
|------------|------------|----------|
| 130.0.0.0 | 645 (100%) | 207 (4%) |
| 130.0.0.1 | 645 (100%) | 220 (4%) |
| 130.0.0.2 | 645 (100%) | 93 (2%) |
| 130.0.0.3 | 645 (100%) | 63 (1%) |
| 130.0.0.4 | 645 (100%) | 0 (0%) |
| 130.0.0.5 | 645 (100%) | 0 (0%) |
| 130.0.0.6 | 645 (100%) | 0 (0%) |
| 130.0.0.7 | 645 (100%) | 0 (0%) |
| 130.0.0.8 | 645 (100%) | 0 (0%) |
| 130.0.0.9 | 645 (100%) | 0 (0%) |
| 130.0.0.10 | 645 (100%) | 0 (0%) |
| 130.0.0.11 | 645 (100%) | 0 (0%) |
| 130.0.0.12 | 645 (100%) | 0 (0%) |
| 130.0.0.13 | 645 (100%) | 0 (0%) |
| 130.0.0.14 | 645 (100%) | 0 (0%) |
| 130.0.0.15 | 645 (100%) | 0 (0%) |
| 130.0.0.16 | 645 (100%) | 0 (0%) |
| 130.0.0.17 | 645 (100%) | 0 (0%) |
| 130.0.0.18 | 645 (100%) | 0 (0%) |
| 130.0.0.19 | 645 (100%) | 0 (0%) |
| 130.0.0.20 | 645 (100%) | 0 (0%) |
| 130.0.0.21 | 645 (100%) | 0 (0%) |
| 130.0.0.22 | 645 (100%) | 0 (0%) |
| 130.0.0.23 | 645 (100%) | 0 (0%) |
| 130.0.0.24 | 645 (100%) | 0 (0%) |
| 130.0.0.25 | 645 (100%) | 0 (0%) |
| 130.0.0.26 | 645 (100%) | 0 (0%) |
| 130.0.0.27 | 645 (100%) | 0 (0%) |
| 130.0.0.28 | 645 (100%) | 0 (0%) |
| 130.0.0.29 | 645 (100%) | 0 (0%) |
| 130.0.0.30 | 645 (100%) | 0 (0%) |
| 130.0.0.31 | 645 (100%) | 0 (0%) |
| 130.0.0.32 | 645 (100%) | 18 (<1%) |
| 130.0.0.33 | 645 (100%) | 18 (<1%) |
| 130.0.0.34 | 645 (100%) | 251 (4%) |
| 130.0.0.35 | 645 (100%) | 252 (4%) |
| 130.0.0.36 | 645 (100%) | 12 (<1%) |
| 130.0.0.37 | 645 (100%) | 4 (<1%) |
| 130.0.0.38 | 645 (100%) | 0 (0%) |
| 130.0.0.39 | 645 (100%) | 0 (0%) |
| 130.0.0.40 | 645 (100%) | 0 (0%) |
| 130.0.0.41 | 645 (100%) | 0 (0%) |
| 130.0.0.42 | 645 (100%) | 0 (0%) |

9.11.4 Subscriber sessions and CP NAT profile

Use the following command to display basic information for a BNG session. The output includes the CP NAT profile which is used as input for other commands.

```
show mobile-gateway bng session
```

Use the following command to display information about the NAT address range and the realm for the subscriber.

```
show mobile-gateway bng session nat
```

Use the following command to display information about a specific CP NAT profile.

```
show mobile-gateway profile bng cp-nat-profile
```

Example: Basic BNG session information to get the CP NAT profile for the session

```
A:MAG-c# show mobile-gateway bng session

=====
BNG Sessions
=====
[IPoE] MAC                : 02:00:00:00:00:01
-----
L2 Access Id              : 1/2/8
S-Vlan                    : 1
C-Vlan                    : 1
MAC                       : 02:00:00:00:00:01

Up Time                   : 0d 00:00:05
Circuit Id                : N/A
Remote Id                 : N/A
Provisioned Addresses     : IPv4
Signaled Addresses       : IPv4

UP Peer                   : 10.0.1.3
Selected APN/DNN          : 10001
Network Realm             : 10001
IPv4 Pool                 : 10001

IPv4 Address              : 172.16.0.2
IPv4 Address Origin       : Local pool
IPv4 Prefix Len           : 28
IPv4 Gateway              : 172.16.0.1
IPv4 Primary DNS          : 0.0.0.0
IPv4 Secondary DNS        : 0.0.0.0
IPv4 Primary NBNS         : 0.0.0.0
IPv4 Secondary NBNS       : 0.0.0.0
DHCPv4 Server IP         : 172.16.0.1
DHCPv4 Lease Time        : 7d 00:00:00
DHCPv4 Renew Time        : 3d 12:00:00
DHCPv4 Rebind Time       : 6d 03:00:00
DHCPv4 Lease End         : 05/29/2023 12:39:39
DHCPv4 Remaining Lease Time : 6d 23:59:55

Subscriber                : auto_sub_1 (1)
Acct-Session-Id           : X000111002832A30B00000001
Acct-Multi-Session-Id    : Y000000012832A30B00000000
State Id                  : 0x646b628b00000000
Sub Profile               : cust_1
Sla Profile               : sla-prof
SAP-Template              : mySapTemplate
Group-itf-template       : myGroupItfTmpl
Number of Framed IPv4 Routes : 0
Number of Framed IPv6 Routes : 0
NAT Profile               : Dut-C_pol_l2a-0-1_1
HTTP Redirect URL         : N/A
Intermediate Destination Id : N/A
Ingress IPv4 filter override : N/A
Egress IPv4 filter override : N/A
Ingress IPv6 filter override : N/A
Egress IPv6 filter override : N/A
Number of QoS Overrides   : 0
```

```

PFCP Local SEID      : 0x0000000000011100
PFCP Remote SEID    : 0x0000000000000002

UE Id                : 0x00011100
PDN Session Id      : 0x00011100
Group/VM             : 2/3
Call-Insight        : disabled

Charging Profile 1   : myChargingProfile
  Charging enabled    : Yes

-----
Number of sessions shown : 1

```

Example: The NAT address range and the realm for the subscriber

```

A:MAG-c# show mobile-gateway bng session nat

=====
BNG Sessions NAT
=====
[IPoE] MAC          : 02:00:00:00:00:01
-----
CP NAT profile      : Dut-C_pol_l2a-0-1_1

NAT pool            : l2a-0-1
UP NAT policy       : l2a-0-1_1
Outside network realm : 123321
Outside IP address   : 130.0.0.25
Outside ports        : 1024-1024
Extended ports       : 1026-1035
                     : 1036-1045
                     : 1046-1055
                     : 1056-1065
                     : 1066-1075
                     : 1076-1085
                     : 1086-1095
                     : 1096-1105
                     : 1106-1115

Number of sessions shown : 1

```

Example: CP NAT profile information

```

A:MAG-c# show mobile-gateway profile bng cp-nat-profile "Dut-C_pol_l2a-0-1_1"

=====
CP NAT profile
=====
Description          : Converted old nat-pool of dut Dut-C into CUPS
                      nat-cp-profile
Access mode           : auto
Last management change : 05/22/2023 12:29:30
-----
Pool                 : l2a-0-1
Description           : nat-policy Dut-C_pol_l2a-0-1_1 with nat-pool
                      l2a-0-1
Mode                  : napt
Port forwarding range : 1 - 1023
Port reservation      : 1 ports
Forwarding port limit : (Not Specified)
Local address assignment pool

```

```
Name : l2a-0-1
Network realm : 123321
UP NAT policy : l2a-0-1_1
Last management change : 05/22/2023 12:29:32
```

9.11.5 Histograms

Histograms provide summarized distribution views for various entities.

Use the following command to display the extended PB allocation in relation to the number of subscribers.

```
show mobile-gateway bng nat histogram port-blocks
```

Use the following command to display the aggregated distribution of extended PBs across IP addresses. Unused IP addresses are not included.

```
show mobile-gateway bng nat histogram extended-port-blocks-per-ip
```

Example: Extended PB allocation in relation to the number of subscribers

In the following example output, 24,241 subscribers have allocated only one extended PB, 229 subscribers have allocated 2 extended PBs, and so on.

```
A:MAG-c# show mobile-gateway bng nat histogram port-blocks cp-nat-profile "Dut-C_pol_poll_
for_601-pool-0" nat-pool "601-pool-0" up 10.0.1.3

=====
Usage histogram UP 10.0.1.3 port blocks per subscriber
=====
Num port blocks          Num subscribers
-----
1                        24241
2                         229
3                         232
4                         230
5                         236
6                         222
7                         200
8                         245
9                         281
10                       112737
-----
No. of entries: 10
```

Example: Aggregated distribution of extended PBs across IP addresses

```
A:MAG-c# show mobile-gateway bng nat histogram extended-port-blocks-per-ip cp-nat-profile
"Dut-C_pol_poll_for_601-pool-0" nat-pool "601-pool-0" up 10.0.1.3 bucket-size 100 num-
buckets 50

=====
Usage histogram UP 10.0.1.3 extended port blocks per IP address
=====
Num extended port blocks  Num IP addresses
-----
1-99                      0
100-199                   0
200-299                   0
```

| | |
|-----------|-----|
| 300-399 | 14 |
| 400-499 | 2 |
| 500-599 | 1 |
| 600-699 | 6 |
| 700-799 | 11 |
| 800-899 | 50 |
| 900-999 | 65 |
| 1000-1099 | 104 |
| 1100-1199 | 171 |
| 1200-1299 | 201 |
| 1300-1399 | 163 |
| 1400-1499 | 60 |
| 1500-1599 | 23 |
| 1600-1699 | 1 |
| 1700-1799 | 0 |
| 1800-1899 | 0 |
| 1900-1999 | 0 |
| 2000-2099 | 0 |
| 2100-2199 | 0 |
| 2200-2299 | 0 |
| 2300-2399 | 0 |
| 2400-2499 | 0 |
| 2500-2599 | 0 |
| 2600-2699 | 0 |
| 2700-2799 | 0 |
| 2800-2899 | 0 |
| 2900-2999 | 0 |
| 3000-3099 | 0 |
| 3100-3199 | 0 |
| 3200-3299 | 0 |
| 3300-3399 | 0 |
| 3400-3499 | 0 |
| 3500-3599 | 0 |
| 3600-3699 | 0 |
| 3700-3799 | 0 |
| 3800-3899 | 0 |
| 3900-3999 | 0 |
| 4000-4099 | 0 |
| 4100-4199 | 0 |
| 4200-4299 | 0 |
| 4300-4399 | 0 |
| 4400-4499 | 0 |
| 4500-4599 | 0 |
| 4600-4699 | 0 |
| 4700-4799 | 0 |
| 4800-4899 | 0 |
| 4900- | 0 |

No. of entries: 50

10 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.

10.1 Geo-redundancy overview



Note: Nokia has introduced the terms active/standby as alternatives to the former master/slave terminology. This document has been updated with these new terms.

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 runtime. 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 33: Geo-redundant MAG-c deployment



10.2 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
```

10.3 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 on the standby system.

```
configure redundancy multi-chassis peer mc-mobile traffic-detection
```

When you specify the **relaxed** option for the **traffic detection** command, the standby system changes to the active role only when it receives a PFCP or a IBCP packet (for BNG).

In **strict** mode, the standby system changes to the active role when it receives a PFCP packet and a GTP-C packet on S11 for 4G FWA or an HTTP/2 packet on N11 for 5G SA FWA.

Use the following command to configure traffic detection behavior on the active system.

```
configure redundancy multi-chassis peer mc-mobile master-traffic-detection
```

When the **master-traffic-detection** command is enabled, the active system performs traffic detection to avoid an active/active scenario when mc-mobile is down.



Note: Nokia recommends enabling the **master-traffic-detection** command and configuring the **traffic-detection** command to the **relaxed** option for BNG deployment and the **strict** option for FWA deployment.

If the number of BNG-UPs is small, you can 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 or FWA-UP:



Note: The BNG-UP and FWA-UP support both the MD-CLI and the classic CLI, while the MAG-c only supports the classic CLI.

- **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
```

10.4 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
```

10.5 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.

10.6 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 (VPRN-over-GRE) service destination point (SDP) for the following types of traffic:

- PFCP
- IBCP
- RADIUS CoA
- GTP-C
- HTTP/2

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.

10.7 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.

10.8 Deploying and configuring geo-redundancy

Deployment guidelines

The following guidelines apply to geo-redundant deployments:

- Configure the PFCP path timers such that the termination of a PFCP path takes longer than the completion of a geo-redundancy switchover, including the detection time and routing convergence. Otherwise, the BNG-UPs or FWA-UPs may terminate a PFCP path and remove all corresponding sessions during a geo-redundancy switchover.

Use the PFCP headless mode to achieve the preceding configuration. See [PFCP connectivity failure](#) for more information about the PFCP path timers.

- When adding a new configuration on a geo-redundant pair, provision the standby system before the active system.
- Use the following command to configure the **relaxed** option for fixed access geo-redundant deployments and the **strict** option for FWA deployments.

```
configure redundancy multi-chassis peer mc-mobile traffic-detection
```

- For fixed access geo-redundant deployments, Nokia recommends increasing the traffic detection timer to 10 seconds using the following command.

```
configure redundancy multi-chassis peer mc-mobile traffic-detection-poll-timer
```

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
```

The following example shows a geo-redundant configuration for an active system, with the relaxed option specified for traffic detection, as recommended for BNG deployments.

Example: Geo-redundant configuration on a primary system (BNG-UP deployment)

```
A:MAG-c>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
-----
A:MAG-c>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

```

The following example shows a geo-redundant configuration for a standby system, with the relaxed option specified for traffic detection, as recommended for BNG deployments.

Example: Geo-redundant configuration on a standby system (BNG-UP deployment)

```

A:BNG-CP>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
-----
A:BNG-CP>config>router>policy-options# 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
```

**Note:**

Configuration on the active and standby systems must remain the same during normal operation, while the sessions are established. Configuration changes that impact existing sessions are not allowed because they impact the state on the active and standby systems. For the cases that require new resources to be added to the system, such as IP pools, you must make the configuration changes on the standby system before applying the same changes to the active system.

Contact your Nokia technical support representative to verify the configuration changes that are allowed during the operational state of active and standby systems.

11 Python support

Customize the MAG-c behavior by applying user-defined Python scripts to inspect and modify control protocol packets.

Python scripts

Sending or receiving specific control protocol packets can trigger a user-defined Python script. Using the packet as input, the script applies a set of Nokia API calls to inspect and modify the packet. The script outputs 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 software includes Nokia-provided APIs and the following standard libraries:

- MicroPython libraries
 - sys
 - uarray
 - ubinascii
 - ucollections
 - uhashlib
 - uio
 - ure
 - ustruct
 - utime



Note: Nokia has modified the implementation of this module. For more information about the use of this module, see the *pySROS API documentation* provided with the [SR OS \(nokia.com\)](#) documentation.

- standard libraries
 - datetime
 - ipaddress

For more information about the Nokia-provided APIs, see *MAG-c TPSDA Python 3 API*.

Supported protocol messages

Python is supported for the following HTTP/2 protocols:

- Request (ingress and egress)
- Response (ingress and egress)

The following tables list the supported protocol message types and direction.

Table 20: Supported direction for RADIUS messages

| Message type | Ingress | Egress |
|------------------|---------|--------|
| Access-Request | | ✓ |
| Access-Accept | ✓ | |
| Access-Reject | ✓ | |
| Account-Request | | ✓ |
| Account-Response | ✓ | |
| Access-Challenge | ✓ | |

Table 21: Supported direction for RADIUS CoA messages

| Message type | Ingress | Egress |
|--------------|---------|--------|
| CoA Request | ✓ | |
| DM Request | ✓ | |
| CoA/DM Reply | | ✓ |

Table 22: Supported direction for PPPoE messages

| Message type | Ingress | Egress |
|--------------|---------|--------|
| PADI | ✓ | |
| PADO | | ✓ |
| PADR | ✓ | |
| PADS | | ✓ |
| PADT | ✓ | ✓ |
| LCP | ✓ | ✓ |
| PAP | ✓ | ✓ |
| CHAP | ✓ | ✓ |
| IPCP | ✓ | ✓ |
| IPv6CP | ✓ | ✓ |

Table 23: Supported direction for DHCPv4 messages

| Message type | Ingress | Egress |
|--------------|-----------|-----------|
| Discover | ✓ | ✓ (relay) |
| Offer | ✓ (relay) | ✓ |
| Request | ✓ | ✓ (relay) |
| Ack | ✓ (relay) | ✓ |
| Decline | ✓ | ✓ (relay) |
| Nak | ✓ (relay) | ✓ |
| Release | ✓ | ✓ (relay) |
| Inform | ✓ | ✓ (relay) |
| Lease-Query | ✓ | ✓ (relay) |



Note: (relay) means supported only for relayed messages.

Table 24: Supported direction for DHCPv6 messages

| Message type | Ingress | Egress |
|---------------|-----------|-----------|
| Solicit | ✓ | ✓ (relay) |
| Advertise | ✓ (relay) | ✓ |
| Request | ✓ | ✓ (relay) |
| Confirm | ✓ | ✓ (relay) |
| Renew | ✓ | ✓ (relay) |
| Rebind | ✓ | ✓ (relay) |
| Reply | ✓ (relay) | ✓ |
| Release | ✓ | ✓ (relay) |
| Decline | ✓ | ✓ (relay) |
| Info-Request | ✓ | ✓ (relay) |
| Relay-Forward | ✓ (LDRA) | ✓ (relay) |
| Relay-Reply | ✓ (relay) | ✓ (LDRA) |



Note:

- (relay) means supported only for relayed messages.
- (LDRA) means supported only for lightweight DHCPv6 relay agent (LDRA).

Operational commands

To check if a Python script is in service, use the following command.

```
show python python-script
```

To enable debugging for Python, use the following commands.

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

To bring a modified script in service, reload it using one of the following options:

- Use the following command.

```
tools perform python-script reload
```

- Use the **shutdown** and **no shutdown** commands for the modified Python script.



Note: When you modify a signed script, use the following command to sign it again.

```
tools perform python-script protect
```

11.1 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 an FTP server.

Step 2. Configure the URL of the script file.

```
configure python python-script
```



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.

```
configure python python-policy
```

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.

```
A:MAG-c>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
-----
A:MAG-c>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
-----

```

11.2 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.

```
tools perform python-script protect
```



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.

```
configure python python-script protection
```

Use the output file and password of the preceding step.

Related topics

[Python support](#)

12 BNG-UP resiliency

The Nokia MAG-c supports a MAG-c-driven BNG-UP resiliency scheme. Learn about this resiliency scheme, the resiliency handling, and deployment use cases.

12.1 Terminology for BNG-UP resiliency


| | |
|---------------------------------|--|
| fate sharing group (FSG) | 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. This maps to the TR-459 concept of a subscriber group (SGRP); FSGs use the PFCP extensions for SGRPs defined in TR-459. In the context of PFCP, the terms FSG and SGRP are interchangeable. |
| active BNG-UP | 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. |
| standby BNG-UP | 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. |
| hot standby | 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. |

12.2 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.


The following figure shows a high-level overview of communication for BNG-UP resiliency.

Figure 34: 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 35: Multiple backup BNG-UPs



To support all use cases, the MAG-c assigns sessions to an 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.

12.3 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 domain 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](#)

12.3.1 Fate sharing group creation

The MAG-c automatically creates FSGs for every 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 mode.

```
configure mobile-gateway pdn bng up-group fsg-mode
```

Currently, the following FSG modes are available:

- **single** – a single FSG is generated for the whole UP group and all sessions matching the group are assigned to that FSG
- **mesh-per-site** – a mesh of FSGs is generated for each site configured in the UP group and sessions are automatically assigned to an FSG (see [FWA UPF mesh](#) for more information on this model)

- **reference to an FSG profile**

Use the following command to configure a reference to an FSG profile.

```
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.



Note: The preferred indicator is not configurable when **fsg-mode** is set to **mesh-per-site**.

- **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](#)


12.3.2 Fixed access with broadcast 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 35: 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 36: Multiple Layer 2 access IDs per UP group



sw1374

```

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 a 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 [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
```

The following figure shows an example for a 1:1 hot standby resiliency with an S-tag per access node.

Figure 37: 1:1 hot standby resiliency example



sw1375

Example for a per S-tag 1:1 hot standby resiliency with an S-tag per access node

This example extends the previous model 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

```

The following figure shows an example of a per S-tag 1:1 hot standby resiliency.

Figure 38: 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](#)

12.3.3 Fixed access with active/standby pseudowire access

BNG-UP resiliency is possible with active/standby pseudowire (A/S PW) access. The access network selects the active and standby PW of an A/S PW, and the BNG-UPs must follow this decision.


The configuration of the FSG-based resiliency mechanism ensures that the FSG state follows the A/S state determined by the access node. This mechanism is modeled using UP groups as follows:

- The A/S PW pair is covered by a single UP group; that is, the Layer 2 access IDs representing the A/S PW cannot be present in any other UP group.
- The UP group covering the A/S PW pair cannot contain any other Layer 2 access IDs.
- The BNG-UP signals to the MAG-c a Layer 2 access ID health of 0% or 100% , which corresponds to the PW being standby (0%) or active (100%) on that BNG-UP. To configure this on a Nokia BNG-UP, see [Configuring an A/S PW port on the Nokia BNG-UP](#).
- The UP group covering the A/S PW pair must not contain any configuration that triggers the MAG-c to enforce an A/S state in the access network; for example, the **drain** or **preferred** configuration.

- The linked FSG profile cannot monitor network health because that can lead to the MAG-c enforcing an A/S decision that does not correspond to the PW A/S state.

The following figure and configuration output show an example of a PW and the associated UP group configuration.

Figure 39: A/S PW deployment



```
up-group "prefer-east"
  l2-access-id pw-1
  up "north"
  exit
  up "south"
  exit
  no shutdown
```



Note:

Because the **drain** command is not supported in this setup, Nokia recommends performing the following procedure when maintenance on a BNG-UP is required: [Maintenance on a Nokia BNG-UP in an A/S PW setup](#).

When the headless PFCP mode is enabled, exercise care to avoid active/active scenarios. When PFCP path management to one node fails, the MAG-c forces a switchover independent of the signaled Layer 2 access ID health. This can occur in the following indistinguishable cases:

- The entire BNG-UP has failed. The decision to switch over is correct. The A/S PW also performs a switchover.
- Only the PFCP path has failed (headless). The decision to switch over may be incorrect because the A/S PW state is independent of the PFCP path.

To handle the preceding cases, the following configuration is required:

- When in headless mode, the BNG-UP must fully bring down the PW. This forces the access node to also switch over because it perceives a failed BNG-UP. To configure this, see [Bringing down the PW in headless conditions](#).



Caution: If both BNG-UPs of an A/S PW pair become headless, forwarding stops but the session state is retained.

- On the MAG-c, PFCP headless mode must be aligned with BFD to make sure that UP and CP go into headless mode at similar intervals. See [Headless mode](#) for more information.
- The BNG-UP must be configured to move an FSG to standby whenever headless operations occur. To configure this, see [Moving an FSG to standby in headless mode](#).

12.3.3.1 Configuring an A/S PW port on the Nokia BNG-UP

Prerequisites

The basic SDP, PW port, and capture SAP configuration must be present.

About this task

Perform the following steps to configure an A/S PW port on the BNG-UP and enable the Nokia BNG-UP signal a health value of 0% or 100% to the MAG-c. For more information, see [Fixed access with active/standby pseudowire access](#).



Note: All commands in this procedure must be executed on the BNG-UP in the MD-CLI mode.

Procedure

Step 1. Configure an operational group to monitor the PW port.

```
configure service oper-group
```

Step 2. Apply the operational group to the PW port.

```
configure pw-port oper-group
```

Step 3. Apply the operational group for health monitoring on the capture SAP created for the PW port.

```
configure service vpls capture-sap pfc up-resiliency monitor-oper-group
```

Step 4. Set the drop in health value to 100 in the context created in step 3. This causes the reported health of the Layer 2 access ID to go from 100% to 0% when the PW becomes standby on the BNG-UP.

```
configure service vpls capture-sap pfc up-resiliency monitor-oper-group health-drop
```

Step 5. Optional: Define a Layer 2 access ID alias to simplify the correlation of the A/S PWs on two different BNG-UPs when provisioning UP groups on the MAG-c.

```
configure service vpls capture-sap pfc l2-access-id-alias
```

12.3.3.2 Maintenance on a Nokia BNG-UP in an A/S PW setup

About this task

Because the **drain** command is not supported in a BNG-UP resiliency setup with A/S PW access, Nokia recommends performing the following procedure for BNG-UP maintenance. For example, perform this procedure for upgrading the BNG-UP or restarting the PFCP association.

Avoid performing maintenance on multiple BNG-UPs in parallel. If multiple BNG-UPs require maintenance, execute the procedure for each BNG-UP one at a time.



Note: Most steps in this procedure contain commands that must be executed on the BNG-UP. Those commands are MD-CLI commands.

Procedure

Step 1. Verify that no FSG changes are ongoing for FSGs related to this BNG-UP; for example, changes as a result of finalizing a maintenance procedure for another BNG-UP.

Execute the following command to list all FSGs and UP groups related to the BNG-UP.

```
show mobile-gateway pdn bng up
```

Execute the following command for each FSG listed in the output of the preceding command.

```
show mobile-gateway pdn bng fsg
```

No changes are ongoing if the output indicates the following:

- **Standby Ready** field displays **yes**
- **Active Changing** field displays **no**
- **Current Hold-Off Delay** field displays **0 msec**
- **Locked** field displays **no**

If any of the preceding fields display a different value, a change is ongoing. Wait for the change to complete. To check progress, periodically execute the preceding command.

Step 2. On the BNG-UP, disable the PW port SDP.

```
configure pw-port sdp admin-state
```

Step 3. If a PFCP association restart is required, disable the PFCP association on the BNG-UP.

```
configure subscriber-mgmt pfcp association admin-state
```

Step 4. Execute the required maintenance operations on the BNG-UP.

Step 5. If a PFCP association restart was required in step 3, re-enable the PFCP association on the BNG-UP.

Step 6. Verify and wait until the MAG-c has synchronized all sessions.

On the MAG-c, execute the following command for all FSGs identified in step 1.

```
show mobile-gateway pdn bng fsg
```

Synchronization is complete if the output displays the following:

- **Standby Ready** field displays **yes**
- **Active Changing** field displays **no**
- **Current Hold-Off Delay** field displays **0 msec**
- **Locked** field displays **no**

If any of the preceding fields display a different value, synchronization is ongoing. Wait for the synchronization to complete. To check progress, periodically execute the preceding command.

Step 7. On the BNG-UP, re-enable the PW port.

```
configure pw-port sdp admin-state
```

Related topics

[Monitoring ongoing FSG changes](#)

12.3.3.3 Bringing down the PW in headless conditions

About this task

The BNG-UP must fully bring down the PW when operating in headless mode. This forces the access node to also switch over because it perceives a failed BNG-UP. To bring down the PW on a Nokia BNG-UP in headless conditions, perform the following steps on the Nokia BNG-UP. For more information, see [Fixed access with active/standby pseudowire access](#).



Note: All commands in this procedure must be executed on the BNG-UP in the MD-CLI mode.

Procedure

Step 1. Configure an operational group to monitor the path between the BNG-UP and the MAG-c.

```
configure service oper-group
```

Step 2. Enable BFD for the configured operational group.

In the following context, use the **interface-name** and **router-instance** commands to match the interface used by PFCP, and use the **dest-ip** command to set the IP address that the MAG-c uses for PFCP.

```
configure service oper-group bfd-liveness
```

Step 3. For the configured operational group, set the hold up time equal to the PFCP heartbeat interval and the hold down time greater than the BFD session recovery time.

```
configure service oper-group hold-time up
configure service oper-group hold-time down
```

Step 4. Enable BFD under the interface used for PFCP.

From the following list of commands, use only the commands that are relevant to the IP stack used for PFCP:

```
configure router interface ipv4 bfd
configure router interface ipv6 bfd
configure service vprn interface ipv4 bfd
configure service vprn interface ipv6 bfd
```

Although the **transmit-interval**, **receive**, and **multiplier** commands in the preceding contexts can be configured according to your needs, Nokia discourages sub-second failure detection to avoid aggressive switchovers upon minor interruptions.

- Step 5.** Set up the configured operational group for monitoring under the SDP linked to the PW port. This causes the system to bring down the SDP and PW port when in headless mode.

```
configure pw-port sdp monitor-oper-group
```

- Step 6.** Enable BFD assisted PFCP headless mode to align PFCP path management with the BFD decision when in headless mode.

```
configure subscriber-mgmt pfcpl association bfd-expedited-path-down
```

12.3.3.4 Moving an FSG to standby in headless mode

About this task

Perform the following steps on the Nokia BNG-UP to automatically move an FSG to standby when headless operations occur. For more information, see [Fixed access with active/standby pseudowire access](#).



Note: The commands that must be executed on the BNG-UP are in the MD-CLI mode.

Procedure

- Step 1.** Create an FSG template on the BNG-UP.

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

- Step 2.** In the configured FSG template on the BNG-UP, set the path restoration state to **standby**.

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

- Step 3.** Apply the FSG template to the FSG on the MAG-c.

- If the configured FSG template name is "default" and the MAG-c does not specify any template, the template is automatically applied and no further configuration is needed.
- If the configured FSG template name is different from "default", perform the following substeps.

- a.** Create an FSG template on the MAG-c.

```
configure mobile-gateway profile bng fsg-profile
```

- b.** In the configured FSG template on the MAG-c, set the BNG-UP FSG template name to the FSG template name configured in step 1.

```
configure mobile-gateway profile bng fsg-profile upf-fsg-template
```

- c.** Apply the MAG-c FSG template to the UP group for the A/S PW port on the MAG-c.


```
configure mobile-gateway pdn bng up-group fsg-profile
```

12.4 FWA UPF mesh

For fixed wireless access (FWA), there is a routed aggregation model between the UE, RAN, and the MAG-u. Consequently, each MAG-u in a single routed network can back up any other MAG-u in that same network. This allows for optimal spread of standby sessions and helps to minimize impact if a failure occurs. To achieve this, each MAG-u must be provisioned with as many FSGs as there are other MAG-u nodes. For an aggregation network with N MAG-u nodes, this means that each MAG-u is provisioned with $N-1$ FSGs for which it is active, for a total of $N \times (N-1)$ FSGs in the network. This forms a mesh of FSGs.

The MAG-c can automatically generate an FSG mesh based on the UP group configuration. When the UP group is configured, the MAG-c generates $N \times (N-1)$ FSGs, and each FSG is assigned two MAG-u nodes, where one is preferred. This automatic preference generation ensures that the UP is active for that FSG when all the UPs are healthy.

Figure 40: FSG mesh basic overview



sw4780

As shown in [Figure 40: FSG mesh basic overview](#), this generation results in two FSGs per UP pair, one FSG that prefers one of the UPs and the other FSG that prefers the other UP.



Note: Because the MAG-c sets the preferred mode automatically, it is no longer configurable with the `up-group` command.

To provision such a mesh, see [Provisioning an FWA UPF mesh](#).

12.4.1 Provisioning an FWA UPF mesh

About this task

For an optimal spread of standby sessions and minimal impact if a failure occurs, configure an FWA UPF mesh.



Note: For the complete list of the **configure** command parameters necessary for configuring an FWA UPF mesh, see the MAG-c CLI Reference Guide.

Procedure

Step 1. Create a new UP group using the following command.

```
configure mobile-gateway pdn bng up-group
```

Step 2. Configure the FSG mode for the UP group to **mesh-per-site** using the following command.

```
configure mobile-gateway pdn bng up-group fsg-mode mesh-per-site
```

Step 3. Configure all the UPs that are to be part of the mesh using the following command.

```
configure mobile-gateway pdn bng up-group up
```

Step 4. Create a new site within the UP group using the following command.

```
configure mobile-gateway pdn bng up-group site
```



Note: While site awareness itself is an optional feature (see [Inter-site redundancy for FWA sessions](#) for more information), the FSG mesh always needs at least one site to be configured. This allows easy migration to site awareness at a later time.

Step 5. Add all the MAG-u nodes to the newly created site using the following command.

```
configure mobile-gateway pdn bng up-group site up
```

Step 6. Configure the GTP-u endpoint addresses to be assigned to the FSG using the following commands.

```
configure mobile-gateway pdn bng up-group gtp-u-endpoints realm ipv4-range  
configure mobile-gateway pdn bng up-group gtp-u-endpoints realm ipv6-prefix
```

For more information about GTP-u endpoint allocation per FSG, see [Traffic steering parameters](#).

Step 7. Configure an FSG profile using the following command.

```
configure mobile-gateway profile bng fsg-profile
```

Step 8. In the configured FSG profile, use the following command to enable the network realms that require health monitoring, for example the GTP-u endpoint realms (S1-u/N3) and the network-facing realms (SGi/N6).

```
configure mobile-gateway profile bng fsg-profile health-calculation network-realm
```



Note: This instructs the MAG-c to take only the health of that network realm into account. It is also necessary to configure the MAG-u to signal a sensible health value for the configured realms. See the 7750 SR and VSR BNG CUPS User Plane Function Guide for more information.

- Step 9.** Configure any other FSG parameters, and when finished, apply that FSG profile to the BNG UP group using the following command.

```
configure mobile-gateway pdn bng up-group fsg-profile
```

- Step 10.** Enable the UP group using the following command.

```
configure mobile-gateway pdn bng up-group no shutdown
```

- Step 11.** Investigate the state of the created UP group by executing the following command with the name of the created UP group as a parameter.

```
show mobile-gateway pdn bng up-group
```

This command lists all the configured UPs and $N*(N-1)$ FSGs, where N is the number of UPs. For each UP, the following information should be stated:

- health of 100
- list $N-1$ FSGs for which the UP is active
- list $N-1$ FSGs for which the UP is on standby

See [Operational commands](#) for more information about this command.

The following example output displays the properties of the created mesh.

Example

Mesh properties

```
A:MAG-c>show>mobile-gateway>pdn>bng>up-group# "fwa_mesh"
=====
Up-Group 'fwa_mesh'
=====
PDN gateway           : 1
Description           : (Not Specified)
Admin State           : up

FSG Profile           : fwa
  Default Standby Type : hot

-----
      FSG Active UP           Standby UP           Sessions  Hold-Off
-----
      2 up-east               up-central (ready)           0          0
      3 up-central            up-east (ready)              0          0
      4 up-west               up-central (ready)           0          0
      5 up-central            up-west (ready)              0          0
      6 up-west               up-east (ready)              0          0
      7 up-east               up-west (ready)              0          0
-----
No. of FSGs: 6
-----

-----
UP                               Pref Drain Health  Fsg  Role
-----
up-central (192.0.2.13)          no  no    100  2    Std
                                   3    Act
                                   4    Std
                                   5    Act
```

```

up-east (192.0.2.11)          no  no    100    2      Act
                             3      Std
                             6      Std
up-west (192.0.2.12)       no  no    100    4      Act
                             5      Std
                             6      Act
                             7      Std
-----
No. of UPs: 3
-----

-----
Site                               Rebalance
-----
main                               no
-----
No. of Sites: 1
=====

```

This output suggests that:

- Six FSGs are created, following the $N \times (N-1) = 3 \times 2$ formula.
- Each FSG has two associated MAG-u nodes.
- Each UP is 100% healthy.
- Each UP pair has two FSGs, for example, up-west and up-east share FSGs six and seven, where up-west is active (preferred) for FSG six and up-east is active (preferred) for seven.

Step 12. After verifying the UP group is correctly provisioned, assign FWA sessions to the UP group and site using the following command.

```

configure mobile-gateway profile authentication-database entry fixed-wireless-access
up-group

```

Example

Mesh generation in UP group with three UPs

```

A:MAG-c>config>mobile-gateway>pdn>bng# info
-----
    up-group "fwa_mesh"
      fsg-mode mesh-per-site
      fsg-profile "fwa"
      gtp-u-endpoints
        realm "to_ran"
        ipv4-range 192.0.2.50 192.0.2.100
        ipv6-prefix 2001:db8::/64
      exit
    exit
  up "up-central"
  exit
  up "up-east"
  exit
  up "up-west"
  exit
  site "main"
    up "up-central"
    up "up-east"
    up "up-west"
  exit

```

```
no shutdown
exit
```

12.4.2 Inter-site redundancy for FWA sessions

It is sometimes desirable to logically segment MAG-u nodes and assign sessions to a specific subset of MAG-u nodes by default. For example, for latency purposes, sessions are assigned to a set of nearby MAG-u nodes. In the context of the MAG-c, such a set of MAG-u nodes is known as a **site**.



Note: The name site implies MAG-u nodes share a physical location. While this is an intended use case, the MAG-c does not require any physical colocation or other distance requirements between MAG-u nodes in a single site. It is merely a logical grouping.

For information about how to configure additional sites and assign MAG-u nodes to them, see [Creating sites with assigned MAG-u nodes](#).

The benefit of using multiple sites within a UP group (instead of multiple UP groups), is it achieves inter-site redundancy. The MAG-c tracks the number of failed MAG-u nodes for each site. If the number of failed MAG-u nodes meets or exceeds a provisioned threshold, the MAG-c stops assigning sessions to FSGs within the failed site. Instead, it assigns sessions directly to MAG-u nodes in other sites, bypassing FSG functionality.

12.4.2.1 Creating sites with assigned MAG-u nodes

About this task

Adding a MAG-u to a site logically segments it, which can be useful for latency and inter-site redundancy purposes.



Note: For the complete list of the **configure** command parameters necessary for creating sites with assigned MAG-u nodes, see the MAG-c CLI Reference Guide.

Procedure

Step 1. Add the MAG-u nodes to the UP group using the following command.

```
configure mobile-gateway pdn bng up-group up
```

Step 2. Create a site within the UP group using the following command.

```
configure mobile-gateway pdn bng up-group site
```

Step 3. Add the MAG-u nodes to the newly created site using the following command.

```
configure mobile-gateway pdn bng up-group site up
```

Step 4. Configure the behavior of inter-site failover using the following command.

```
configure mobile-gateway pdn bng up-group site inter-site-failover
```

Depending on your preferences, use the following options:

- **failed-up-threshold**

A site is declared as failed when the number of failed MAG-u nodes within it is bigger than or equal to the number of UPs in the site or the value set with this command, whichever is smaller. A MAG-u itself is declared as failed based on its health signaling and the failure threshold configured with the following command.

```
configure mobile-gateway profile bng fsg-profile health-calculation failure-
threshold
```

See [BNG-UP health determination](#) for more information.

- **backup-up**

When a site fails, the MAG-c assigns a session to any non-failed MAG-u within all other sites by default. This command overrides this function and instead provides an explicit list of backup UPs to be considered.

- **pool-overrides**

The commands in this context apply to sessions that were originally assigned to a site that failed, and are now installed on a MAG-u of this site. Instead of using the pools assigned to the session during authentication, the pools in this context are used. This can be useful, for example, if dedicated pools are used per site and the original pools are not available on this site.

12.4.3 Anti-affinity to avoid FSG creation

In certain cases, it is undesirable to automatically provision FSGs between two MAG-u nodes. For example, if a subset of MAG-u nodes is likely to fail because of strict colocation, it is possible to prevent the creation of FSGs for those MAG-u nodes by applying an anti-affinity label to each of them using the following command.

```
configure mobile-gateway pdn bng up-group up anti-affinity-label
```

Example: Configuration of single shared affinity label

The MAG-c does not create any FSGs between UPs with the same label. The following example shows a single affinity label shared between the MAG-u nodes "up-east" and "up-west".

```
A:MAG-c>config>mobile>pdn>bng# info
-----
      up-group "fwa_mesh"
        fsg-mode mesh-per-site
        fsg-profile "fwa"
        gtp-u-endpoints
          realm "to_ran"
          ipv4-range 192.0.2.50 192.0.2.100
          ipv6-prefix 2001:db8::/64
        exit
      exit
      up "up-central"
      exit
      up "up-east"
        anti-affinity-label "avoid_fsg"
      exit
      up "up-west"
        anti-affinity-label "avoid_fsg"
      exit
```

```

site "main"
  up "up-central"
  up "up-east"
  up "up-west"
exit
no shutdown
exit
-----

```

Example: FWA mesh information for UP groups

The output of the following example, and [Figure 41: Mesh anti-affinity](#), display the result, which is that two less FSGs have been created and there is no FSG present that contains both those MAG-u nodes.

```


A:MAG-c>show>mobile-gateway>pdn>bng>up-group# "fwa_mesh"
=====
Up-Group 'fwa_mesh'
=====
PDN gateway                : 1
Description                 : (Not Specified)
Admin State                 : up
FSG Profile                 : fwa
  Default Standby Type     : hot
-----
      FSG Active UP          Standby UP          Sessions  Hold-Off
-----
      8 up-east             up-central (ready)          0          0
      9 up-central         up-east (ready)             0          0
     10 up-west            up-central (ready)          0          0
     11 up-central         up-west (ready)             0          0
-----
No. of FSGs: 4
-----

UP                          Pref Drain Health   Fsg   Role
-----
up-central (192.0.2.13)     no  no      100   8     Std
                             9     Act
                             10    Std
                             11    Act
up-east (192.0.2.11)       no  no      100   8     Act
                             9     Std
up-west (192.0.2.12)      no  no      100  10     Act
                             11    Std
-----
No. of UPs: 3
-----

Site                          Rebalance
-----
main                          no
-----
No. of Sites: 1
=====

```

Figure 41: Mesh anti-affinity



sw4779

12.4.4 Session rebalance and revert

Sub-optimal session distributions can occur in multiple scenarios, such as the following:

- New sessions are not installed on a subset of FSGs because the associated MAG-u nodes failed. This leads to a load imbalance across FSGs.
- New sessions ignore load-based FSG selection but are steered to specific FSGs for a number of possible reasons (see [Session-to-FSG mapping and load-balancing](#) for more information). This leads to a load imbalance across FSGs.
- New sessions are installed on another site because of site failures (see [Inter-site redundancy for FWA sessions](#) for more information). This leads to sessions not being present on the preferred site.

To resolve this, the MAG-c offers revert and rebalancing functions. Both these functions disconnect specific sessions to resolve the sub-optimal session distribution. The sessions are signaled with NAS cause code #39 (Reactivation Requested) on 5G systems and with GTP cause value #8 (Reactivation Requested) on 4G systems, to trigger the UE to immediately reconnect.

The full list of revert and rebalancing functions is as follows:

- **revert**

The revert function disconnects all sessions that should be set up on an FSG of a specific target site but are active on a UP from another site.

Use the following command to trigger a revert action in the specified site.

```
clear mobile-gateway pdn bearer-context up-group target-site action revert
```



Note: See the MAG-c CLI Reference Guide for the complete list of available parameters for the **clear** command.

- **rebalance**

The rebalance function periodically disconnects a number of sessions in the specified target site until the difference between the least loaded UP and the most loaded UP is less than or equal to the specified difference. If no difference is specified, it defaults to 5%.

Use the following command to trigger the rebalance action in the specified site.

```
clear mobile-gateway pdn bearer-context up-group target-site action rebalance
```



Note: See the MAG-c CLI Reference Guide for the complete list of available parameters for the **clear** command.

To see if a rebalance action is in progress, use the following command, with or without the **detail** parameter.

```
show mobile-gateway pdn bng up-group
```

For example:

```
A:MAG-c>show>mobile-gateway>pdn>bng>up-group# "fwa_mesh"
...
-----
Site                                     Rebalance
-----
main                                     yes
-----
```

If you add the **detail** parameter, you can also see the current loop and the max loops that the system executes.

For example:

```
A:MAG-c>show>mobile-gateway>pdn>bng>up-group# "fwa_mesh" detail
...
-----
Site main
-----
up-east (192.0.2.11)
up-west (192.0.2.12)
up-central (192.0.2.13)

Inter Site Failover
  Failed UP Threshold           : 8

Rebalance
  Ongoing                       : yes
  Current Loop                   : 6
  Max Loops                      : 50
-----
```

- **cancel-rebalance**

To cancel an ongoing rebalance action, use the following command.

```
clear mobile-gateway pdn bearer-context up-group target-site action cancel-rebalance
```

- **time-of-day**

Optionally, it is possible to set up the MAG-c to automatically execute both revert and rebalance actions at a specified time of day using the following command.

```
configure mobile-gateway profile bng fsg-profile auto-revert-rebalance time-of-day
```

All the commands related to the MAG-c revert and rebalancing functions are subject to a deletion rate. Nokia recommends to use the lowest possible deletion rate to minimize impact on other call flows.

You can provision the deletion rate using the following command.

```
configure mobile-gateway pdn session-deletion-rate
```

12.4.5 Adding a MAG-u to an existing UP group for FWA sessions

Prerequisites

In a geo-redundant MAG-c deployment, you must apply all the configuration commands on the standby node before applying them on the active node.

Procedure

Step 1. Add the MAG-u to the UP group using the following command.

```
configure mobile-gateway pdn bng up-group up
```



Note: When configuring an IPv4 address, make sure there are sufficient addresses in the IPv4 ranges configured with the **gtp-u-endpoints** command; otherwise, the system may reject adding a new UP. This is not the case when configuring IPv6 addresses, because a prefix is used instead of a range.

Step 2. Apply any configuration, such as an anti-affinity-label, to the UP.



Note: Most UP configurations cannot be changed after adding the UP to a site.

Step 3. Add the MAG-u to a site using the following command.

```
configure mobile-gateway pdn bng up-group site up
```

What to do next

After adding a MAG-u to an existing UP group, the MAG-c starts creating new FSGs to restore the full mesh, and new sessions are assigned to those FSGs based on load, as usual. To expedite rebalancing sessions to the new FSGs, you can execute a rebalance action as described in [Session rebalance and revert](#).

12.4.6 Removing a MAG-u from an existing UP group

Prerequisites

Execute steps 1 through 3 of [Maintenance on a Nokia BNG-UP](#).

Procedure

Step 1. Remove the MAG-u from the site by using the following command.

```
configure mobile-gateway pdn bng up-group site no up
```

Step 2. Remove the MAG-u from the UP group by using the following command.

```
configure mobile-gateway pdn bng up-group no up
```

What to do next

The MAG-c does not automatically remove all FSGs associated with the UP. Instead, it marks them as lingering and stops assigning sessions to them. After all the sessions of an FSG are removed, the FSG itself is removed.

To expedite this process, use the following command to clear all the sessions that are attached to lingering FSGs.

```
clear mobile-gateway pdn bearer-context up-group target-site action clear-lingering-fsgs
```

12.5 Maintenance on a Nokia BNG-UP

About this task

Nokia recommends performing the following procedure for BNG-UP maintenance where session stability cannot be guaranteed; for example, for upgrading the BNG-UP or restarting the PFCP association.

Avoid performing maintenance on multiple BNG-UPs in parallel. If multiple BNG-UPs require maintenance, execute the procedure for each BNG-UP one at a time.



Note:

In a geo-redundant MAG-c deployment, you must apply all configuration commands first on the standby node, and then the active node.

Execute the **admin save** command after each configuration step.

Procedure

Step 1. Verify that no FSG changes are ongoing for FSGs related to the BNG-UP; for example, changes as a result of finalizing a maintenance procedure for another BNG-UP.

List all FSGs and UP groups related to the BNG-UP.

```
show mobile-gateway pdn bng up
```

Execute the following command for each FSG listed in the output of the preceding command.

```
show mobile-gateway pdn bng fsg
```

No changes are ongoing if the output indicates the following:

- **Standby Ready** field displays **yes**
- **Active Changing** field displays **no**

- **Current Hold-Off Delay** field displays **0 msec**
- **Locked** field displays **no**

If any of the preceding fields display a different value, a change is ongoing. Wait for the change to complete. To check progress, periodically execute the preceding command.

If any of the **Drain** fields display **yes**, a BNG-UP is set to **drain** mode. This may indicate that a maintenance procedure for another BNG-UP is ongoing. Verify if this is intended and, if necessary, correct by completing the maintenance procedure for the specific BNG-UP.

Step 2. Enable the **drain** mode for the BNG-UP in all UP groups identified in step 1.

```
configure mobile-gateway pdn bng up-group up drain
```

Step 3. Verify that the BNG-UP is no longer active in any FSG and that the FSGs are stable.

List the BNG-UP state in all FSGs.

```
show mobile-gateway pdn bng up
```

The BNG-UP is no longer active if the **Role** column is different from **Act** for all entries in the list. If this is not the case, wait until the value in the **Role** column is different from **Act** for all entries. Periodically execute the preceding command to check this information.

Execute the following command for each FSG identified in step 1.

```
show mobile-gateway pdn bng fsg
```

No changes are ongoing if the output indicates the following:

- **Active Changing** field displays **no**
- **Current Hold-Off Delay** field displays **0 msec**
- **Locked** field displays **no**

If any of the preceding fields display a different value, a change is ongoing. Wait for the change to complete. To check progress, periodically execute the preceding command.

Step 4. Execute the required maintenance operations (for example, an upgrade or PFCP Association Restart) on the BNG-UP.

Step 5. After the BNG-UP maintenance operations have fully completed, perform step 3 to verify that the FSGs are stable.

Step 6. Disable the **drain** mode for the BNG-UP in all its related UP groups.

```
configure mobile-gateway pdn bng up-group up no drain
```

Related topics

[Monitoring ongoing FSG changes](#)

12.6 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 and load-balancing](#)).
- 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. The MAG-c may execute multiple subsequent FSG changes to accomplish this.

Related topics

[Fate sharing group creation](#)

12.6.1 Session-to-FSG mapping and load-balancing

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.

When setting up a fixed wireless access in mesh mode (see [FWA UPF mesh](#)), the MAG-c derives a UP group and site using the **up-group** command in the following context.

```
configure mobile-gateway profile authentication-database entry fixed-wireless-access
```

Both the UP group and the site parameters must be provisioned to map an FWA session to an FSG.

Based on the configuration of this command, the MAG-c assigns the session to an FSG using the following criteria, in order of priority:

1. If the mapped site fails, the session is not assigned to an FSG; instead, it is set up in a non-resilient way on a UP in another site (see [Inter-site redundancy for FWA sessions](#) for more information). In this case, the same logic applies to the subsequent criteria, but a direct UP is chosen instead of an FSG.
2. In some cases, session setup parameters are taken into account to prefer a specific FSG or set of FSGs. In such cases, a specific FSG is selected or the number of eligible FSGs is reduced. For example:
 - When a second session of a UE is set up, the MAG-c prefers the same FSG as in the previous session.
 - When a session restarts and the held IP is recovered (see [Session lifetime and held addresses](#)), the MAG-c prefers the FSG that holds the ODSA micro-net for that held IP address.
 - When a session is started with a static IP, and unnumbered interfaces are not used, the MAG-c prefers the FSG that holds the ODSA micro-net for that static IP address.

- When a preferred UP is provisioned for a session, the MAG-c prefers FSGs that are preferred for that UP.
 - When both the active and standby UPs fail for an FSG, the FSG is excluded from selection.
3. The least loaded FSG is chosen, where load is based on the signaled load of the preferred UP for that FSG. To enable learning load from the UP, use the following command.

```
configure mobile-gateway pdn pfc-p-load-control enable-pfc-p-load-control
```

4. If all else is equal, the MAG-c distributes sessions evenly over the remaining FSGs.

12.6.2 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.

Fixed access

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 six FSGs with two subnets each, distributed over three 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 42: Example of the relationship between FSGs, MAC addresses, and subnets



sw1378

FWA

For FWA sessions, the MAG-c generates a unique set of GTP-u endpoint addresses per FSG. The active MAG-u can attract the traffic similar to ODSA prefixes. The MAG-c allocates the addresses based on how the following commands are configured.

```
configure mobile-gateway pdn bng up-group gtp-u-endpoints realm ipv4-range
configure mobile-gateway pdn bng up-group gtp-u-endpoints realm ipv6-prefix
```

For each realm configured, the MAG-c allocates either an IPv4 address, an IPv6 address, or both, as follows:

- If an IPv4 range is configured, an address from that range is assigned to the FSG.
- If an IPv6 prefix is configured, the FSG ID is directly encoded in the last 16 bits of the IPv6 prefix to create a unique /128 IPv6 address.



Caution: The list of realms provisioned in this context must match all the realms provisioned using the following commands.


```
configure mobile-gateway pdn s11 interface gtp-c interface-realm
configure mobile-gateway pdn s5 interface gtp-c interface-realm
configure mobile-gateway pdn sba-server-services nsmf-pdusession n3-interface-realm
```

[Figure 43: Example of the relationship between FSGs, GTP-u endpoint addresses, and subnets](#) shows the GTP-u endpoint addressing for six FSGs, each with two ODSA subnets and distributed over three BNG-UPs. A single GTP-u endpoint realm is configured with IPv6 prefix 2001:db8::/112, and IPv4 range 192.0.2.50 to 192.0.2.100. The following list identifies the relationship between the FSGs, GTP-u endpoint addresses, and subnets:

- FSG 1
 - GTP-u endpoint IPv6 2001:db8::1
 - GTP-u endpoint IPv4 192.0.2.50
 - session subnet 10.1.1.0/24
 - session subnet 10.1.2.0/24
- FSG 2
 - GTP-u endpoint IPv6 2001:db8::2
 - GTP-u endpoint IPv4 192.0.2.51
 - session subnet 10.2.1.0/24
 - session subnet 10.2.2.0/24
- FSG 3
 - GTP-u endpoint IPv6 2001:db8::3
 - GTP-u endpoint IPv4 192.0.2.52
 - session subnet 10.3.1.0/24
 - session subnet 10.3.2.0/24
- FSG 4
 - GTP-u endpoint IPv6 2001:db8::4

- GTP-u endpoint IPv4 192.0.2.53
- session subnet 10.4.1.0/24
- session subnet 10.4.2.0/24
- FSG 5
 - GTP-u endpoint IPv6 2001:db8::5
 - GTP-u endpoint IPv4 192.0.2.54
 - session subnet 10.5.1.0/24
 - session subnet 10.5.2.0/24
- FSG 6
 - GTP-u endpoint IPv6 2001:db8::6
 - GTP-u endpoint IPv4 192.0.2.55
 - session subnet 10.6.1.0/24
 - session subnet 10.6.2.0/24

Figure 43: Example of the relationship between FSGs, GTP-u endpoint addresses, and subnets



sw4778

Related topics

[ODSA](#)

12.6.3 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%.



Note: The per Layer 2 access ID status value is not used for FWA BNG-UP resiliency, where network instance health is also used toward the access.

- 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 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
```



Note: For FWA BNG-UP resiliency, no Layer 2 access ID is configured in the UP group and no Layer 2 access ID health is tracked independent of this configuration.

- 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 does not count that entity in the aggregation calculation.

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 recovered 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.

The following table 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 25: Summary of BNG-UP states

| State | Description | Sources |
|--------|---|--|
| health | A value between 0% and 100% or the special value -1 (unavailable) | Aggregation of the per-logical-port and per-network-realm health reports from the BNG-UP |

| State | Description | Sources |
|---------------------|---|--|
| | Indicates the health of the BNG-UP | <p>PFCP path management state (for example, headless).</p> <p>The drain mode configured with the following command.</p> <pre>configure mobile-gateway pdn bng up-group up drain</pre> |
| failed indicator | <p>An indicator that considers the BNG-UP failed if its health is less than the failure threshold</p> <p>Enables switchovers in more restrictive (for example, non-revertive) scenarios</p> | <p>Based on the health state and the threshold configured with the following command.</p> <pre>configure mobile-gateway profile bng fsg-profile health-calculation failure-threshold</pre> |
| not-ready indicator | <p>An indicator on the standby BNG-UP that does not have all hot standby sessions installed</p> <p>Kept until the standby BNG-UP has installed the hot standby sessions</p> | <p>Set for each new standby BNG-UP or a standby BNG-UP whose health becomes unavailable (-1).</p> <p>Removed after the first successful FSG change when the health is 0% or higher.</p> |
| lockout | <p>A failure state in which the BNG-UP cannot be made active or standby</p> <p>Kept until the BNG-UP is no longer active or standby and a lockout timer has expired</p> | <p>Applied automatically for multiple failure scenarios, see UP lockout for more information.</p> |

12.6.4 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.

12.6.5 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 startup 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.
- **time-of-day**
The MAG-c uses the revertive behavior each day during a specified time period. This option is useful to allow revertive mode at times when impact is expected to be low. In addition, the MAG-c automatically triggers a reselection procedure of all FSGs at the start of this period to allow changes under relaxed revertiveness conditions.

When non-revertive mode applies, it is possible to force a reselection and bypass the revertive mode using the following command.

```
tools perform mobile-gateway pdn bng up-resiliency force-reselection up-group
```

This can be useful when the non-revertive mode is preferable to minimize traffic impact, but it is also necessary to manually trigger a revertive change at a specific time.

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 that has priority over any backup BNG-UP to avoid 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 also 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.

12.6.6 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:
 - **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. The following figure 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 44: GARP race conditions



sw1379

12.6.7 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](#))



Note: This only applies to a full PFCP path down, and not to headless mode. For more information about the differences, see [Headless mode](#). For more information about the interaction of BNG-UP resiliency with headless mode, see [Interaction with headless mode](#).

The MAG-c sees placing 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 performs 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 one time and starts from a clean slate before it can be made active or standby again.

12.6.8 Monitoring ongoing FSG changes

About this task

At any time, an FSG change can be in progress or a reevaluation of the active and standby BNG-UPs for the FSG can be queued. In specific cases, for example, before executing maintenance procedures, it is important to know if an FSG is undergoing changes.

Procedure

To monitor ongoing FSG changes, inspect the fields in the output of the following command.

```
show mobile-gateway pdn bng fsg
```



Note: When using model-driven interfaces, the same output can be retrieved using the following context.

```
state mobile-gateway pdn bng fsg
```

See [YANG state](#) for more information.

Use the following fields in the displayed output to determine whether an FSG change is ongoing or upcoming:

- **Current Hold-Off Delay**
This field indicates if a received change trigger is subject to a hold-off period. If the displayed value is **0 msec**, no change triggers are held off. See [Active/standby selection triggers](#) for more information about change triggers and the hold time.
- **Active Changing**

This field indicates if a change of the active BNG-UP is ongoing. If the displayed value is **no**, the active BNG-UP is not being changed; if it is **yes**, the active BNG-UP is being changed.

- **Locked**

This field indicates if a BNG-UP is currently in lockout state. If the displayed value is **no**, the BNG-UP is not in lockout state. Any other value indicates that the BNG-UP is in lockout state and the duration it will remain in the lockout state. This is an important indicator because a locked out BNG-UP periodically attempts to get out of lockout, which triggers an FSG change. See [UP lockout](#) for more information about lockout.

- **Standby Ready**

This field indicates if a change of the standby BNG-UP is ongoing. If the displayed value is **yes**, the standby BNG-UP is not being changed; if it is **no**, the standby BNG-UP is being changed or is in another condition where it is unfit to replace the active BNG-UP, for example, because it is in drain state. In case of the hot standby resiliency model, a change of the standby BNG-UP can take time if sessions have to be recreated on the standby BNG-UP. Use the following command to estimate how many sessions still have to be created.

```
show mobile-gateway bng session standby-state hot-to-create count
```

Example: FSG with two BNG-UPs

```
A:MAG-c# show mobile-gateway pdn bng fsg 1
=====
Fate Sharing Group 1
=====
PDN gateway                : 1
Up-Group                   : broadcast_domain
No. Of Sessions            : 0
Current Hold-Off Delay     : 0 msec
Virtual MAC                 : 02:00:5e:00:00:01

FSG Profile                 : basic
  Default Standby Type     : hot

Active UP                   : up-east (192.0.2.11)
  Active Changing          : no
  Locked                   : no
  Preferred                 : yes
  Drain                    : no
  Health                   : 100
  S-Vlan Range             : N/A
  C-Vlan Range             : N/A

Standby UP                  : up-west (192.0.2.12)
  Standby Ready            : yes
  Locked                   : no
  Preferred                 : no
  Drain                    : no
  Health                   : 100
  S-Vlan Range             : N/A
  C-Vlan Range             : N/A
=====
```

12.7 Hot standby

In the context of BNG-UP resiliency, hot standby is a per-session concept that defines how a session is handled on the standby BNG-UP. Hot standby sessions are precreated on the standby BNG-UP, which can start forwarding traffic for those sessions as soon as it becomes active. This offers significantly reduced forwarding loss during switchovers, and 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 (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 removes the full session.

12.8 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, uses a heuristic process to determine whether to keep FSGs active or make them standby during headless operations. In rare cases, the BNG-UP may keep an FSG active while the MAG-c has successfully made another BNG-UP active. As a result, 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:

- Uplink 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.



Note: Downlink QoS can still be guaranteed when the non-headless BNG-UP announces routes with a higher preference than the headless BNG-UP to consistently forward downlink traffic through the non-headless BNG-UP. Additionally, if the access network updates its uplink forwarding based on downlink traffic, uplink traffic is forwarded through the non-headless BNG-UP.

- 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. However, this is extremely unlikely as the FSG MAC is most likely known at any point in time.

For more information about the headless heuristics and the downlink routing differentiation, see the *7750 SR and VSR BNG CUPS User Plane Function Guide*.

To avoid the unwanted consequences of the active/active state, configure the Nokia BNG-UP 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.

12.9 Operational commands

The same operational commands can be used for resilient sessions as for nonresilient 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.

13 KPI and KCI performance monitoring

The MAG-c supports periodic collection of Key Performance Indicator (KPI) and Key Capacity Indicator (KCI) statistics for control plane signaling and system resource usage. The MAG-c writes the KPI and KCI data as XML per collection interval.

KPI and KCI counters are used for post-processing of system behavior for control plane (CP) signaling and system resource usage during peak and non-peak hours. Statistical data about system behavior helps users understand network activity and plan for signaling load and resources needed.

The MAG-c system periodically collects statistical data and stores KPI and KCI counters in files generated at the end of each collection interval. The system writes the KPI and KCI files in XML format in compact flash directories. Counters are grouped in counter modules (tables of counter values). The CLI configuration, reference points, and features determine if a table is produced in a KPI or KCI file. To save space, if all the counters in a module are zero for a reporting period, the MAG-c omits the module from the XML file.

The MAG-c reports statistical measurements in a variety of dimensions. For example:

- per ref-point interface per peer
- per APN/DNN

The dimension reporting is fixed for each type of counter module. Available system resources that can be scaled may also report statistics, (for example, per SM-VM) or system-wide (for example, for all OAM-VMs).

Most KPI statistics are reported as deltas (except for system statistics and KCI statistics are absolute values).

The XML files contain a large set of raw statistics covering MAG-c functional areas and features. Post-processing collection and analytics tools can aggregate or consume these statistics to compute system-level KPIs and KCIs. For example:

- Success% = (Success) / (Success + Failures)
- Attach Rate = (Success) / Collection interval

13.1 KPI and KCI XML file naming

The MAG-c applies the following convention for the KPI and KCI measurement result filenames.

```
<Type><Startdate>.<Starttime>[-<offset>]-[<Enddate>].<Endtime>[-<offset>][_<jobId>][_<Unique Id>]
```

The value of <Type> depends on the collection interval and rollover interval. The possible values of <Type> are:

- **A** – the collection interval equals the rollover interval
- **C** – the collection interval is lower than the rollover interval

The optional <Enddate> variable is only included when the <type> equals **C**.

Example: KPI or KCI XML filename

```
A20250524.0345-0500-0400-0500_-022150_MAG-NODE-1.xml
```

- 20250524 = Start date (YYYYMMDD format)
- 0345 = interval start time
- 0400 = interval end time
- 0500 = UTC offset
- 022150 = jobId
- MAG-NODE-1 = UniqueId (epc-node-name)

13.2 Storage and retrieval of KPI and KCI XML files

The MAG-c stores generated XML files in the local file system, from which an external management system can retrieve them using SFTP or FTP.

Nokia recommends using the CF2: card for storing the KPI and KCI XML files. This is because the system primarily uses the CF1: card to store configuration files and images. If the location is not defined, CF2: is the default location for accounting policies.

The system compresses the XML files using the gzip format and writes them temporarily in the local directory. After closing the files (rollover), the system moves them to a common directory for storage. After the rollover, an SNMP trap is sent to inform the external management system that a new set of files is available for retrieval.

13.3 Content and format of KPI and KCI XML files

The performance measurements are contained in XML files.

The XML file format complies with the following 3GPP standards:

- The performance measurement XML file uses the following:
 - naming conventions specified in 3GPP TS 32.432
 - format specified in 3GPP TS 32.401
- TS 32.401 Concept and requirements
- TS 32.404 Performance measurements - definitions and template
- TS 32.406 Performance measurements Core Network (CN) Packet Switched (PS)

Records inside the XML file are stored by reporting interval. All records are contained in the same file for that reporting interval for a specific accounting policy. If the rollover interval is larger than the collection interval, a file contains statistical values from multiple collection intervals.

13.4 Record types and counter module mappings

Users can find the list of KPI and KCI counters with associated modules and record types in the *MAG-c KPI and KCI Counters Spreadsheet*.

The MAG-c supports the following unique record types that apply for MAG-c statistics:

- kpi-system
- kpi-bearer-mgmt
- kpi-ref-point
- kpi-path-mgmt
- kci-system
- kci-bearer-mgmt
- kci-path-mgmt
- complete-kpi
- complete-kci
- kpi-ref-path-group
- kpi-kci-bearer-mgmt
- kpi-kci-path-mgmt
- kpi-kci-system
- complete-kpi-kci
- kpi-ref-pt-failure-cause-code

To obtain the list of all record types, use the following command.

```
show log accounting-records
```

See the 7450 ESS, 7750 SR, 7950 XRS, and VSR Clear, Monitor, Show, and Tools CLI Command Reference Guide for more information about the command.

Output example

```
A:MAG-c# show log accounting-records
=====
Accounting Policy Records
=====
Record # Record Name                               Def. Interval
-----
1      service-ingress-octets                           5
2      service-egress-octets                             5
3      service-ingress-packets                           5
4      service-egress-packets                            5
5      network-ingress-octets                            15
6      network-egress-octets                             15
7      network-ingress-packets                           15
8      network-egress-packets                            15
9      compact-service-ingress-octets                    5
10     combined-service-ingress                           5
11     combined-network-ing-egr-octets                    15
```

| | | |
|----|------------------------------------|----|
| 12 | combined-service-ing-egr-octets | 5 |
| 13 | complete-service-ingress-egress | 5 |
| 14 | combined-sdp-ingress-egress | 5 |
| 15 | complete-sdp-ingress-egress | 5 |
| 16 | complete-subscriber-ingress-egress | 5 |
| 23 | custom-record-subscriber | 5 |
| 24 | custom-record-service | 5 |
| 25 | custom-record-aa-sub | 15 |
| 26 | queue-group-octets | 15 |
| 27 | queue-group-packets | 15 |
| 28 | combined-queue-group | 15 |
| 29 | combined-mpls-lsp-ingress | 5 |
| 30 | combined-mpls-lsp-egress | 5 |
| 31 | combined-ldp-lsp-egress | 5 |
| 32 | saa | 5 |
| 33 | video | 10 |
| 34 | kpi-system | 5 |
| 35 | kpi-bearer-mgmt | 5 |
| 37 | kpi-ref-point | 5 |
| 38 | kpi-path-mgmt | 5 |
| 40 | kci-system | 5 |
| 41 | kci-bearer-mgmt | 5 |
| 42 | kci-path-mgmt | 5 |
| 43 | complete-kpi | 5 |
| 44 | complete-kci | 5 |
| 46 | kpi-ref-path-group | 5 |
| 47 | kpi-kci-bearer-mgmt | 5 |
| 48 | kpi-kci-path-mgmt | 5 |
| 49 | kpi-kci-system | 5 |
| 50 | complete-kpi-kci | 5 |
| 52 | complete-ethernet-port | 15 |
| 53 | extended-service-ingress-egress | 5 |
| 54 | complete-network-ing-egr | 15 |
| 56 | complete-pm | 5 |
| 61 | kpi-ref-pt-failure-cause-code | 5 |

=====

13.5 Measurement objects groups and indexes

Each measurement object, for example, the resource in the system that produces measurement values (<measValue> elements in the XML reports), is uniquely identified by the Measurement Information ID, the Group Name, and a set of indexes.

The measurement information ID is found in the measInfoId attribute of the container <measInfo> element. For example:

```
<measInfo measInfoId="KPISystemCP-ISA">
```

The Measurement Information ID identifies the counter module that generates the measurement values. The <measInfo> element is the container for multiple measurement values.

The group name is found inside the measObjLdn attribute of the <measValue> element. For example: <measValue measObjLdn="KPI=System,GroupName=CP-ISA,group=1,slot=3,mda=1">.

The group name identifies the system entity that contributes values to this counter module. Multiple groups may contribute to one counter module, and a group may contribute to more than one counter module.

Groups can be:

- VM types (CP-ISA, CPM)

- Reference point or interface type (S11, PFCP, Namf, Nchf, Npcf)

The set of indexes of a measurement value formulates the identifier of the resource or measurement object that generates the counter values. The indexes can be numeric or textual. The set of indexes used in a measurement value is specific to the counter module.

They can be found in the local distinguished name of the measurement object (measObjLdn) in the <measValue> element. For example:

```
<measValue measObjLdn="KPI=BearerManagement,GroupName=CP-ISA,group=1,slot=3,mda=1">
```

The set of indexes are located right after the GroupName in the measObjLdn attribute value.

Example set of indexes:

- slot=<value>
- group=<value>,slot=<value>,mda=<value>
- apn=<value>,slot=<value>,mda=<value>

See the *MAG-c KPI and KCI Counters Spreadsheet* for more information about the measurement information ID, group name, and indexes for each supported module.

Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)