

TR-459
Multi-Service Disaggregated BNG with CUPS.
Reference Architecture, Deployment Models,
interface, and Protocol Specifications

Issue: 2
Issue Date: April 2023

Notice

The Broadband Forum is a non-profit corporation organized to create guidelines for broadband network system development and deployment. This Technical Report has been approved by members of the Forum. This Technical Report is subject to change. This Technical Report is owned and copyrighted by the Broadband Forum, and all rights are reserved. Portions of this Technical Report may be owned and/or copyrighted by Broadband Forum members.

Intellectual Property

Recipients of this Technical Report are requested to submit, with their comments, notification of any relevant patent claims or other intellectual property rights of which they may be aware that might be infringed by any implementation of this Technical Report, or use of any software code normatively referenced in this Technical Report, and to provide supporting documentation.

Terms of Use

1. License

Broadband Forum hereby grants you the right, without charge, on a perpetual, non-exclusive and worldwide basis, to utilize the Technical Report for the purpose of developing, making, having made, using, marketing, importing, offering to sell or license, and selling or licensing, and to otherwise distribute, products complying with the Technical Report, in all cases subject to the conditions set forth in this notice and any relevant patent and other intellectual property rights of third parties (which may include members of Broadband Forum). This license grant does not include the right to sublicense, modify or create derivative works based upon the Technical Report except to the extent this Technical Report includes text implementable in computer code, in which case your right under this License to create and modify derivative works is limited to modifying and creating derivative works of such code. For the avoidance of doubt, except as qualified by the preceding sentence, products implementing this Technical Report are not deemed to be derivative works of the Technical Report.

2. NO WARRANTIES

THIS TECHNICAL REPORT IS BEING OFFERED WITHOUT ANY WARRANTY WHATSOEVER, AND IN PARTICULAR, ANY WARRANTY OF NONINFRINGEMENT AND ANY IMPLIED WARRANTIES ARE EXPRESSLY DISCLAIMED. ANY USE OF THIS TECHNICAL REPORT SHALL BE MADE ENTIRELY AT THE USER'S OR IMPLEMENTER'S OWN RISK, AND NEITHER THE BROADBAND FORUM, NOR ANY OF ITS MEMBERS OR SUBMITTERS, SHALL HAVE ANY LIABILITY WHATSOEVER TO ANY USER, IMPLEMENTER, OR THIRD PARTY FOR ANY DAMAGES OF ANY NATURE WHATSOEVER, DIRECTLY OR INDIRECTLY, ARISING FROM THE USE OF THIS TECHNICAL REPORT, INCLUDING BUT NOT LIMITED TO, ANY CONSEQUENTIAL, SPECIAL, PUNITIVE, INCIDENTAL, AND INDIRECT DAMAGES.

3. THIRD PARTY RIGHTS

Without limiting the generality of Section 2 above, BROADBAND FORUM ASSUMES NO RESPONSIBILITY TO COMPILE, CONFIRM, UPDATE OR MAKE PUBLIC ANY THIRD PARTY ASSERTIONS OF PATENT OR OTHER INTELLECTUAL PROPERTY RIGHTS THAT MIGHT NOW OR IN THE FUTURE BE INFRINGED BY AN IMPLEMENTATION OF THE TECHNICAL REPORT IN ITS CURRENT, OR IN ANY FUTURE FORM. IF ANY SUCH RIGHTS ARE DESCRIBED ON THE TECHNICAL REPORT, BROADBAND FORUM TAKES NO POSITION AS TO THE VALIDITY OR INVALIDITY OF SUCH ASSERTIONS, OR THAT ALL SUCH ASSERTIONS THAT HAVE OR MAY BE MADE ARE SO LISTED.

All copies of this Technical Report (or any portion hereof) must include the notices, legends, and other provisions set forth on this page.

Issue History

| Issue Number | Approval Date | Publication Date | Issue Editor | Changes |
|--------------|---------------|------------------|---|---|
| 1 | 9 June 2020 | 9 June 2020 | Kenneth Wan, Nokia | Original |
| 2 | 24 April 2023 | 24 April 2023 | Kenneth Wan, Nokia Nagaraj Turaiyur, Juniper Networks | Subscriber Group Resiliency Subscriber Group Prefix Assignment DHCP relay ACL solution BBF PFCP node reports and error notifications Call Flow updates |

Comments or questions about this Broadband Forum Technical Report should be directed to info@broadband-forum.org.

Editor(s): Kenneth Wan, Nokia
Nagaraj S Turaiyur, Juniper Networks

Work Area Director(s): David Sinicrope, Ericsson
Jonathan Newton, Vodafone

Project Stream Leader(s): Jonathan Newton, Vodafone

Table of Contents

| | |
|---|----|
| Executive Summary | 15 |
| 1 Purpose and Scope | 16 |
| 1.1 Purpose | 16 |
| 1.2 Scope | 16 |
| 2 References and Terminology..... | 17 |
| 2.1 Conventions..... | 17 |
| 2.2 References | 17 |
| 2.3 Definitions..... | 19 |
| 2.4 Abbreviations..... | 20 |
| 3 Technical Report Impact..... | 24 |
| 3.1 Energy Efficiency..... | 24 |
| 3.2 Security..... | 24 |
| 3.3 Privacy..... | 25 |
| 4 Introduction..... | 26 |
| 4.1 MS-BNG Functional Architecture | 27 |
| 4.1.1 MS-BNG Functions | 28 |
| 4.1.2 MS-BNG Interfaces | 30 |
| 4.2 DBNG Functional Architecture | 32 |
| 4.2.1 DBNG-CP Functions | 33 |
| 4.2.1.1 DBNG-CP Northbound Interfaces | 34 |
| 4.2.2 DBNG-UP Functions | 34 |
| 4.2.2.1 DBNG-UP Interfaces..... | 34 |
| 4.2.3 Interfaces between DBNG-CP and DBNG-UP..... | 35 |
| 4.2.3.1 Management Interface | 35 |
| 4.2.3.2 Control Packet Redirection Interface | 35 |
| 4.2.3.3 State Control Interface | 37 |
| 4.2.4 DBNG High level Architecture..... | 40 |
| 4.2.4.1 Traffic Steering Function | 40 |
| 4.2.4.2 Lawful Intercept Function..... | 42 |
| 4.3 Deployment models..... | 43 |
| 4.3.1 Deployment model: Geographical separation of DBNG-CP and DBNG-UP | 43 |
| 4.3.2 Deployment model: Non-Geographical separation of DBNG-CP and DBNG-UP..... | 44 |
| 4.4 Subscriber Session Resiliency and IP Prefix Management..... | 45 |

| | | |
|----------|--|-----|
| 4.4.1 | <i>Differences between Active and Backup Session State</i> | 46 |
| 4.4.2 | <i>Introduction to Subscriber Groups</i> | 46 |
| 4.4.3 | <i>Use of the Subscriber Group for resilience capability</i> | 47 |
| 4.4.4 | <i>Mapping Subscriber Groups to DBNG-UP</i> | 47 |
| 4.4.5 | <i>Switchover triggers</i> | 47 |
| 4.4.6 | <i>Virtual MAC Address</i> | 48 |
| 4.4.7 | <i>DBNG-UP independent actions</i> | 48 |
| 4.4.8 | <i>SGRP and Logical Port Relationship</i> | 49 |
| 4.4.9 | <i>CP Logical port resilience groups</i> | 49 |
| 4.4.10 | <i>Resilience Information attributed to the Subscriber Group</i> | 49 |
| 4.4.10.1 | <i>Partial State Properties</i> | 50 |
| 4.4.11 | <i>Use of Subscriber Groups for assignment and management of IP Prefixes</i> | 50 |
| 4.5 | Call Flows | 52 |
| 4.5.1 | <i>Control Plane and User Plane Association</i> | 52 |
| 4.5.2 | <i>Initial Control Packet Redirection Rule</i> | 53 |
| 4.5.2.1 | <i>Programming of Control Packet Redirection Interface for IpoE based subscriber sessions</i> | 54 |
| 4.5.2.2 | <i>Programming of the Control Packet Redirection Interface for PPPoE based subscriber sessions</i> | 56 |
| 4.5.3 | <i>External DHCP Server Control Packet Redirection Rule</i> | 58 |
| 4.5.4 | <i>IpoE DHCPv4 Immediate Session Creation</i> | 61 |
| 4.5.5 | <i>IpoE DHCPv4 Delayed Session Creation</i> | 62 |
| 4.5.6 | <i>IpoE DHCPv4 Relay Call Flows</i> | 63 |
| 4.5.6.1 | <i>IpoE DHCPv4 Relay Immediate Session Creation (via DBNG-UP)</i> | 63 |
| 4.5.6.2 | <i>IpoE DHCPv4 Relay Immediate Session Creation (via DBNG-CP)</i> | 65 |
| 4.5.6.3 | <i>IpoE DHCPv4 Relay Delayed Session Creation (via DBNG-UP)</i> | 67 |
| 4.5.6.4 | <i>IpoE DHCPv4 Relay Delayed Session Creation (via DBNG-CP)</i> | 69 |
| 4.5.7 | <i>IpoE DHCPv6 Immediate Session Creation</i> | 71 |
| 4.5.8 | <i>IpoE DHCPv6 Delayed Session Creation</i> | 72 |
| 4.5.9 | <i>IpoE DHCPv6 Relay Call Flows</i> | 74 |
| 4.5.9.1 | <i>IpoE DHCPv6 Relay Immediate Session Creation (via DBNG-UP)</i> | 74 |
| 4.5.9.2 | <i>IpoE DHCPv6 Relay Immediate Session Creation (via DBNG-CP)</i> | 76 |
| 4.5.9.3 | <i>IpoE DHCPv6 Relay Delayed Session Creation (via DBNG-UP)</i> | 78 |
| 4.5.9.4 | <i>IpoE DHCPv6 Relay Delayed Session Creation (via DBNG-CP)</i> | 80 |
| 4.5.10 | <i>IpoE SLAAC</i> | 82 |
| 4.5.11 | <i>IpoE Data Trigger</i> | 83 |
| 4.5.12 | <i>IpoE Dual Stack immediate session creation</i> | 84 |
| 4.5.13 | <i>IpoE Dual Stack delayed session creation</i> | 86 |
| 4.5.14 | <i>IpoE SLAAC and DHCPv6 PD immediate session creation</i> | 88 |
| 4.5.15 | <i>IpoE SLAAC and DHCPv6 PD delayed session creation</i> | 89 |
| 4.5.16 | <i>PPPoE immediate session creation</i> | 90 |
| 4.5.17 | <i>PPPoE delayed session creation</i> | 92 |
| 4.5.18 | <i>PPPoEv6 immediate session creation</i> | 94 |
| 4.5.19 | <i>PPPoEv6 delayed session creation</i> | 96 |
| 4.5.20 | <i>PPPoE Dual Stack immediate session creation</i> | 98 |
| 4.5.21 | <i>PPPoE Dual Stack delayed session creation</i> | 100 |
| 4.5.22 | <i>LAC immediate session creation</i> | 102 |
| 4.5.23 | <i>LAC delayed session creation</i> | 104 |
| 4.5.24 | <i>LNS – PPPoEv4 immediate session creation</i> | 106 |
| 4.5.25 | <i>LNS – PPPoEv4 delayed session creation</i> | 108 |
| 4.5.26 | <i>LNS – Dual Stack immediate session creation</i> | 109 |
| 4.5.27 | <i>LNS – Dual Stack delayed session creation</i> | 111 |
| 4.5.28 | <i>Public Wi-Fi Access</i> | 112 |
| 4.5.29 | <i>Public Wi-Fi Layer 3 Access</i> | 113 |
| 4.5.30 | <i>TWAG Call Flows</i> | 114 |
| 4.5.30.1 | <i>S2a initial attach based on layer 2 trigger: Ipv4 based on DHCPv4</i> | 115 |

| | | |
|----------|---|-----|
| 4.5.30.2 | S2a initial attach based on layer 2 trigger: Ipv6 prefix based on SLAAC | 117 |
| 4.5.30.3 | S2a initial attach based on layer 3 trigger: Ipv4 based on DHCPv4..... | 118 |
| 4.5.31 | Hybrid Access Gateway | 120 |
| 4.5.32 | Lawful Intercept call flows | 122 |
| 4.5.32.1 | Lawful Intercept while subscriber is online | 123 |
| 4.5.32.2 | Lawful Intercept while subscriber is offline | 124 |
| 4.5.32.3 | Lawful Intercept triggered by AAA authentication..... | 125 |
| 4.5.33 | Subscriber session modification..... | 126 |
| 4.5.34 | DHCPv4 Release | 128 |
| 4.5.35 | DHCPv4 Relay Release | 129 |
| 4.5.35.1 | DHCPv4 Relay Release (via DBNG-UP)..... | 129 |
| 4.5.35.2 | DHCPv4 Relay Release (via DBNG-CP)..... | 130 |
| 4.5.36 | DHCPv6 Release | 131 |
| 4.5.37 | DHCPv6 Relay Release | 132 |
| 4.5.37.1 | DHCPv6 Relay Release (via DBNG-UP)..... | 132 |
| 4.5.37.2 | DHCPv6 Relay Release (via DBNG-CP)..... | 134 |
| 4.5.38 | Lease and Lifetime timeouts | 135 |
| 4.5.39 | PPPoE termination due to LCP echo timeout..... | 136 |
| 4.5.40 | PPPoE Client initiated PPP LCP termination..... | 137 |
| 4.5.41 | PPPoE Server initiated PPP LCP termination..... | 138 |
| 4.5.42 | RG initiated PPP LCP termination for L2TP session | 139 |
| 4.5.43 | LAC initiated termination for L2TP session..... | 141 |
| 4.5.44 | LAC initiated termination for L2TP session on LNS..... | 143 |
| 4.5.45 | DBNG-CP initiated termination | 145 |
| 4.5.46 | DBNG subscriber statistics..... | 146 |
| 4.5.47 | Resiliency Call Flows | 151 |
| 4.5.47.1 | Establishment of a resilient session..... | 151 |
| 4.5.47.2 | DBNG-CP Managed Resilience Switchover | 153 |
| 4.5.47.3 | Establishment of a resilient session with track-logical-port | 155 |
| 4.5.47.4 | Independent DBNG-UP Switchover..... | 157 |
| 5 | Technical Requirements | 159 |
| 5.1 | State Control Interface requirements..... | 159 |
| 5.2 | Requirements to support Control packet redirect interface | 162 |
| 5.3 | Management Interface requirements..... | 164 |
| 5.4 | Disaggregated MS-BNG control plane requirements | 164 |
| 5.5 | Disaggregated MS-BNG user plane requirements | 166 |
| 5.6 | Disaggregated MS-BNG functional requirements | 167 |
| 5.7 | Requirements for 3GPP PFCP node messages | 168 |
| 5.8 | Requirements for 3GPP PFCP session messages | 169 |
| 6 | PFCP CUPS protocol | 170 |
| 6.1 | PFCP messages..... | 170 |
| 6.1.1 | PFCP node messages | 170 |
| 6.1.2 | PFCP session messages..... | 174 |

| | | |
|---------|--|-----|
| 6.1.3 | <i>PFCP information elements</i> | 180 |
| 6.2 | PDR Matching Rules | 185 |
| 6.2.1 | <i>Per-session, per-logical-port, and default CPRi matching precedence</i> | 185 |
| 6.2.2 | <i>ARP and NS rules</i> | 186 |
| 6.2.3 | <i>Unicast Traffic for Local DBNG-UP IP addresses</i> | 186 |
| 6.2.4 | <i>Priority Between Network Redirect and Per-session Rules</i> | 187 |
| 6.3 | PFCP Connectivity Requirements | 187 |
| 6.4 | General PFCP information exchanges for a subscriber session | 189 |
| 6.4.1 | <i>General PFCP rules for control packet redirection</i> | 190 |
| 6.4.2 | <i>General PFCP rules for data packet forwarding</i> | 190 |
| 6.4.3 | <i>General PFCP rules for Server Control Packet Redirection</i> | 191 |
| 6.4.4 | <i>General Information PFCP Filter IEs</i> | 193 |
| 6.4.5 | <i>General information on NSH header</i> | 193 |
| 6.4.5.1 | <i>GTP Tunnel Endpoint ID</i> | 193 |
| 6.4.5.2 | <i>NSH header insertion option on CPRi</i> | 194 |
| 6.5 | PFCP use case and information exchanges | 194 |
| 6.5.1 | <i>Overview of default CPR and per logical port CPR tunnel</i> | 194 |
| 6.5.1.1 | <i>Default Control Packet Redirect Rule(s)</i> | 194 |
| 6.5.1.2 | <i>Per logical port Control Packet Redirect Rule(s)</i> | 195 |
| 6.5.1.3 | <i>Use case: Default control packet redirection for immediate session Creation</i> | 195 |
| 6.5.1.4 | <i>Use case: Default control packet redirection for delayed session creation</i> | 196 |
| 6.5.2 | <i>Use case: IPoE</i> | 197 |
| 6.5.2.1 | <i>PFCP Control Packet redirection rule</i> | 197 |
| 6.5.2.2 | <i>PFCP Data Packet Forwarding rule</i> | 197 |
| 6.5.3 | <i>Use case: PPPoE</i> | 197 |
| 6.5.3.1 | <i>PFCP Control Packet redirection rule</i> | 198 |
| 6.5.3.2 | <i>PFCP Data Packet Forwarding rule</i> | 198 |
| 6.5.4 | <i>Use Case: L2TP LAC</i> | 198 |
| 6.5.4.1 | <i>PFCP session for L2TP tunnel setup</i> | 199 |
| 6.5.4.2 | <i>PFCP update for L2TP session</i> | 199 |
| 6.5.5 | <i>Use Case: L2TP LNS</i> | 199 |
| 6.5.5.1 | <i>PFCP session for L2TP tunnel setup</i> | 199 |
| 6.5.5.2 | <i>PFCP Control Packet redirection rule</i> | 200 |
| 6.5.5.3 | <i>PFCP Data Packet Forwarding rule</i> | 200 |
| 6.5.6 | <i>Use Case: TWAG</i> | 200 |
| 6.5.6.1 | <i>DBNG-UP TEID assignment (optional)</i> | 200 |
| 6.5.6.2 | <i>PFCP Control Packet redirection rule</i> | 201 |
| 6.5.6.3 | <i>PFCP Data Packet Forwarding rule</i> | 201 |
| 6.5.7 | <i>Subscriber Group and Subscriber Prefix Use Case</i> | 201 |
| 6.6 | BBF PFCP Information Element Summary | 203 |
| 6.7 | PFCP Grouped IE extensions for Node Related Messages | 217 |
| 6.7.1 | <i>PFCP Association Setup Request</i> | 217 |
| 6.7.2 | <i>PFCP Association Setup Response</i> | 217 |
| 6.7.3 | <i>PFCP Association Update Request</i> | 217 |
| 6.7.3.1 | <i>BBF-node-info create</i> | 218 |
| 6.7.3.2 | <i>BBF-node-info modify</i> | 219 |
| 6.7.3.3 | <i>BBF-node-info delete</i> | 219 |
| 6.7.4 | <i>PFCP Association Update Response</i> | 219 |
| 6.7.5 | <i>PFCP Node Report Request</i> | 220 |

| | | |
|---------|---|-----|
| 6.7.5.1 | 3GPP Node Report Type | 221 |
| 6.7.5.2 | BBF Logical Port Report IE | 221 |
| 6.7.5.3 | BBF SGRP Notification Report IE | 222 |
| 6.7.5.4 | BBF Network Instance Report IE | 224 |
| 6.7.6 | BBF SGRP | 225 |
| 6.7.6.1 | SGRP Modify | 227 |
| 6.7.6.2 | SGRP Delete | 227 |
| 6.7.7 | BBF UP subscriber prefix | 228 |
| 6.7.7.1 | BBF UP Subscriber Prefix Modify | 229 |
| 6.7.7.2 | BBF UP Subscriber Prefix Deletion | 229 |
| 6.8 | PFCP Grouped IE extensions for Session Related Messages | 230 |
| 6.8.1 | PFCP Session Establishment Request | 230 |
| 6.8.1.1 | Create PDR | 230 |
| 6.8.1.2 | PDI | 230 |
| 6.8.1.3 | Ethernet Packet Filter | 231 |
| 6.8.1.4 | Forwarding Parameters | 232 |
| 6.8.1.5 | Create Traffic Endpoint | 232 |
| 6.8.2 | PFCP Session Establishment Response | 233 |
| 6.8.3 | PFCP Session Modification Request | 233 |
| 6.8.3.1 | Update Forwarding Parameters | 234 |
| 6.8.3.2 | Ethernet Packet Filter | 234 |
| 6.8.4 | PFCP Session Modification Response | 235 |
| 6.8.5 | PPP LCP Connectivity | 235 |
| 6.8.6 | L2TP Tunnel | 235 |
| 6.8.7 | BBF ACL IE | 236 |
| 6.9 | BBF PFCP IE extensions | 237 |
| 6.9.1 | BBF UP Function Features | 237 |
| 6.9.2 | Logical Port | 239 |
| 6.9.3 | BBF Outer Header Creation | 239 |
| 6.9.3.1 | NSH header information | 241 |
| 6.9.4 | BBF Outer Header Removal | 241 |
| 6.9.5 | PPPoE Session ID | 242 |
| 6.9.6 | PPP Protocol | 242 |
| 6.9.7 | Verification Timers | 243 |
| 6.9.8 | PPP LCP Magic Number | 243 |
| 6.9.9 | MTU | 244 |
| 6.9.10 | L2TP Tunnel Endpoint | 244 |
| 6.9.11 | L2TP Session ID | 245 |
| 6.9.12 | L2TP type | 246 |
| 6.9.13 | BBF Direction IE | 246 |
| 6.9.14 | BBF Family IE | 246 |
| 6.9.15 | BBF SGRP Identifier | 247 |
| 6.9.16 | BBF SGRP State | 247 |
| 6.9.17 | BBF SGRP Flags | 247 |
| 6.9.18 | BBF Operational Condition IE | 248 |
| 6.9.19 | BBF IPv4 Prefix | 248 |
| 6.9.20 | BBF IPv6 Prefix | 249 |
| 6.9.21 | BBF Prefix Tag | 249 |
| 6.9.22 | BBF Error Code | 250 |
| 6.9.23 | BBF Error Message | 251 |
| 6.9.24 | BBF Maximum ACL Chain Length | 251 |
| 6.9.25 | BBF Forwarding Capability | 251 |
| 6.9.26 | BBF Connectivity Status | 251 |

| | | |
|-------------|---|-----|
| 6.9.27 | <i>Vendor-Specific Node Report Type</i> | 252 |
| 6.9.28 | <i>BBF C-Tag Range</i> | 252 |
| 6.9.29 | <i>BBF S-Tag Range</i> | 253 |
| 6.9.30 | <i>BBF Apply Action</i> | 253 |
| 6.10 | General DBNG-CP and DBNG-UP interoperability recommendations..... | 254 |
| Annex A: | Use Cases | 255 |
| A.1 | Multi-access MS-BNG CUPS use case..... | 255 |
| A.2 | Wireline-Access disaggregated MS-BNG use case | 256 |
| A.3 | Data-trigger service use case..... | 256 |
| A.4 | Wholesale Retail model with L2TP use case | 256 |
| A.5 | IPoE DHCP and DHCPv6 Relay Service Model..... | 257 |
| A.5.1 | <i>DHCPv4 Relay</i> | 258 |
| A.5.2 | <i>DHCPv6 Relay</i> | 259 |
| A.6 | Subscriber Session ACL and ACL Chaining | 259 |
| Appendix I. | Alternative Call Flows for PPPoE delayed session creation..... | 260 |
| I.1 | PPPoE delayed session creation | 260 |
| I.2 | PPPoEv6 delayed session creation..... | 263 |
| I.3 | PPPoE Dual Stack delayed session creation | 265 |
| I.4 | LAC delayed session creation | 267 |
| I.5 | LNS – PPPoEv4 delayed session creation..... | 268 |
| I.6 | LNS - Dual Stack delayed session creation | 270 |

Table of Figures

| | |
|--|-----|
| Figure 1: MS-BNG Functional blocks and interfaces | 27 |
| Figure 2: MS-BNG functions separating into Control Plane and User Plane | 32 |
| Figure 3: Management Interface | 35 |
| Figure 4: Example of Control and User Plane control message exchange..... | 36 |
| Figure 5: Control Packet Redirect Interface | 36 |
| Figure 6: Example of Control Plane pushing forwarding rules to the User Plane | 37 |
| Figure 7: State Control Interface | 37 |
| Figure 8: Example of User Plane combinations | 38 |
| Figure 9: State Control Interface for Access to Network direction..... | 38 |
| Figure 10: State Control Interface for Network to Access direction..... | 39 |
| Figure 11: DBNG Lawful Intercept Function..... | 42 |
| Figure 12: Geographically distributed deployment model | 43 |
| Figure 13: Non-Geographically distributed deployment model | 44 |
| Figure 14: DBNG subscriber session resilience | 45 |
| Figure 15: Example SGRP to DBNG UP mapping..... | 47 |
| Figure 16: DBNG-CP and DBNG-UP association | 52 |
| Figure 17: Programming of the Control Packet Redirection interface for IpoE based subscriber sessions..... | 54 |
| Figure 18: Programming of the Control Packet Redirection Interface for PPPoE based subscriber sessions | 56 |
| Figure 19: External DHCP server Control Packet Redirection rule | 59 |
| Figure 20: IpoE DHCPv4 immediate session creation call flow | 61 |
| Figure 21: IpoE DHCPv4 delayed session creation call flow | 62 |
| Figure 22: IpoE DHCPv4 Relay Immediate Session Creation (via DBNG-UP) call flow..... | 63 |
| Figure 23: IpoE DHCPv4 Relay Immediate Session Creation (via DBNG-CP) call flow..... | 65 |
| Figure 24: IpoE DHCPv4 Relay Delayed Session Creation (via DBNG-UP) call flow | 67 |
| Figure 25: IpoE DHCPv4 Relay Delayed Session Creation (via DBNG-CP) call flow | 69 |
| Figure 26: IpoE DHCPv6 immediate session creation call flow | 71 |
| Figure 27: IpoE DHCPv6 delayed session creation call flow | 72 |
| Figure 28: IpoE DHCPv6 Relay Immediate Session Creation (via DBNG-UP) call flow..... | 75 |
| Figure 29: IpoE DHCPv6 Relay Immediate Session Creation (via DBNG-CP) call flow..... | 76 |
| Figure 30: IpoE DHCPv6 Relay Delayed Session Creation (via DBNG-UP) call flow | 78 |
| Figure 31: IpoE DHCPv6 Relay Delayed Session Creation (via DBNG-CP) call flow | 80 |
| Figure 32: IpoE SLAAC call flow | 82 |
| Figure 33: IpoE Data Trigger call flow | 83 |
| Figure 34: IpoE Dual Stack immediate session creation call flow | 84 |
| Figure 35: IpoE Dual Stack delayed session creation call flow | 87 |
| Figure 36: IpoE SLAAC and DHCPv6 PD immediate session creation call flow | 88 |
| Figure 37: IpoE SLAAC and DHCPv6 PD delayed session creation call flow | 89 |
| Figure 38: PPPoE immediate session creation call flow | 90 |
| Figure 39: PPPoE delayed session creation call flow | 92 |
| Figure 40: PPPoEv6 immediate session creation call flow | 94 |
| Figure 41: PPPoEv6 delayed session creation call flow | 96 |
| Figure 42: PPPoE Dual Stack immediate session creation call flow..... | 98 |
| Figure 43: PPPoE Dual Stack delayed session creation call flow..... | 100 |
| Figure 44: LAC immediate session creation call flow | 102 |
| Figure 45: LAC Delayed Session Creation call flow | 104 |
| Figure 46: LNS PPPoEv4 immediate session creation call flow | 106 |
| Figure 47: LNS Ipv4 Delayed Session Creation call flow | 108 |
| Figure 48: LNS Dual Stack immediate session creation call flow | 109 |
| Figure 49: LNS Dual Stack Delayed Session Creation call flow | 111 |
| Figure 50: Public Wi-Fi Access call flow | 112 |
| Figure 51: Public Wi-Fi Layer 3 Access call flow..... | 113 |
| Figure 52: S2a initial attached based on layer 2 trigger: DHCPv4..... | 115 |
| Figure 53: S2a initial attached based on layer 2 trigger: SLAAC | 117 |

| | |
|--|-----|
| Figure 54: S2a initial attached based on layer 3 trigger: DHCPv4..... | 118 |
| Figure 55: Hybrid Access Gateway L3 network-based Tunneling call flow..... | 120 |
| Figure 56: Example of Lawful intercept Request..... | 123 |
| Figure 57: Lawful intercept while subscriber is offline..... | 124 |
| Figure 58: Lawful intercept triggered by AAA authentication..... | 125 |
| Figure 59: Subscriber session modification call flow..... | 126 |
| Figure 60: DHCPv4 Release Call flow..... | 128 |
| Figure 61: IpoE DHCPv4 Relay Release Call Flow (via DBNG-UP)..... | 129 |
| Figure 62: IpoE DHCPv4 Relay Release Call Flow (via DBNG-CP)..... | 130 |
| Figure 63: DHCPv6 Release Call Flow..... | 131 |
| Figure 64: IpoE DHCPv6 Relay Release Call Flow (via DBNG-UP)..... | 132 |
| Figure 65: IpoE DHCPv6 Relay Release Call Flow (via DBNG-CP)..... | 134 |
| Figure 66: Lease and Lifetime timeout..... | 135 |
| Figure 67: PPPoE termination due to LCP echo timeout..... | 136 |
| Figure 68: PPPoE client initiated PPP LCP termination..... | 137 |
| Figure 69: PPPoE Server initiated PPP LCP termination..... | 138 |
| Figure 70: RG initiated PPP LCP termination for L2TP session..... | 139 |
| Figure 71: LAC initiated termination for L2TP session..... | 141 |
| Figure 72: LAC initiated termination for L2TP session on LNS..... | 143 |
| Figure 73: DBNG-CP initiated termination..... | 145 |
| Figure 74: DBNG subscriber statistics for IpoE..... | 147 |
| Figure 75: DBNG Subscriber Statistics for PPPoE..... | 149 |
| Figure 76: Establishment of a resilient session call flow..... | 151 |
| Figure 77: DBNG-CP Managed Resilience Switchover call flow..... | 153 |
| Figure 78: Establishment of a resilient session call flow with Track-Logical-Port..... | 155 |
| Figure 79: Independent DBNG-UP Switchover call flow..... | 157 |
| Figure 80: Subscriber Group and Subscriber Prefix..... | 201 |
| Figure 81: 3GPP Node Report Type..... | 221 |
| Figure 82: 3GPP Vendor-Specific Information Element Format Reference..... | 237 |
| Figure 83: BBF UP Function Features..... | 237 |
| Figure 84: Logical Port..... | 239 |
| Figure 85: BBF Outer Header Creation..... | 240 |
| Figure 86: NSH header information..... | 241 |
| Figure 87: BBF Outer Header Removal..... | 241 |
| Figure 88: PPPoE Session ID..... | 242 |
| Figure 89: PPP Protocol..... | 243 |
| Figure 90: Verification Timers..... | 243 |
| Figure 91: PPP LCP Magic Number..... | 244 |
| Figure 92: MTU..... | 244 |
| Figure 93: L2TP Tunnel Endpoint..... | 245 |
| Figure 94: L2TP Session ID..... | 245 |
| Figure 95: L2TP type..... | 246 |
| Figure 96: BBF Direction IE..... | 246 |
| Figure 97: BBF Family IE..... | 246 |
| Figure 98: BBF SGRP Identifier..... | 247 |
| Figure 99: BBF SGRP State..... | 247 |
| Figure 100: BBF SGRP Flags..... | 248 |
| Figure 101: BBF Operational Condition..... | 248 |
| Figure 102: BBF IPv4 Prefix..... | 249 |
| Figure 103: BBF IPv6 Prefix..... | 249 |
| Figure 104: BBF Prefix tag..... | 250 |
| Figure 105: BBF Error Code..... | 250 |
| Figure 106: BBF Error Message..... | 251 |
| Figure 107: Maximum ACL Chain Length..... | 251 |
| Figure 108: BBF Forwarding Capability..... | 251 |

| | |
|--|-----|
| Figure 109: BBF Connectivity Status..... | 252 |
| Figure 110: Vendor-Specific Node Report Type | 252 |
| Figure 111: C-Tag Range | 252 |
| Figure 112: S-Tag Range | 253 |
| Figure 113: BBF Apply action Protocol..... | 253 |
| Figure 114: Multi-access MS-BNG | 255 |
| Figure 115: Multi-access DBNG | 256 |
| Figure 116: L2TP deployment model | 257 |
| Figure 117: DBNG-UP Connected External DHCP Server Deployment Model..... | 258 |
| Figure 118: ACL chaining | 260 |
| Figure 119: PPPoE call flow | 261 |
| Figure 120: PPPoEv6 call flow | 263 |
| Figure 121: Delayed PPPoE Dual Stack call flow | 265 |
| Figure 122: LAC call flow – Delayed Session Creation..... | 267 |
| Figure 123: LNS IPv4 call flow Delayed Session Creation..... | 269 |
| Figure 124: LNS Dual Stack call flow Delayed Session Creation | 272 |

Table of Tables

| | |
|--|-----|
| Table 1: Functional Blocks of a MS-BNG | 28 |
| Table 2: MS-BNG access interfaces | 30 |
| Table 3: MS-BNG network interfaces | 31 |
| Table 4: MS-BNG control and management interfaces..... | 31 |
| Table 5: Examples of traffic detection and traffic forwarding rules..... | 160 |
| Table 6: Mandatory 3GPP PFCP Node messages for TR-459..... | 170 |
| Table 7: 3GPP PFCP Node Messages IEs Applicable to TR-459 | 171 |
| Table 8: Mandatory 3GPP PFCP Session messages for TR-459..... | 174 |
| Table 9: 3GPP PFCP Session Establishment Messages IEs Applicable to TR-459 | 175 |
| Table 10: 3GPP PFCP Session Establishment Response Messages IEs Applicable to TR-459..... | 176 |
| Table 11: 3GPP PFCP Session Modification Request Messages IEs Applicable to TR-459 | 177 |
| Table 12: 3GPP PFCP Session Modification Response Messages IEs Applicable to TR-459 | 179 |
| Table 13: 3GPP PFCP Session Deletion Request Messages IEs Applicable to TR-459 | 179 |
| Table 14: 3GPP PFCP Session Deletion Response Messages IEs Applicable to TR-459 | 179 |
| Table 15: 3GPP PFCP Session Report Request IEs Applicable to TR-459 | 180 |
| Table 16: 3GPP PFCP Session Report Response IEs Applicable to TR-459 | 180 |
| Table 17: 3GPP PFCP IEs that must be supported for TR-459 | 181 |
| Table 18: 3GPP PFCP IE flags that must be supported for TR-459 | 183 |
| Table 19: 3GPP PFCP IE that must be supported for TR-459 and are conditional based on the 3GPP CP function features flag..... | 184 |
| Table 20: Example of a PDR for Control Packet Redirection from DBNG-UP to DBNG-CP | 190 |
| Table 21: Example of a PDR for Control Packet redirection from DBNG-CP to DBNG-UP | 190 |
| Table 22: Example of a PDR for upstream data packet forwarding through the DBNG-UP | 191 |
| Table 23: Example of a PDR for downstream data packet forwarding through the DBNG-UP | 191 |
| Table 24: PDR for DHCPv4 Server Control Packet Redirection from DBNG-UP to DBNG-CP | 192 |
| Table 25: PDR for DHCPv6 Server Control Packet Redirection from DBNG-UP to DBNG-CP | 192 |
| Table 26: PDR for DHCP Server Control Packet Redirection from DBNG-CP to DBNG-UP | 192 |
| Table 27: Example of a PDR for Control Packet Redirection from DBNG-UP to DBNG-CP | 196 |
| Table 28: Example of a PDR for Common Control Packet redirection from DBNG-CP to DBNG-UP..... | 197 |
| Table 29: BBF extended Information Elements for Node Association, Control Packet Redirect Sessions ... | 204 |
| Table 30: BBF extended Information Elements for IPoE/ TWAG..... | 206 |
| Table 31: BBF extended Information Elements for PPPoE..... | 208 |
| Table 32: BBF extended Information Elements for L2TP LAC..... | 210 |
| Table 33: BBF extended Information Elements for L2TP LNS..... | 212 |
| Table 34: BBF UP Function Features and required BBF extended PFCP IEs | 213 |
| Table 35: 3GPP PFCP Information Element Types and applicability | 214 |
| Table 36: BBF extended Information Element(s) in a PFCP Association Setup Request | 217 |
| Table 37: BBF extended Information Element(s) in a PFCP Association Setup Response | 217 |
| Table 38: BBF extended Information Element(s) in a PFCP Association Update Request | 218 |
| Table 39: Information Elements in an BBF-node-info-create | 218 |
| Table 40: Information Elements in an BBF-node-info-modify..... | 219 |
| Table 41: Information Elements in an BBF-node-info-delete | 219 |
| Table 42: BBF extended Information Element(s) in a PFCP Association Update Response | 220 |
| Table 43: Information Elements in PFCP Node Report Request | 220 |
| Table 44: Information Elements in BBF Logical Port Report..... | 221 |
| Table 45: Information Elements in BBF SGRP Notification Report..... | 222 |
| Table 46: Information Elements in BBF SGRP Error | 223 |
| Table 47: Information Elements in BBF Prefix Error | 223 |
| Table 48: Information Elements in BBF Network Instance Report..... | 224 |
| Table 49: BBF SGRP..... | 225 |
| Table 50: BBF UP Subscriber Prefix | 228 |
| Table 51: BBF extended Information Element(s) in a PFCP Session Establishment Request..... | 230 |

| | |
|---|-----|
| Table 52: BBF extended Create PDR IE(s) within PFCP Session Establishment Request..... | 230 |
| Table 53: BBF extended PDI IE within PFCP Session Establishment Request | 231 |
| Table 54: BBF extended Ethernet Packet Filter IE(s) within PFCP Session Establishment Request | 231 |
| Table 55: BBF extended Forwarding Parameters IE in FAR..... | 232 |
| Table 56: BBF extended Create Traffic Endpoint IE(s) within PFCP Session Establishment Request..... | 233 |
| Table 57: BBF extended Information Element(s) in a PFCP Session Establishment Response..... | 233 |
| Table 58: BBF extended Information Element(s) in a PFCP Session Modification Request | 234 |
| Table 59: BBF extended Update Forwarding Parameters IE(s) in FAR..... | 234 |
| Table 60: BBF extended Ethernet Packet Filter IE(s) within PFCP Session Modification Request..... | 234 |
| Table 61: PFCP Session Modification Response..... | 235 |
| Table 62: PPP LCP Connectivity | 235 |
| Table 63: L2TP Tunnel | 235 |
| Table 64: BBF ACL..... | 236 |
| Table 65: BBF UP Function Features..... | 238 |
| Table 66: BBF Outer Header Creation Description | 240 |
| Table 67: BBF Outer Header Removal Description | 242 |

Executive Summary

This Technical Report specifies the architecture and requirements for a Disaggregated Broadband Network Gateway (DBNG). The separation of the control plane and user plane in the DBNG enables more efficient use of resources and simplifies Operations. TR-459 Issue 1 [23] standardizes the DBNG CUPS architecture to support traditional broadband MS-BNG use cases. TR-459 Issue 2 further enriches the DBNG CUPS specification by defining DBNG-UP resiliency and dynamic programming of subscriber prefixes from DBNG-CP to multiple DBNG-UPs, as well as a number of optional features such as ACL programming and more flexibility in programming common and dedicated control packet redirection interfaces.

To allow multi-vendor interoperability for all current and future use case, the 3GPP Packet Forwarding Control Protocol (PFCP) is specified as the State Control Interface (SCI) protocol for programming subscriber forwarding state between the control plane and user plane in TR-459 Issue 1 [23]. PFCP is an extensible protocol, it has been extended to support traditional broadband use cases in TR-459 Issue 1 [23] and now further extended to support new use cases defined in TR-459 Issue 2.

1 Purpose and Scope

1.1 Purpose

This document specifies the architecture and requirements for a control plane and user plane separation of a Multi-Service Disaggregated Broadband Network Gateway (MS-DBNG). The architecture designates Broadband Network Gateway (BNG) functions to either the control plane or user plane and defines the interfaces between the control plane and user plane. Requirements on both interfaces and protocols help ensure interoperability between different vendors' control planes and user planes. Use cases and deployment models are also captured. Although BNG control plane and user plane are separated, the goal is to ensure traditional broadband service offerings are maintained. In addition, new capabilities can be realized through control plane and user plane separation such as independent control plane and user plane scaling, independent control and user plane life cycle management, and simplifying operations by centralized control plane for configuration.

1.2 Scope

In scope for TR-459 Issue 2:

- Fix errors found in TR-459 [23]
- Update to TR-459 [23] section 4.5 specifically:
 - o Call flows related to session modification such as RADIUS Change of Authorization (CoA).
 - o Call flows related to session termination such as network triggered termination and abrupt session termination
- Update to TR-459 [23] section 6 specifically:
 - o Material related to PFCP IEs required for session modification or termination. The PFCP IEs can be new or existing.
 - o Material related to session statistics reporting and applicable PFCP IEs.
 - o Additional material should be specifically on:
 - PFCP Node message, focusing on:
 - Node initialization and PFCP association.
 - PFCP IE relevant for TR-459 [23] PFCP association
 - Reference table(s) for list of 3GPP IEs and BBF IEs mandatory for PFCP association (use cases).
 - Reference table(s) for list of 3GPP IEs and BBF IEs mandatory to each UP-feature flag (use cases).
 - List the 3GPP IEs which are applicable to TR-459 [23] different use case that would be associated to new or existing BBF specific feature flag.
 - o Clarification that additional 3GPP IEs beyond the one specified in TR-459 [23] are up to the vendor for implementation and experiment.
 - o Clarifications on specific 3GPP PFCP IEs which are applicable to wireline use case.
 - o Interoperability section
 - o Further details on procedures that can assist PFCP implementation with respect to TR-459 [23]
 - o Rearranging existing material for better reading consumption.
- If there is consensus to revisit material, then new/existing section(s) would be updated.

Not in scope:

- Section 4.5 specifically:
 - o Redefining existing call flows will not be modified or changed.
- Section 6 specifically:
 - o Redefining any existing 3GPP PFCP procedures, 3GPP PFCP IEs, BBF defined PFCP IEs.
 - o Modification of any existing PFCP IEs or the list of IEs to be included per use case (for backward compatibility purpose).

2 References and Terminology

2.1 Conventions

In this Technical Report, several words are used to signify the requirements of the specification. These words are always capitalized. More information can be found in RFC 2119 [33].

| | |
|------------|---|
| MUST | This word, or the term “REQUIRED”, means that the definition is an absolute requirement of the specification. |
| MUST NOT | This phrase means that the definition is an absolute prohibition of the specification. |
| SHOULD | This word, or the term “RECOMMENDED”, means that there could exist valid reasons in particular circumstances to ignore this item, but the full implications need to be understood and carefully weighed before choosing a different course. |
| SHOULD NOT | This phrase, or the phrase “NOT RECOMMENDED” means that there could exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications need to be understood and the case carefully weighed before implementing any behavior described with this label. |
| MAY | This word, or the term “OPTIONAL”, means that this item is one of an allowed set of alternatives. An implementation that does not include this option MUST be prepared to inter-operate with another implementation that does include the option. |

2.2 References

The following references are of relevance to this Technical Report. At the time of publication, the editions indicated were valid. All references are subject to revision; users of this Technical Report are therefore encouraged to investigate the possibility of applying the most recent edition of the references listed below.

A list of currently valid Broadband Forum Technical Reports is published at www.broadband-forum.org.

| Document | Title | Source | Year |
|-----------------------------|--|--------|------|
| [1] TR-25 | Core Network Architecture for Access to Legacy Data Network over ADSL | BBF | 1999 |
| [2] TR-59 | DSL Evolution - Architecture Requirements for the Support of QoS-Enabled IP Services | BBF | 2003 |
| [3] TR-69 Amendment 6 | CPE WAN Management Protocol | BBF | 2018 |
| [4] TR-101 Issue 2 | Migration to Ethernet-Based Broadband Aggregation | BBF | 2011 |
| [5] TR-134 Corrigendum 1 | Broadband Policy Control Framework (BPCF) | BBF | 2013 |
| [6] TR-145 | Multi-service Broadband Network Functional Modules and Architecture | BBF | 2012 |
| [7] TR-146 | Subscriber Sessions | BBF | 2013 |
| [8] TR-147 | Layer 2 Control Mechanism for Broadband Multi-Service Architectures | BBF | 2008 |
| [9] TR-177 Corrigendum 1 | IPv6 in the Context of TR-101 | BBF | 2017 |

| | | | |
|-------------------------------|---|------|-----------|
| [10]TR-178 Issue 2 | Multi-service Broadband Network Architecture and Nodal Requirements | BBF | 2017 |
| [11]TR-187 Issue 2 | IPv6 for PPP Broadband Access | BBF | 2013 |
| [12]TR-291 | Nodal Requirements for Interworking between Next Generation Fixed and 3GPP Wireless Access | BBF | 2014 |
| [13]TR-300 | Policy Convergence for Next Generation Fixed and 3GPP Wireless Networks | BBF | 2014 |
| [14]TR-321 | Public Wi-Fi Access in Multi-service Broadband Networks | BBF | 2015 |
| [15]TR-341 | Radius Attributes Catalog | BBF | 2016 |
| [16]TR-348 | Hybrid Access Broadband Network Architecture | BBF | 2016 |
| [17]TR-378 | Nodal Requirements for Hybrid Access Broadband Networks | BBF | 2019 |
| [18]TR-384 | Cloud Central Office (CloudCO) Reference Architectural Framework | BBF | 2018 |
| [19]TR-145 | Multi-service Broadband Network Functional Modules and Architecture | BBF | 2012 |
| [20]TR-459.2 | Multi-Service Disaggregated BNG with CUPS: Integrated Carrier Grade NAT function. Reference Architecture, Deployment Models, Interface, and Protocol Specifications | BBF | Oct 2021 |
| [21]TR-459.3 | Multi-Service Disaggregated BNG with CUPS: IPTV Multicast function - Reference Architecture, Deployment Models, Interface and Protocol Specifications | BBF | Aug 2021 |
| [22]MR-459.2 | Improving Service Resilience through BNG Disaggregation | BBF | Mar 2020 |
| [23]TR-459 | Control and User Plane Separation for a disaggregated BNG | BBF | June 2020 |
| [24]3GPP TS 23.007 | Restoration procedures | 3GPP | 2022 |
| [25]3GPP TS 23.214 | Architecture enhancements for control and user plane separation of EPC nodes | 3GPP | 2019 |
| [26]3GPP TS 23.401 | General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access | 3GPP | Dec 2019 |
| [27]3GPP TS 23.402 | Architecture enhancements for non-3GPP accesses | 3GPP | 2019 |
| [28]3GPP TS 29.244 | Interface between the Control Plane and the User Plane nodes | 3GPP | Dec 2019 |
| [29]3GPP TS 29.274 | 3GPP Evolved Packet System (EPS); Evolved General Packet Radio Service (GPRS) Tunneling Protocol for Control plane (GTPv2-C); Stage 3 | 3GPP | Dec 2019 |
| [30]3GPP TS 33.402 | 3GPP System Architecture Evolution (SAE); Security aspects of non-3GPP accesses | 3GPP | Jun 2018 |
| [31]RFC 791 | INTERNET PROTOCOL | IETF | 1981 |
| [32]RFC 1661 | The Point-to-Point Protocol (PPP) | IETF | 1994 |
| [33] RFC 2119 | Key words for use in RFCs to Indicate Requirement Levels | IETF | 1997 |
| [34]RFC 2516 | A Method for Transmitting PPP Over Ethernet (PPPoE) | IETF | 1999 |
| [35]RFC 2661 | Layer Two Tunneling Protocol "L2TP" | IETF | 1999 |
| [36]RFC 2865 | Remote Authentication Dial In User Service (RADIUS) | IETF | 2000 |
| [37]RFC 2866 | RADIUS Accounting | IETF | 2000 |
| [38]RFC 4301 | Security Architecture for the Internet Protocol | IETF | 2005 |
| [39]RFC 4364 | BGP/MPLS IP Virtual Private Networks (VPNs) | IETF | 2006 |
| [40]RFC 4381 | Analysis of the Security of BGP/MPLS IP Virtual Private | IETF | 2006 |

| | | | |
|---------------|---|------|------|
| | Networks (VPNs) | | |
| [41] RFC 4443 | Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification | IETF | 2006 |
| [42] RFC 5515 | Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions | IETF | 2009 |
| [43] RFC 5880 | Bidirectional Forwarding Detection (BFD) | IETF | 2010 |
| [44] RFC 5881 | Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop) | IETF | 2010 |
| [45] RFC 6973 | Privacy Considerations for Internet Protocols | IETF | 2013 |
| [46] RFC 7432 | BGP MPLS-Based Ethernet VPN | IETF | 2015 |
| [47] RFC 8300 | Network Service Header (NSH) | IETF | 2018 |

2.3 Definitions

The following terminology is used throughout this Technical Report.

| | |
|------------------------|--|
| AAA Client function | A logical entity that sends authenticating, authorizing and accounting requests to an AAA Server function. An example of an AAA client function contained with a network node is a Network Access Server (NAS) as described in RFCs 2865 [36] and 2866 [37]. Examples of AAA client function in BBF TRs are the BRAS of TR-059 [2] and the BNG of TR-101 [4]. (TR-134 [5]) |
| AAA Server | A physical device that contains an AAA Server functions. (TR-134 [5]) |
| AAA Server Function | A logical entity in the client-server relationship that replies to AAA Client Authentication, Authorization and Accounting requests. The AAA server function is typically responsible for receiving user connection requests, authenticating the user, and replying to the AAA Client function with an Accept or Deny response. The AAA Server function can return, as part of this reply, some or all of the configuration information necessary for the AAA client function to deliver service to the user. An AAA server function can act as a proxy client to other AAA server functions or other kinds of authentication servers. The AAA Server function does not contain any business logic other than basic authentication. (TR-134 [5]) |
| ACL | An ACL is an ordered set of rules used to filter traffic in a device. Each rule is used to find a match on a packet and define actions that will be performed on the packet. |
| ACL chaining | Ordered list of two or more ACLs that are applied to subscriber session |
| Cold Standby | In general terms, a Cold Standby is a computing resource available in a DBNG with access to executable software and current configuration information. |
| C-Tag | Customer Tag |
| DBNG | Disaggregated BNG where the subscriber management function is separated between control plane and user plane. |
| Framed Route | This Attribute provides routing information to be configured for the user on the NAS. It is used in the Access-Accept packet and can appear multiple times. (RFC 2865 [36]) |
| HAG | Hybrid Access Gateway. A logical function in the operator network implementing the network side mechanisms for simultaneous use of both fixed broadband and 3GPP access networks. (TR-378 [17]) |
| Hot Standby | The protecting node is assigned to a working entity as its backup. The secondary node is configured, activated, and it maintains full operational state information so as to promptly take over the duties of the primary. To do so, it receives near real-time protocol information about that state. When a failure of the working instance is detected, the protecting node is able to work as a participant in all protocols. The restoration time is generally faster than warm standby. |

| | |
|---|--|
| Logical Port | An opaque name that represents an entity on which Ethernet traffic is received from and forwarded to subscribers on the DBNG. Examples of Logical ports that carry Ethernet traffic include physical links, Link Aggregation Group interfaces, MPLS Pseudowires or virtual LAN interfaces over physical links based on IEEE 802.1q, IEEE 802.1ad, IEEE 802.3ad Link Aggregation Group bundle and MPLS Pseudowires. |
| MS-BNG | TR-178 introduces the Multi-Service BNG (MS-BNG), which extends the capabilities of a traditional BNG to offer services to both residential and business customers as well as to allow mobile backhaul deployments. To achieve this, it performs Ethernet Aggregation and can either forward packets via MPLS or through IP Aggregation/routing. A MS-BNG is part of a TR-145 network architecture and can be deployed in a hierarchical BNG architecture. (TR-178 [10]) |
| Network Instance | Network Instance identifies an IP forwarding and routing context on a DBNG-UP. A Network Instance may include one or more address-families (IPv4, IPv6, labeled unicast, etc) and is an abstraction of the exact technology used: for example, it may identify an IPv4 or an IPv6 forwarding table (i.e., the default forwarding table or a VRF as defined in RFC 4364 [39]), or an IPsec VPN as defined in RFC 4301 [38], or a DNS name for APN. |
| Preferred Resilience Status Protecting Instance | A generic term for the DBNG to indicate the required status to the Access Network. An instance that has been assigned to a particular working instance or a group of working instances. |
| S-Tag | Service Tag |
| TWAG | The trusted WLAN access gateway (TWAG) is the logical entity responsible for the 3GPP UE IP mobility service on the data plane between a Trusted BBF Access and 3GPP network. |
| Warm Standby | The protecting node is assigned to a working entity as its backup. The secondary node is configured, receives configuration updates, but it is not actively engaged in any protocol. Instead, the state information may be sent from the primary instance from time to time. This restoration time is generally faster than cold standby. |
| WLAN | Wireless Local Area Network. |
| Working Instance | An instance of DBNG, whether DBNG-CP or DBNG-UP, that maintains both the current configuration and operational state. |

2.4 Abbreviations

This Technical Report uses the following abbreviations:

| | |
|------|--|
| 3GPP | 3 rd Generation Partnership Project |
| AAA | Authentication, Authorization & Accounting |
| AC | Access Controller |
| ACL | Access Control List |
| AN | Access Node |
| AP | Access Point |
| APN | Access Point Name |
| BBF | Broadband Forum |
| BFD | Bidirectional Forwarding Detection |
| BGP | Border Gateway Protocol |
| BNG | Broadband Network Gateway |
| CO | Central Office |
| CP | Control Plane |
| CPE | Customer Premises Equipment |
| CPR | Control Packet Redirect |

| | |
|--------|---|
| CUPS | Control and User Plane Separation |
| DBNG | Disaggregated BNG |
| DF | Delivery Function |
| DHCP | Dynamic Host Configuration Protocol |
| DHCPv6 | Dynamic Host Configuration Protocol version 6 |
| DSL | Digital Subscriber Line |
| DSLAM | Digital Subscriber Line Access Multiplexer |
| EAP | Extensible Authentication Protocol |
| EMS | Element Management System |
| EPC | Enhanced Packet Core |
| ETH | Ethernet |
| FAR | Forward Action Rule |
| FFS | For Further Study |
| F-TEID | Fully Qualified Tunnel Endpoint Identifier |
| GRE | Generic Routing Encapsulation |
| GTP | GPRS Tunnelling Protocol |
| GTP-c | GTP for control plane |
| GTP-u | GTP user plane |
| GW | Gateway |
| HAG | Hybrid Access Gateway |
| HCPE | Hybrid CPE |
| HSS | Home Subscriber Server |
| ICCN | Incoming Call Connected |
| ICMP | Internet Control Message Protocol |
| ICMPv6 | Internet Control Message Protocol version 6 |
| ICRP | Incoming Call Reply |
| ICRQ | Incoming Call Request |
| IE | Information Element |
| IEEE | Institute of Electrical and Electronics Engineers |
| IETF | Internet Engineering Task Force |
| IGMP | Internet Group Management Protocol |
| IGP | Interior Gateway Protocol |
| IMS | IP Multimedia Subsystem |
| IMSI | International Mobile Subscriber Identity |
| IP | Internet Protocol |
| IPCP | IP Control Protocol |
| IpoE | IP over Ethernet |
| IpoEv4 | IpoE version 4 |
| IpoEv6 | IpoE version 6 |
| IPTV | IP Television |
| Ipv6 | Internet Protocol version 6 |
| Ipv6CP | Ipv6 Control Protocol |
| JSON | JavaScript Object Notation |
| KA | Keep-Alive |
| L2 | Layer 2 |
| L2TP | Layer 2 Tunneling Protocol |
| L2TS | Layer 2 Tunneling switching |
| L2VPN | Layer 2 VPN |
| L3 | Layer 3 |

| | |
|----------|--|
| LAC | L2TP Access Concentrator |
| LCP | Link Control Protocol |
| LEA | Law Enforcement Agency |
| LI | Lawful Intercept |
| LNS | L2TP Network Server |
| LTE | Long Term Evolution |
| MAC | Media Access Control |
| MANO | Management and Orchestration |
| Mi | Management Interface |
| MD | Mediation Device |
| MLD | Multicast Listener Discovery |
| MME | Mobile Management Entity |
| MPLS | Multi-Protocol Label Switching |
| MRU | Maximum Receive Unit |
| MS-BNG | Multi-Service BNG |
| MTU | Maximum Transfer Unit |
| NAI | Network Access Identifier |
| NCP | Network Control Protocol |
| ND | Neighbor Discovery |
| NETCONF | Network Configuration Protocol |
| OAM | Operations Administration and Maintenance |
| OLT | Optical Line Terminal |
| OTT | Over the Top |
| PDI | Packet Detection Information |
| PDN | Packet Data Network |
| PDP | Policy Decision Point |
| PDR | Packet Detection Rule |
| PEP | Policy Enforcement Point |
| PFCP | Packet Forwarding Control Protocol |
| PGW | PDN GW |
| PIM | Protocol Independent Multicast |
| PPP | Point to Point Protocol |
| PPPoE | PPP over Ethernet |
| PPPoEv4 | PPPoE version 4 |
| PPPoEv6 | PPPoE version 6 |
| PTA | PPP Termination and Aggregation |
| QoS | Quality of Service |
| RA | Router Advertisement |
| RADIUS | Remote Authentication Dial In User Service |
| RESTCONF | Representational State Transfer Configuration Protocol |
| RG | Residential Gateway |
| SCCCN | Start Control Connection Connected |
| SCCRP | Start Control Connection Reply |
| SCCRQ | Start Control Connection Request |
| SCi | State Control interface |
| SGRP | Subscriber Group |
| SGW | Serving GW |
| SLA | Service Level Agreement |
| SLAAC | Ipv6 Stateless Address Autoconfiguration |
| SNMP | Simple Network Management Protocol |

| | |
|---------|-------------------------------|
| SSS | Subscriber Session steering |
| TEID | Traffic Endpoint Identifier |
| TLV | Type Length Value |
| TR | BBF Technical Report |
| TWAG | Trusted WLAN Access Gateway |
| TWAN | Trusted WLAN Access Network |
| TWAP | Trusted WLAN AAA Proxy |
| UP | User Plane |
| UPSF | User Plane Selection Function |
| VLAN | Virtual Local Area Network |
| VLL | Virtual Leased Line |
| VoIP | Voice over IP |
| VPLS | Virtual Private LAN Service |
| VPN | Virtual Private Network |
| VSA | Vendor-Specific Attribute |
| Wi-Fi | Wireless Fidelity |
| WLAN | Wireless Local Area Network |
| WLAN-GW | WLAN Gateway |
| WT | BBF Working Text |
| XML | Extensible Markup Language |
| YANG | Yet Another Next Generation |
| ZLB | Zero Length Body |

3 Technical Report Impact

3.1 Energy Efficiency

Energy Efficiency may be impacted migrating from standalone MS-BNG to DBNG. DBNG allows dynamic scaling of both the control and user plane according to the subscriber population which may lead to optimized equipment deployment and therefore, energy consumption. However, energy consumption can also increase depending on deployment factors such as power per node, geo dispersion of user plane, and use of generic hardware not optimized for specific network functions.

Regulatory differences related to electrical power, Heating, Ventilation and Air Conditioning (HVAC) and fire protection between traditional Central Offices and datacenters are out-of-scope for this document.

3.2 Security

The DBNG is subject to the security concerns applicable to the MS-BNG, and particularly to those functions and interfaces included as “in scope” in the “Scope” section (1.2) of this document.

Security concerns for these functions and interfaces are described in the following TRs:

- TR-101 [4] “Migration to Ethernet-Based Broadband Aggregation” (L2 Security Considerations, section 3.7);
- TR-134 [5] “Broadband Policy Control Framework (BPCF)” (Security, section 3.3);
- TR-145 [6] “Multi-service Broadband Network Functional Modules and Architecture” (Security considerations for converged, multi-service networks, section 4.6.2);
- TR-146 [7] “Subscriber Sessions” (Security, section 3.3);
- TR-177 [9] “Ipv6 in the context of TR-101” (Security, section 3.3; Ipv6 Security Considerations, section 4.8; and L2 Security Considerations, section 5.5);
- TR-178(i2) [10] “Multi-service Broadband Network Architecture and Nodal Requirements” (Security, sections 3.3 and 5.4.8; and Security Requirements, section 5.6.3.6);
- TR-187 [11] “Ipv6 for PPP Broadband Access (Security, section 3.3);
- TR-291 [12] “Nodal Requirements for Interworking between Next Generation and 3GPP Wireless Access” (Security, section 3.3);
- TR-321 [14] “Public Wi-Fi Access in Multi-service Broadband Networks” (section 3.3);
- TR-384 [18] “Cloud Central Office Reference Architectural Framework” (section 3.3);

Security concerns for relevant technologies are documented in several additional Standards. See, for example:

- Std. IEEE 802.1-2014 –
 - clause 8 (“Principles of Bridge Operations” – which includes some description of MAC layer security mechanisms),
 - clause 17.4 (“Security Considerations” relating to Bridge management),
 - clause 27.20 (“Security considerations” relating to Shortest Path Bridging);
- “Security Considerations” section of RFC 4364 [39] – “BGP/MPLS VPNs;”
- RFC 4381 [40] – “Analysis of the Security of BGP/MPLS IP Virtual Private Networks (VPNs);”
- “Security Considerations” section of RFC 7432 [46] – “BGP MPLS-Based Ethernet VPN.”

Where not explicitly described in directly relevant standards, implementations of DBNG should include a description of security considerations, possibly expanding on related referenceable standards.

In addition, any implementation of DBNG functionality should include documentation of special requirements that apply specifically to that implementation – such as measures needed to ensure that only authorized

parties can access, or modify, configuration for underlying network infrastructure and that traffic associated with one subscriber cannot be intercepted by other subscribers.

Finally, because of the distribution of control and data plane functions defined for DBNG, the CUPS protocols must include capabilities for providing secure and authenticated communication between distributed components of the DBNG, specifically for the in-scope communication between CP and UP components, as indicated in Requirements [R-11] and [R-12].

Operators should consider using these capabilities as an important means for preventing attacks intended (for example) to divert or disable data forwarding capabilities through control plane impersonation.

3.3 Privacy

The DBNG is subject to the privacy concerns applicable to the MS-BNG, and particularly to those functions and interfaces included as “in scope” in the “Scope” section 1.2 of this document.

The privacy of connection specific information in transit between components of a DBNG is provided through encryption of the data exchange.

In network protocols, privacy concerns, beyond the protection of potentially private data, focus on two aspects:

- 1) The potential for tracking of users through exposure of Personal Identifying Information (PII);
- 2) The potential for correlation of user activity over time through persistent use of network identifiers.

Privacy concerns for generic MS-BNG functions and interfaces are described in the following TRs:

- TR-134 [5] “Broadband Policy Control Framework (BPCF)” (Privacy, section 3.4);
- TR-291 [12] “Nodal Requirements for Interworking between Next Generation and 3GPP Wireless Access” (Privacy, section 3.4);
- TR-384 [18] “Cloud Central Office Reference Architectural Framework” (Privacy, section 3.4)

Because of its distributed nature, the same (or highly correlated) identifying information may be seen at several points in the network, allowing for identification of a target subscriber or DBNG component. This increases the potential exposure to privacy violations.

In addition to security considerations described in section 3.3, DBNG implementers should include information as to what privacy protection is provided in the implementation, (e.g., – avoiding direct or inferable relationships between subscriber PII and network identifiers, avoiding persistent use of identifiers during different stages of subscriber activation, use and deactivation, minimizing the extent to which PII is included in the protocol, or stored at DBNG components, etc.) and explicitly include details of any unavoidable (or required) use and/or storage of PII.

DBNG component implementations should include privacy considerations such as those listed in related standards and similar activities such as:

- IEEE 802 current work to document recommended practices for creating new standards, as well as suggesting what to consider in implementing and deploying network technologies. This work may be published as early as sometime during the year 2019.
- IETF publication (in July, of 2013) of an Information RFC (RFC 6973 [45]) – entitled “Privacy Considerations for Internet Protocols” – that includes some of the history relating to privacy considerations, and suggests “legally generic” (e.g., – recognizing that the definitions and handling of privacy differ across legal jurisdictions) guidance that can be used as “food for thought” in designing network protocols independent of specific legal framework(s).

4 Introduction

The MS-BNG is an essential device that grants subscribers access to the internet and other private networks. It provides critical subscriber management functions, such as: authentication, IP address assignment, bandwidth allocation, and accounting. The MS-BNG also terminates various access types including: fixed wireline, fixed wireless, public Wi-Fi, and hybrid access.

The broadband services that MS-BNG supports include: VoIP, IPTV multicast, OTT video streaming, video game streaming, business VPN services, and many other IP services. As a result, the MS-BNG is taking on exponential subscriber growth as well as increased bandwidth demand which brings forth the following challenges:

- Over-utilizing a MS-BNG
- Under-utilizing a MS-BNG
- Managing and maintaining geographically distributed MS-BNG deployments
- Service provisioning across all deployed MS-BNGs
- Time to market services

The DBNG tackles these challenges by separating the subscriber management CP and UP. DBNG allows independent scaling of the CP and the UP to keep up with subscriber growth and subscriber bandwidth demands.

The CUPS architecture creates the opportunity to simplify operations by maintaining a single management interface on the CP to manage all Ups and provides the chance of interoperable multi-vendor architecture. This approach requires the standardized data model for the Mi interface to be defined, and is therefore for future study.

This document provides the architecture, requirements, and call flows of a DBNG. In addition, DBNG utilizes 3GPP defined Packet Forwarding Control Protocol (PFCP) for programming subscriber forwarding state on the UP from the CP. Section 6 provides the PFCP information element exchanges for typical MS-BNG use cases and PFCP Information Element extensions required to support wireline use cases.

4.1 MS-BNG Functional Architecture

Figure 1 illustrates the set of MS-BNG functions and interfaces that have been defined in BBF Technical Reports (TRs). Operators utilize a combination of MS-BNG functions to provide different types of broadband service(s). It should be noted that certain interfaces are not required depending on the selected MS-BNG functions (please refer to respective TRs for more detail on interface dependency). The MS-BNG consists of access, network, control, and management interfaces. The control and management interfaces of an MS-BNG are commonly known as north bound interfaces. The access and network interfaces are user plane interfaces.

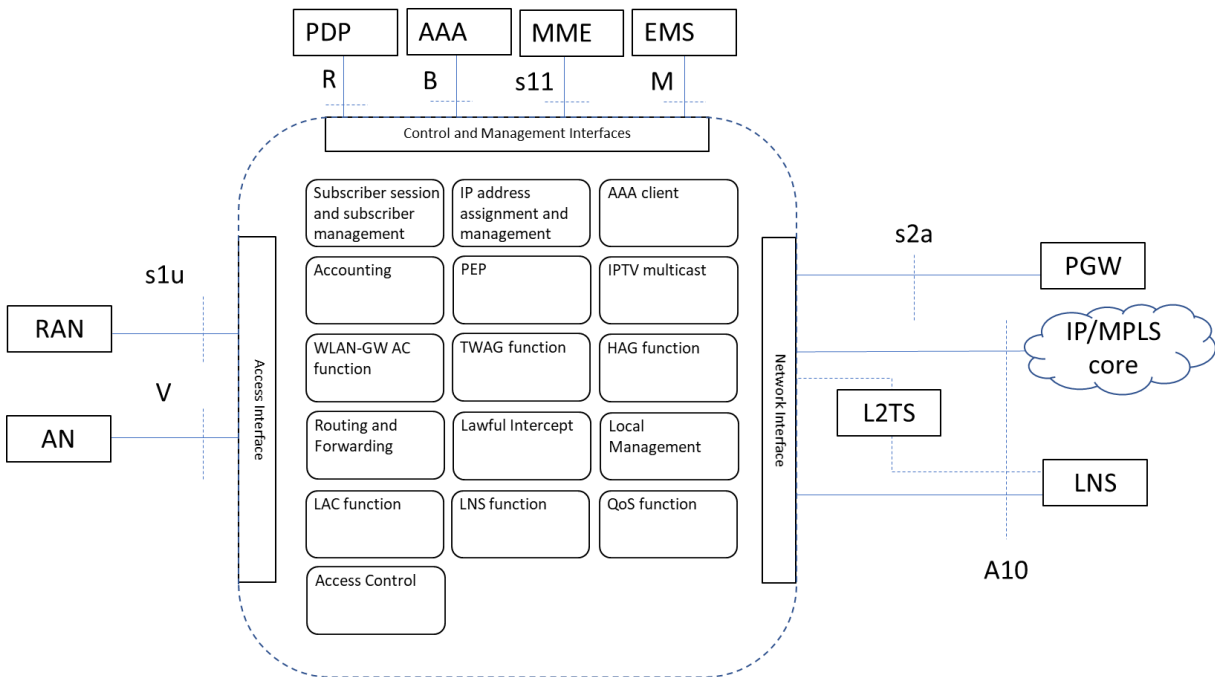


Figure 1: MS-BNG Functional blocks and interfaces

4.1.1 MS-BNG Functions

The MS-BNG utilizes different functional blocks to provide subscriber management functions including: AAA authentication, IP address assignment, policy enforcement, and accounting. In addition, the MS-BNG integrates other functional blocks such as Trusted Wi-Fi Access Gateway (TWAG) and Hybrid Access Gateway (HAG) to provide enhanced broadband services. Table 1 lists the functional blocks within the BNG:

Table 1: Functional Blocks of a MS-BNG

| MS-BNG Functions | BBF TR references | MS-BNG functions specified in TRs |
|--|--|---|
| Subscriber session and subscriber management | TR-145, TR-25 | IpoE, PPPoE session, and L2TP LAC session |
| IP address assignment and management | TR-101, TR-177, TR-178, TR-187, TR-341 | IP address/prefix assignment through DHCPv4, DHCPv6, SLAAC, PPPoE, DHCPv6 over PPPoE, and SLAAC over PPPoE. Note that, for DHCPv4 and DHCPv6, address assignment may be performed by an external DHCP/DHCPv6 server. |
| AAA client | TR-101, TR-177, TR-178, TR-187, TR-134, TR-300, TR-341 | Subscriber authentication and authorization |
| Accounting | TR-101, TR-177, TR-178, TR-187, TR-134, TR-300, TR-341 | Subscriber accounting |
| Policy Enforcement Point (PEP) | TR-101, TR-177, TR-178, TR-187, TR-134, TR-300, TR-341 | Subscriber filters and QoS rules |
| IPTV multicast | TR-101, TR-177, TR-178, TR-187, TR-134, TR-300 | Subscriber IGMP/MLD processing |
| WLAN-GW Access Controller (AC) | TR-321 | AC for Access Point (AP) management |
| TWAG | TR-291 | Hand off broadband access to 3GPP |
| HAG | TR-348 | Hybrid access |
| Routing and Forwarding | TR-101, TR-177, TR-178, TR-187 | IGP, BGP, MPLS, PIM |
| Lawful Intercept (LI) | TR-178 | Mirroring |
| Local Management | | Local Management |
| L2TP Access Concentrator (LAC) | TR-25, TR-187 | PPP based wholesale retail function |
| L2TP Network Server (LNS) | TR-25, TR-187 | PPP based wholesale retail function |
| QoS | TR-101, TR-177, TR-178, TR-187, TR-134, TR-300, TR-341 | QoS enforcement |
| Access Control | TR-147 | Access Control can be part of an overall policy |

| | | |
|--|--|---|
| | | management framework. This includes aspects of QoS control, conditional access, as described in e.g., section 6.2 |
|--|--|---|

4.1.2 MS-BNG Interfaces

The following interfaces are defined in various BBF TRs:

- V-interface: Defined in TR-101 [4] as the Ethernet interface between the Access Node and the MS-BNG. Further, it includes the following capabilities: traffic aggregation, class of service distinction, and user isolation and traceability
- A10-interface: Defined in TR-25 [1] as the interface between the Regional Broadband network and the network service provider Point of Presences (POPs).
- R-interface: Defined in TR-134 [5] as the interface between the Policy Enforcement Point (PEP) and the Policy Decision Point (PDP). The PEP is a function integrated into the MS-BNG.
- B-interface: Defined in TR-134 [5] as the interface between the PEP and the AAA server. The PEP is a function integrated in the MS-BNG, where it can receive/activate/modify/delete policies from AAA server.
- s1u: Defined in TR-378 [17] as the interface between the enhanced Node B (eNodeB) and the HAG as per 3GPP TS23.002
- s11: Defined in TR-378 [17] as the interface between the 3GPP Mobility Management Entity (MME) and the HAG.
- s2a: Defined in TR-291 [12] as the reference point between the TWAG and the 3GPP PDN GW. It is used for interworking between a Trusted BBF Access and 3GPP network and for supporting IP Network-Based Mobility. It conveys mobility and policy control from the 3GPP domain towards the TWAG.
- M-interface: Defined in TR-145 [6] and TR-178 [10] as the management interface between the MS-BNG and the Operator’s EMS, which is used to monitor, control, analyze, and manage the node.

The access interfaces on the MS-BNG terminates various access types such as broadband and fixed mobile connections. Table 2 specifies the MS-BNG access interfaces cross referenced to relevant TRs and its respective protocol stacks.

Table 2: MS-BNG access interfaces

| Interfaces | BBF TR references | Protocol Stacks | TR defined functions with respect to the access interfaces |
|------------|-------------------|-------------------------------|---|
| V | TR-25, TR-101 | IpoEv4, PPPoEv4 | MS-BNG terminates PPPoE and IpoEv4 sessions. For LAC, the MS-BNG will terminate PPPoE sessions. |
| V | TR-177 | IpoEv6 | Based on TR-101, the MS-BNG terminates IpoEv6 sessions |
| V | TR-178 | PPPoEv4, IpoEv4, IpoMPLS | Based on TR-101, the MS-BNG terminates IpoMPLS sessions |
| V | TR-187 | PPPoEv6 | Based on TR-178, the MS-BNG terminates PPPoEv6 sessions |
| V | TR-291 | IpoE, L2oIPGRE | Based on TR-178, the MS-BNG integrates TWAG functions and terminates IpoE sessions |
| V | TR-321 | IpoE, L2oIPGRE | Based on TR-178, the MS-BNG integrates access controller functions for public Wi-Fi and terminates IpoE sessions |
| V and s1u | TR-378 | GTP (s1u), IpoE and PPPoE (V) | Based on TR-178, the MS-BNG integrates Hybrid Access Gateway functions and terminates PPPoE, GTP, and IpoE sessions |

The MS-BNG connects subscriber IpoE or PPPoE session to the network core with a network interface. As highlighted in Figure 1, the network interface also provides connectivity to wholesale service via L2TP or via MPLS. In addition, broadband service can also be offered via the 3GPP core network. Table 3 specifies the MS-BNG network interface cross referenced to relevant BBF TRs and their protocol stacks.

Table 3: MS-BNG network interfaces

| Interfaces | BBF TR references | Protocol stacks | TR defined functions with respect to the network interfaces |
|-------------------|--------------------------|------------------------|---|
| A10 | TR-25 | L2TP | L2TP handoff to LTS or LNS |
| A10 | TR-178 | IpoEv4, MPLS | MS-BNG interfaces with the network core through IP/MPLS |
| A10 | TR-187 | IpoEv6 | MS-BNG interfaces with the network core through Ipv6 |
| S2a | TR-291 | GTP | MS-BNG that integrated TWAG function and interfaces with the 3GPP PGW |

MS-BNG manages subscribers through control and management message exchanges with external elements. Typically, these include AAA server, policy server, accounting servers, EMS or a DHCP / DHCPv6 server. Table 4 specifies the MS-BNG control and management interfaces cross referenced to relevant BBF TRs and their protocol stacks.

Table 4: MS-BNG control and management interfaces

| Interfaces | TR references | Protocol stacks | TR defined functions with respect to the control and management interfaces |
|-------------------|----------------------|---------------------------------|---|
| R | TR-134, TR-300 | RADIUS, Diameter | Policy control framework |
| B | TR-134 | RADIUS, Diameter | AAA service |
| S11 | TR-378 | GTP-c | Control interface between 3GPP MME and MS-BNG for session creation and session management |
| M | TR-145, TR-178 | NETCONF, SSH, Telnet, SNMP, ... | Management |

4.2 DBNG Functional Architecture

Utilizing the baseline MS-BNG in Figure 1, the MS-BNG interfaces and functional blocks are separated between the control plane and user plane. The functional architecture of the DBNG is shown in Figure 2.

Please note:

- The Traffic Steering Function in Figure 2 is described in detail in section 4.2.4.1.
- Interfaces between the DBNG control plane (CP) and DBNG user plane in Figure 2 are described in detail in section 4.2.3.

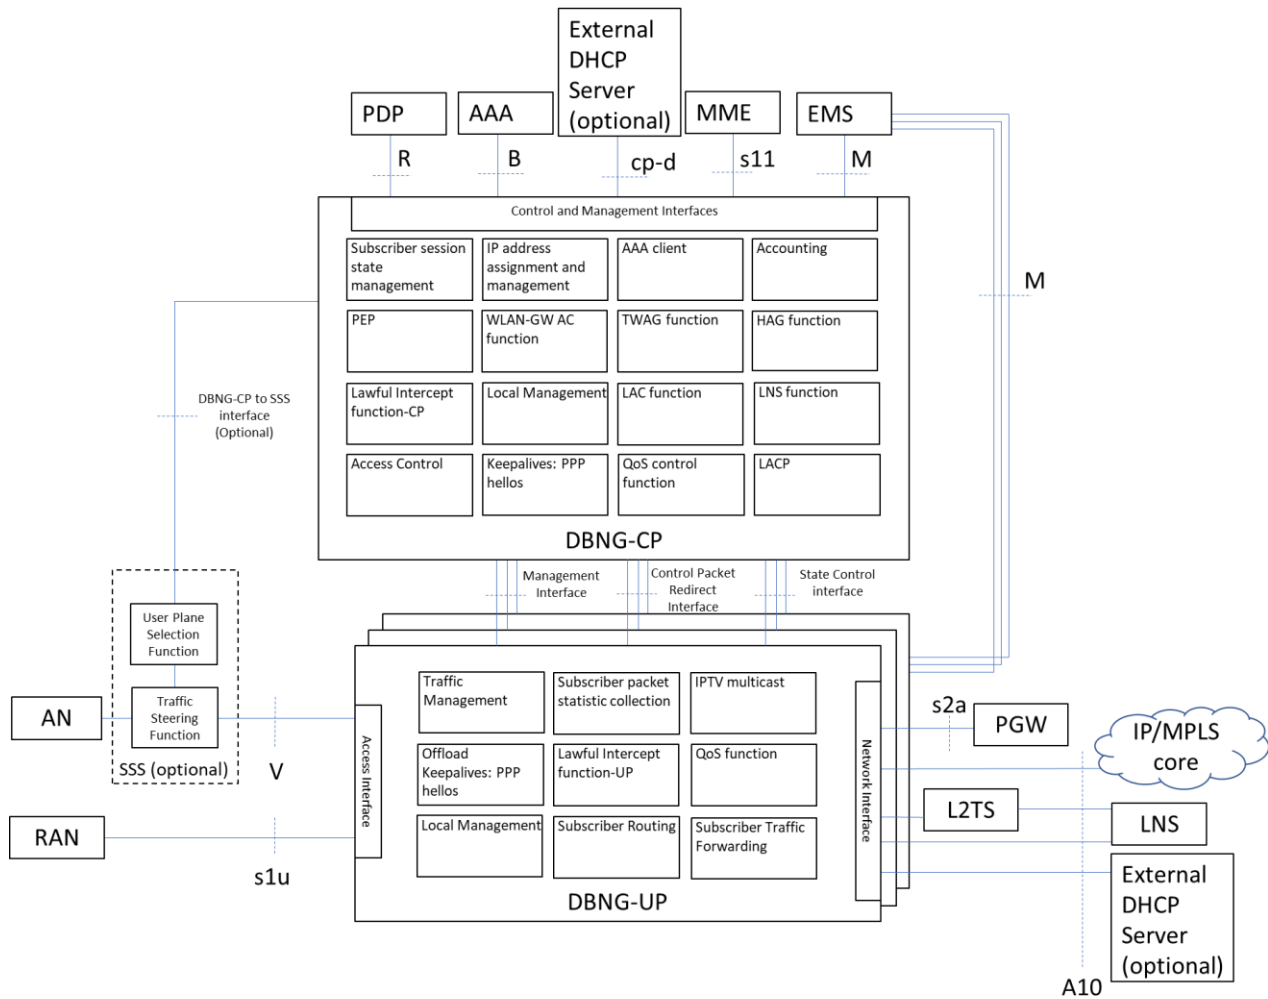


Figure 2: MS-BNG functions separating into Control Plane and User Plane

A combination of CP functions is referred to as a control plane of the DBNG (DBNG-CP). Similarly, a combination of UP specific functions is referred to as a user plane of the DBNG (DBNG-UP). As specified in scope section 1.2, this document only focuses on routing function (both routing control and data forwarding) residing on the DBNG-UP. The DBNG-UP is responsible to route all subscriber data traffic. The DBNG-CP is able to support routing capabilities to interact with external systems.

4.2.1 DBNG-CP Functions

The DBNG-CP performs functions such as authentication via AAA servers and assigns IP address(es) to the subscriber. After the DBNG establishes a subscriber session, accounting updates are sent periodically to an accounting server. The following functions residing in the DBNG-CP include:

- Subscriber session state management
 - Managing PPPoE, lpoE, and L2TP session state.
- IP address assignment and management
 - Utilizing a PPPoE and an lpoE session to assign lpv4 address, lpv6 address, and/or lpv6 prefix are control plane related functions. For lpoE the address may be assigned by an external DHCP or DHCPv6 server. After the address assignment has completed the control plane will signal forwarding rules to the user plane
- Accounting
 - Subscriber usage information is exchanged with accounting servers
- AAA client
 - Control message exchange with AAA servers
 - Authorize IP data-trigger subscribers
 - PPPoE, DHCPv4, DHCPv6, and L2TP messages will trigger authentication with the AAA server
- PEP
 - After receiving the policies from a Policy Decision Point (PDP), the control plane will push these policies onto the DBNG-UP for enforcement.
- Subscriber Group Management
 - Grouping subscribers together for the purpose of prefix assignment and subscriber state resiliency
- Control Protocol Processing
 - Example such as PPPoE, DHCP, DHCPv6, and L2TP

In addition, TR-291 [12], TR-321 [14], and TR-384 [16] integrate further functions into the DBNG-CP.

- WLAN-GW: As explained in TR-321 [14], the access controller is integrated into the MS-BNG. Prior to address assignment, the end device performs EAP authentication with the AC. After a successful authentication, the MS-BNG will assign an IP address for the end device. The authentication of the AC and address assignment of the MS-BNG are all control plane processing. After the address assignment is completed the control plane will signal forwarding rules to the user plane.
- TWAG: As explained in TR-291 [12], TWAG performs wireline authentication with an lpoE subscriber. After a successful authentication, the TWAG function will signal the PGW to create a GTP tunnel. After the tunnel setup has completed, the address assignment will then be completed by the MS-BNG. The MS-BNG authentication/addressing function and the TWAG signaling are all control plane processing. After the address assignment has completed, the control plane will signal forwarding rules to the user plane.
- HAG: As explained in TR-348 [16], the hybrid access consists of a broadband connection and a 3GPP connection. The broadband connection is provided through traditional MS-BNG functions. The 3GPP connection follows 3GPP procedures where the eNodeB establishes tunnels to the HAG using procedure defined in 3GPP TS 29.274 [29]. Further address allocation for the 3GPP connections follows the PDP procedures. After the address assignment is completed the control plane will signal forwarding rules to the user plane for both the broadband and 3GPP connection.

Other DBNG-CP functions include:

- Lawful Intercept Control Plane Function: DBNG-CP receives instruction from a lawful intercept mediation element and instructs the DBNG-UP lawful intercept actions.
- Management Interface: DBNG-CP has a management interface that allows management of the DBNG-CP and many of its DBNG-Ups. The management of DBNG-UP via DBNG-CP has not been specified in this document.

- **Keepalives:** The DBNG-CP must be able to process and generate subscriber session PPP keep-alive messages (including the generation of and the response to LCP Echo-Requests and the processing of LCP Echo-Reply messages). However, the DBNG-CP should offload the PPP LCP Echo message processing and generation to the DBNG-UP, which is recommended when trying to achieve high scalability. The Keepalive function is offloaded to the DBNG-UP on a per subscriber basis and hence the DBNG-CP must still be able to process the LCP echo-Replies until the DBNG-UP takes over the offload function for a subscriber.
- **DBNG-UP Selection:** with regards to the Subscriber Session Steering (SSS) described in section 4.2.4.1, the DBNG-CP will be able to interact with the user plane selection function of the SSS and may optionally support this function. This feature and the interaction of this feature has not been specified in this document.

4.2.1.1 DBNG-CP Northbound Interfaces

The northbound interfaces R, cp-d, B, M, and S11 are terminated on the DBNG-CP. Also the optional interface with the SSS (UPSF), if present, is terminated on the DBNG-CP.

4.2.2 DBNG-UP Functions

Once the DBNG-CP completes authentication and address assignment, the DBNG-UP creates state related to forwarding for the subscriber session. The DBNG-UP performs forwarding, traffic management, and policy enforcement on the subscriber traffic. The following functions are part of DBNG-UP:

- **Traffic Management:** follows the instruction from the DBNG-CP on forwarding and related state for each subscriber session. For example, Layer 2 and Layer 3 QoS marking, header modification, subscriber session termination, forwarding rules, filtering rules, and QoS rules.
- **Subscriber packet statistic collection:** the ability to collect per subscriber packet forwarding statistics
- **Integrated Network Address Translation:** Please refer to TR-459.2 [20] for more details

In addition, other DBNG functions that require immediate response may be located on the local DBNG-UP to improve scaling.

- **IPTV Multicast:** Please refer to TR-459.3 [21] for more details
- **NAT function:** Please refer to TR-459.2 [20] for more details
- **HTTP Redirection:** For public Wi-Fi use cases
- **Subscriber Routing:** Routing protocol, Routing control, IP forwarding, and subscriber route management
- **Keepalive Messages:** This is the recommended mode of PPP LCP echo generation and processing. The offload improves scalability and/or failure detection times. In this case, the DBNG-UP will need to inform the DBNG-CP of any relevant changes in keepalive state.
- **Lawful Intercept:** The user plane component of lawful intercept where mirrored packet must adhere to local country standard LI format.
- **Local Control Plane:** Process all the messages mentioned: IGMP, MLD, IGP, BGP, and keepalives.

4.2.2.1 DBNG-UP Interfaces

The access interface (V) and network interface (A10) that forward subscriber data traffic belong to the DBNG-UP. Optional management interface (M) for connecting to external EMS. The (M) interface is based on YANG data models, but the specification of the required data models is for future study. Optional interface to the SSS.

4.2.3 Interfaces between DBNG-CP and DBNG-UP

With the separation of the control and user plane, interfaces are required to facilitate communication between the DBNG-CP and DBNG-UP.

4.2.3.1 Management Interface

The Management Interface (Mi) shown in Figure 3 between the DBNG-CP and the DBNG-UP allows the DBNG-CP to manage its associated DBNG-Ups supporting pushing configurations and retrieving operational state and status to and from the DBNG-Ups. For reference, this is similar to “configuration” and “show” commands for a traditional MS-BNG. Some examples of the configurations that are pushed to the DBNG-Ups might be general routing protocol configurations and QoS policy templates.

It must be noted that a DBNG-CP can be deployed with a variety of DBNG-Ups from different vendors. While traditional MS-BNGs have vendor proprietary internal representation of system resources, hardware resources and physical configurations are transparent to the DBNG-CP.

The Mi supports these functionalities, for example:

- Publishing of DBNG-UP operational data and resource information
- Notification of events and alarms between DBNG-CP and DBNG-UP
- Resource constructs such as interfaces described in TR-178 [10]




Figure 3: Management Interface

The Mi interface will be based upon YANG data models, but the specification of the required data models for the Mi is for future study.

4.2.3.2 Control Packet Redirection Interface

A separate interface is required to forward and tunnel control packets such as DHCP, L2TP Control packets, PPP/PPPoE, and Router Solicit packets through the user plane to the control plane. Figure 4 is a flow diagram that provides an example of the control messages from Residential Gateway (RG) that are tunneled through the DBNG-UP to the DBNG-CP through the default Control Packet Redirection interface (CPR interface). For the CPR interface there are three types:

- A single default CPR interface per DBNG-UP for initial access control packet redirection which is described in detail in section 4.5.1.
- A per logical port CPR interface described in detail in section 4.5.2.
- A per subscriber session based CPR interface, an example of such is captured in section 4.5.4.

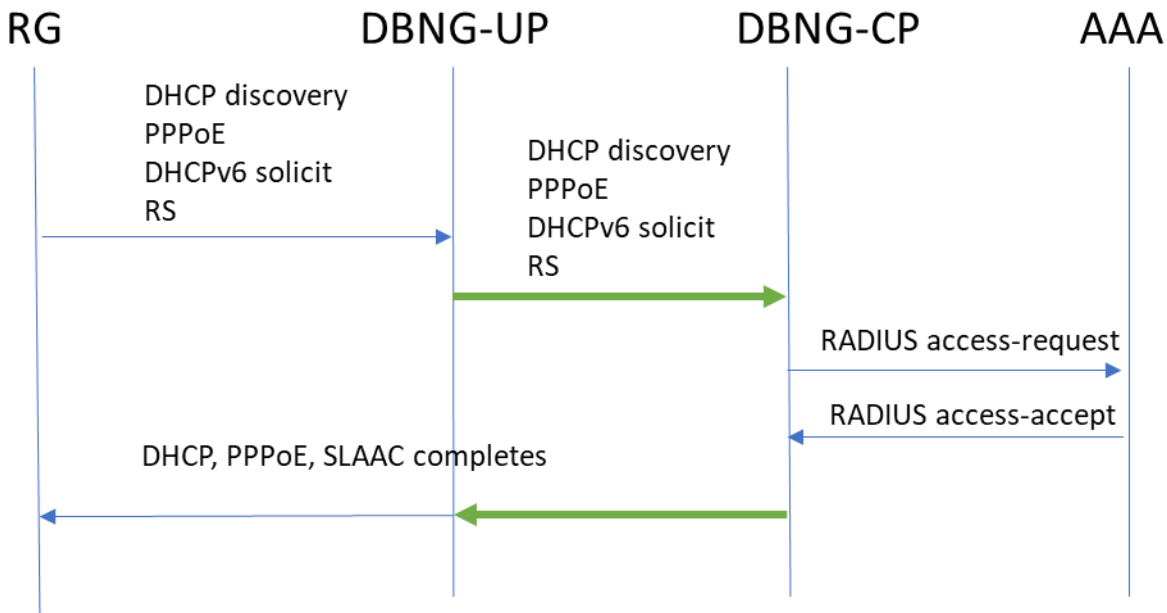


Figure 4: Example of Control and User Plane control message exchange

Figure 4 shows that a Control Packet Redirect Interface (CPR Interface) between the DBNG-UP and DBNG-CP is required for triggering subscriber authentication. With DBNG, the DBNG-CP and the DBNG-UP each become a separate network function. The network separation between DBNG-CP and DBNG-UP can vary from a small layer 2 domain to a layer 3 multi hop network. Therefore, control packets are sent over a tunnel. Figure 5 below illustrates at a high level the functionality of a CPR Interface.

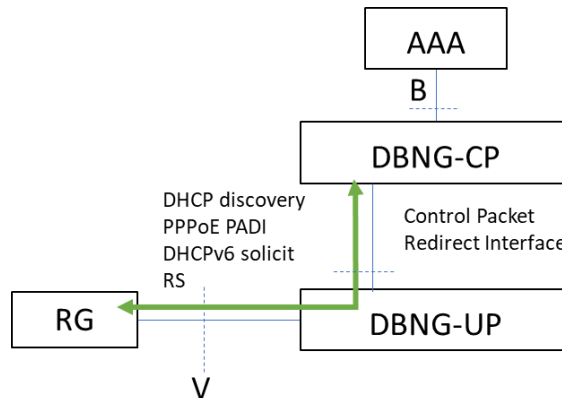


Figure 5: Control Packet Redirect Interface

A subscriber session typically starts with control messages from subscriber access, which may request one or more addresses. The DBNG-UP must redirect these control messages to the DBNG-CP. This default redirect rule would be signaled by the DBNG-CP to DBNG-UP after a successful association between the DBNG-CP and DBNG-UP. The DBNG-UP can have:

- Session context for a subscriber (a known subscriber on the DBNG-UP)
- No session context for a subscriber (an unknown subscriber on the DBNG-UP, for example, new subscriber connecting to the network for the first time).

Since the DBNG-CP is decoupled from the DBNG-UP, the DBNG-CP has no access circuit information (ex. Logical port, etc.). Therefore, the DBNG-UP is required to include data plane information as meta-data when redirecting control packets.

4.2.3.3 State Control Interface

The State Control Interface (SCi) is used to program a default rule to redirect control packets between user plane and control plane. Once the subscriber has successfully authenticated, the DBNG-CP will install traffic forwarding rules and related states to the DBNG-UP as shown in Figure 6. Successful installation of the traffic rules will be indicated by an acknowledgement from the DBNG-UP. After the traffic rules are programmed onto the DBNG-UP, the DBNG-UP will forward the subscriber data traffic according to the rules.

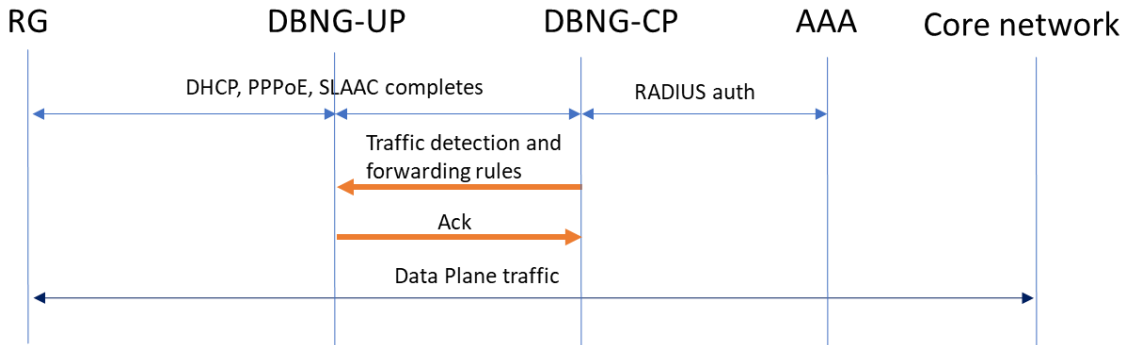


Figure 6: Example of Control Plane pushing forwarding rules to the User Plane

Figure 7 illustrates the third interface, SCi, that is required to program traffic detection and forwarding rules.

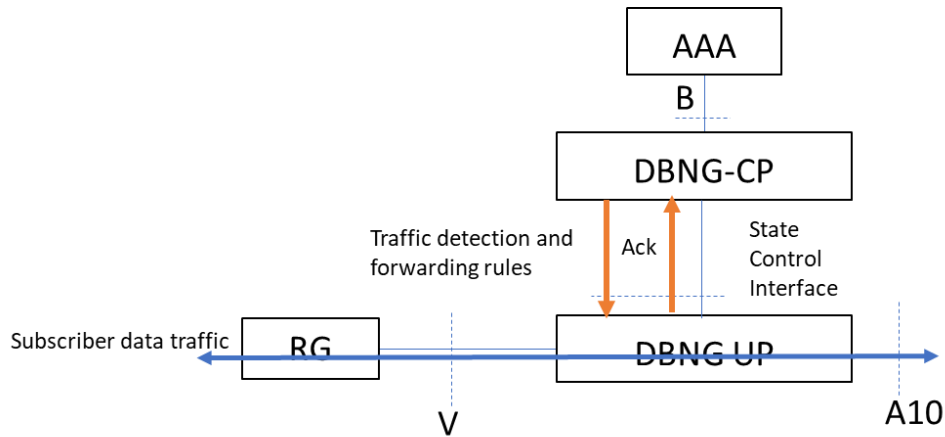


Figure 7: State Control Interface

Since MS-BNG can support different access types and protocols on the V interface, the traffic detection and forwarding rules must be flexible. Figure 8 below shows the different user plane combinations with respect to the MS-BNG functions as described in their respective TRs.

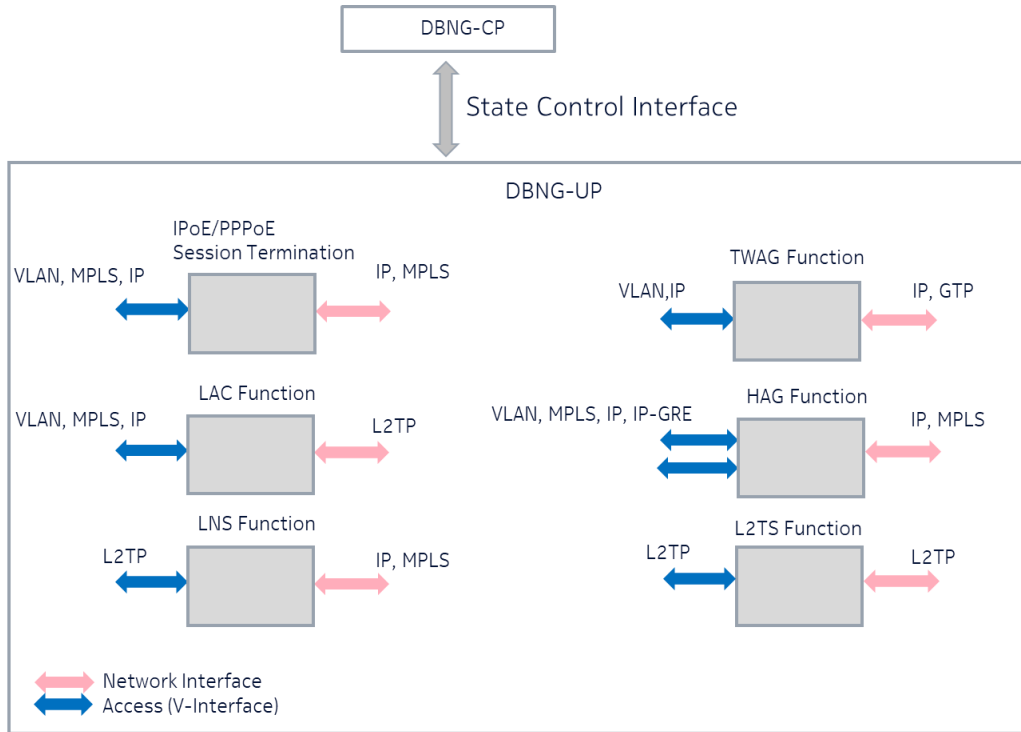


Figure 8: Example of User Plane combinations

The basic traffic detection and forwarding rules in the upstream direction (e.g., access to network) and the downstream direction (e.g., network to access) follows the same pattern and fundamentally consists of session identification followed by one or more actions. Figure 9 and Figure 10 are logical representation of DBNG-CP sending directives to DBNG-UP, instructing the DBNG-UP to install basic forwarding state for fixed L2 access (e.g., access from DSLAM or OLTs over Ethernet).

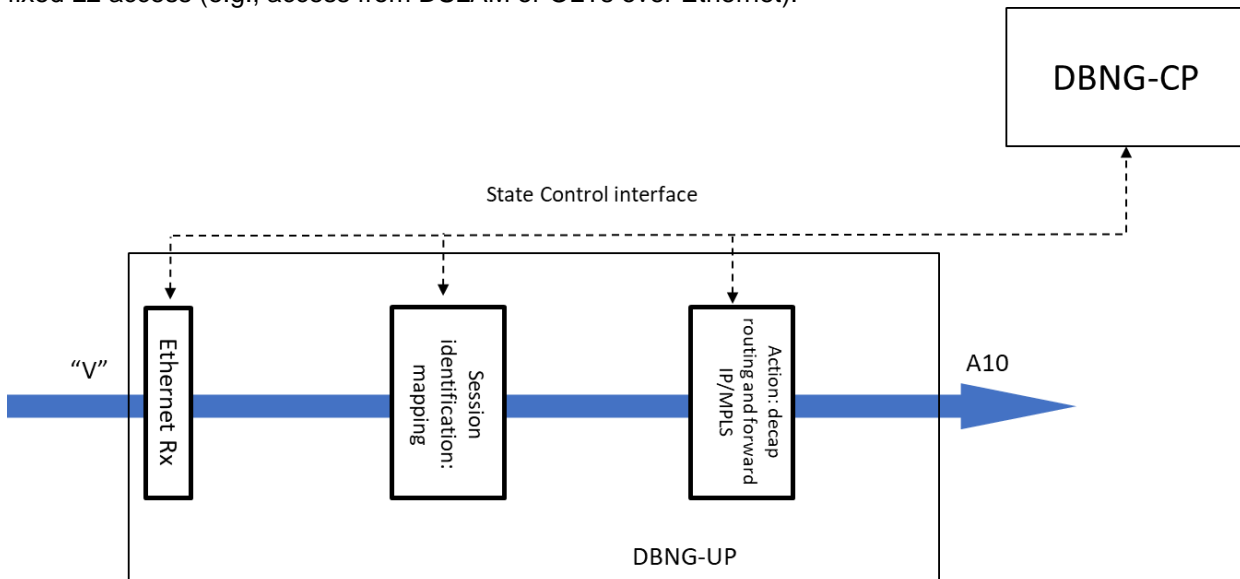


Figure 9: State Control Interface for Access to Network direction

Direction Upstream – Access to Network:

- Session-identification: Port/VLAN-Tag(s) + Subscriber-MAC
- Action: Remove encapsulation, IP FIB lookup, Forward to network.

Please note: In this example, the session-identification is only applicable to L2 subscriber. LNS Session-identification for access facing interface may differ with this example.

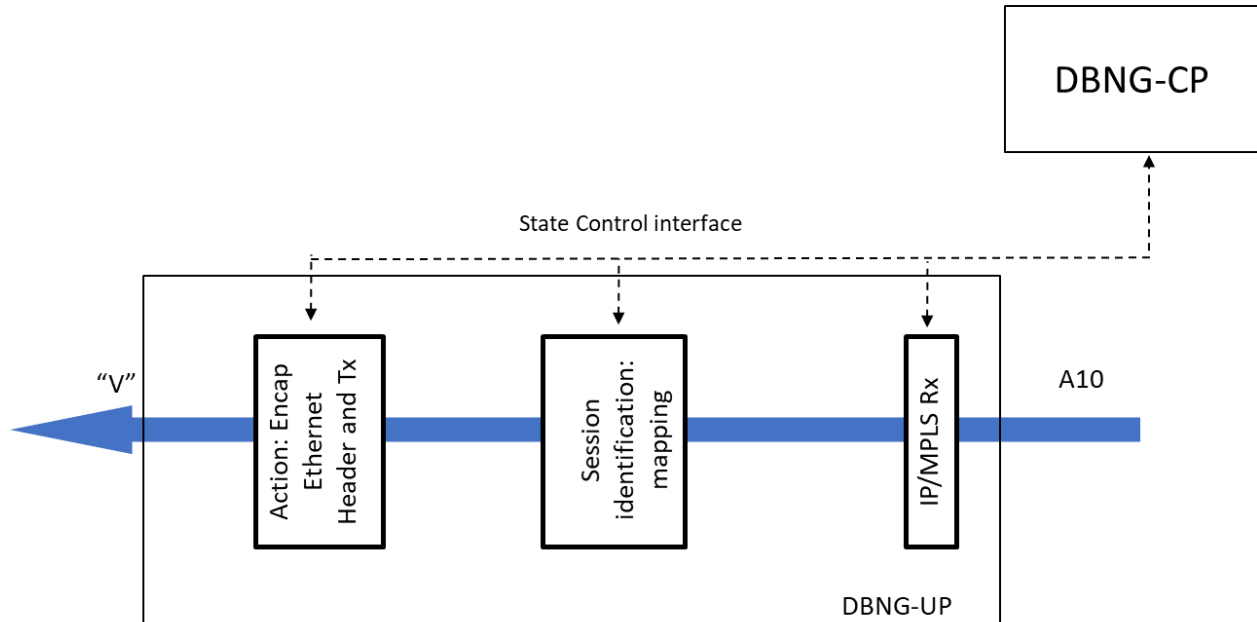


Figure 10: State Control Interface for Network to Access direction

Direction Downstream – Network to Access:

- Session-identification: IP address and VPN instance
- Action: Lookup IP destination address, build encapsulation using Port/VLAN-Tag(s)+subscriber-MAC, Forward to access.

The session state on the DBNG-UP is always controlled by the DBNG-CP i.e., the DBNG-UP just follows the directive from the DBNG-CP to install, modify and delete the session state. In addition to basic forwarding state, the DBNG-CP can also associate, update, and disassociate rules such as related to the following:

- Filtering
- SLA management
- Statistics collection
- Credit control
- Traffic mirroring
- Application aware policies

Cases where DBNG-CP would trigger session state change:

1. DBNG-CP triggers a new session state on the DBNG-UP. E.g., when the subscriber successfully authenticates
2. DBNG-CP triggers an update of session state on the DBNG-UP. E.g., triggered by an update from the policy-server or upon receiving a RADIUS Change of Authorization (CoA)
3. DBNG-CP triggers the deletion of session state. E.g., based on administrative action, session termination or a disconnect-message initiated from the AAA server.

The State Control interface also has the ability to report events and alarms from DBNG-UP to the DBNG-CP for individual subscriber session and as well as group of sessions (e.g., subscriber groups and subscriber prefix).

4.2.4 DBNG High level Architecture

In summary, the following are the identified interfaces required for DBNG-CP and DBNG-UP communication.

1. Management Interface (Mi)
2. Control Packet Redirect Interface (CPR interface)
3. State Control Interface (SCi)

Together with the DBNG functions and interfaces, Figure 2 illustrates a high level architecture of CUPS for a DBNG and the new defined interfaces between the DBNG-CP and DBNG-UP.

4.2.4.1 Traffic Steering Function

The separation of control and user plane introduces the option for multiple DBNG-Ups to be under the control of a single DBNG-CP. Multiple DBNG-Ups may be deployed to achieve the desired network scale, to support different service levels, to provide Multi-Access Edge Compute (MEC) services to a subset of subscribers, or to provide DBNG-UP functions within different network slices.

The Subscriber Session Steering (SSS) is a new network function, which is being specified by BBF, that will provide the capability to steer (i.e., to direct) customer sessions to the most appropriate DBNG-UP, allowing load balancing of sessions across the available DBNG-Ups and/or the selection of a specific DBNG-UP to accomplish specific services and allow the access to specific data networks and platforms. It will allow to overcome the static nature of the connectivity between the Access Node and the BNG where subscribers on a particular access node are mostly connected to the same BNG, with the decision and configuration being done only when the network is deployed or upgraded.

Among the key functions of the SSS Function, a user plane selection function is included: this is responsible for making the real time decision regarding which DBNG-UP the subscriber session should be connected to. With regards to the relation between SSS and DBNG, in general, the SSS will be defined such that its interaction with the DBNG will be abstractly specified and:

- DBNG-CP will mainly interact with the Control Plane of the SSS;
- DBNG-UP will mainly interact with the User Plane of the SSS.

The requirements for DBNG that will ensure a standard way of interacting with the SSS are FFS and will follow the publication of the SSS specification. At this point in time, the main interactions that are foreseen between the DBNG and the SSS are the following (not exhaustive list):

- The DBNG-CP will request to SSS to select the DBNG and its specific DBNG-UP where the subscriber session has to be anchored, possibly providing it with relevant policy received by AAA.
- The SSS will indicate the selected DBNG-UP to the DBNG-CP.
- The DBNG-CP will signal to the SSS the capabilities supported by its DBNG-Ups.
- The DBNG-UP will be connected (directly or indirectly) to the user plane of the SSS.

When the SSS is deployed in the network, the selection of DBNG-UP may be performed at the session setup and/or during the session lifetime. The following are examples of selection criteria that the SSS might apply to assist DBNG-CP:

- The current load of the DBNG-Ups under control of the DBNG-CP.
- The knowledge of the service platforms co-located with the various DBNG-Ups.
- The knowledge of the group of subscribers that may have dedicated DBNG-UP (network slicing).
- The need of executing a maintenance operation, such as removing a DBNG-UP from service.
- The need of guaranteeing a certain service level agreement.

The SSS user plane selection function might be implemented as a stand-alone network function, as a part of an SDN controller or even as an add-on of a DBNG-CP. In the latter case, the DBNG-CP will be of course required to support additional requirements which are being specified for the SSS. Notice that the implementation of the SSS user plane selection function on a DBNG may be sub-optimal in case there are more than one DBNG-CP deployed in the Operator's network.

4.2.4.2 Lawful Intercept Function

Figure 11 provides an overview of the DBNG Lawful Intercept functionality. The Lawful Intercept function-CP and Lawful Intercept function-UP have been already described in sections 4.2.1 and 4.2.2.

The external function Lawful Enforcement Agency Mediation Device (LEA MD) provides the interface used to set up and provision the lawful intercept on the DBNG, generating requests to DBNG-CP for traffic-mirroring and removing existing traffic-mirroring targets. In the following call flows, the LEA MD plays the role of LI requester.

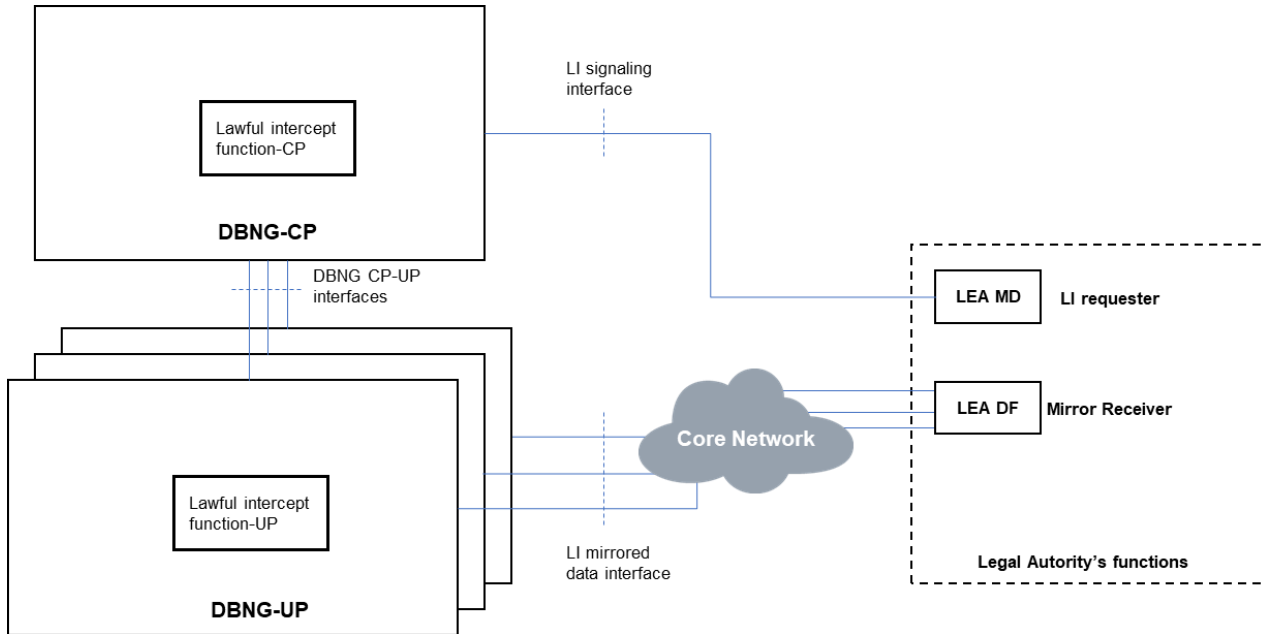


Figure 11: DBNG Lawful Intercept Function

The external function Lawful Enforcement Agency Delivery Function (LEA DF) has the role to collect the traffic mirrored by the DBNG-UP. In the 4.5.32 call flows, the LEA DF plays the role of mirror receiver. Both LEA MD and DF belong to the Legal Authority domain. Generally, they are combined in a single MD/ DF function.

LI Signaling interface is used by LEA MD to request interceptions, modify, and delete previous interception requests; it is also used by DBNG-CP to notify normal events, such as the start / end of mirroring or extraordinary events such as mirroring failures, reboots, and switchovers.

4.3 Deployment models

The following section describes different deployment scenarios.

4.3.1 Deployment model: Geographical separation of DBNG-CP and DBNG-UP

The DBNG system can cover the services and subscribers in a broad geographic region. The control plane can be deployed centrally with the user planes deployed in different areas closer to subscribers as shown in Figure 12.

For example, take three areas A, B, and C, each with its own user plane. The three user planes share one control plane. The control plane can be deployed in a centralized location at a Core Data Center as opposed to the user planes which are deployed at city edge datacenters closer to subscribers. In this design, data centers and edge data centers are selected based on their location and responsibilities.

The centralized control plane communicates with outside subsystems and user planes for control and management. Under the control plane's instruction, users' traffic is forwarded by DBNG-UP to the Internet. Although not depicted in the diagram, each DBNG-UP may have an interface to the EMS for management purposes.

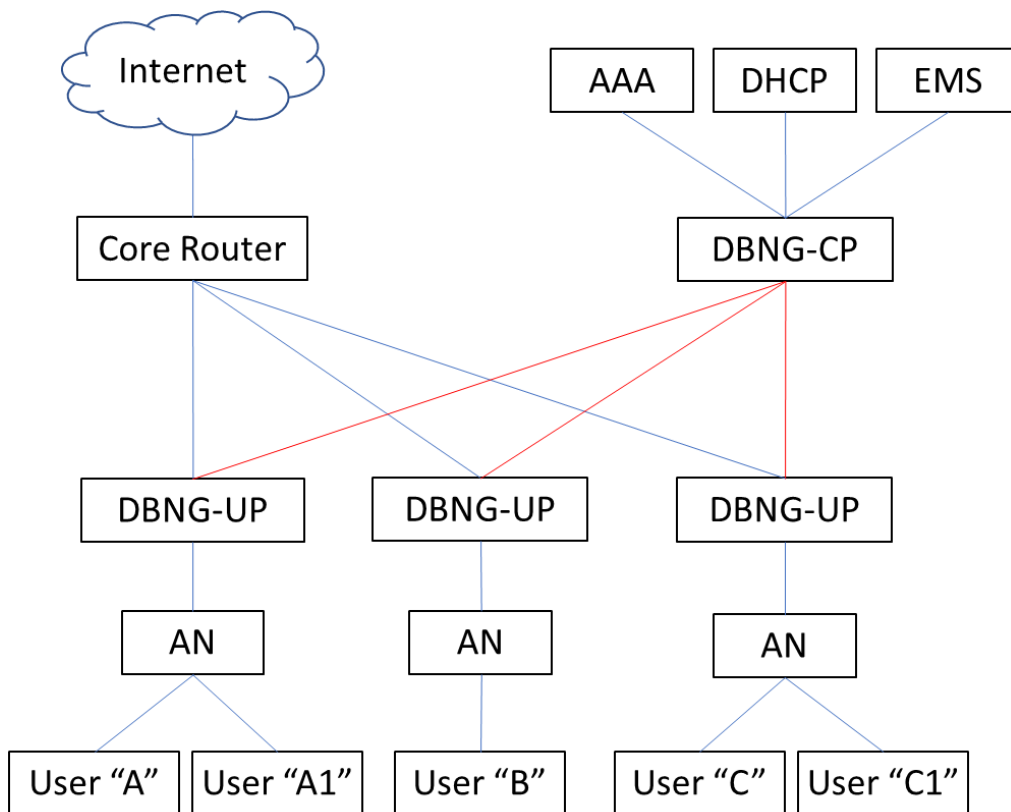


Figure 12: Geographically distributed deployment model

4.3.2 Deployment model: Non-Geographical separation of DBNG-CP and DBNG-UP

Due to scalability and security requirements, the DBNG-CP and DBNG-UP may be placed together at the same central office (CO). In this case, the CO installation does not geographically distribute the DBNG-UP as shown in Figure 13. A Point of Delivery (POD) may consist of multiple DBNG-Ups and a single DBNG-CP. In this deployment model, the same deployment principles described in 4.3.1 apply, but since the DBNG-CP and DBNG-Ups are within the same secured domain certain security requirements can now be removed. Although not depicted in the diagram, each DBNG-UP may have an interface to the EMS for management purposes.

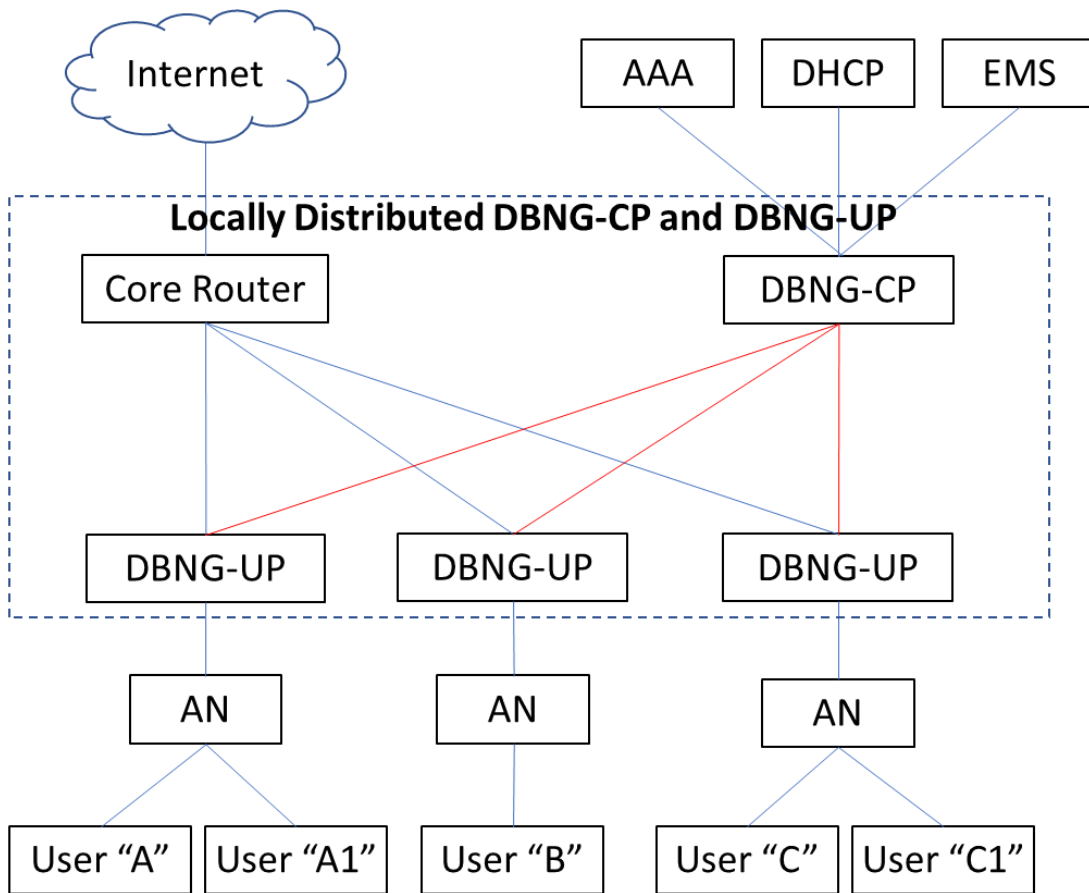


Figure 13: Non-Geographically distributed deployment model

4.4 Subscriber Session Resiliency and IP Prefix Management

This WT proposes a UP resiliency solution for the following type of access: IpoE(Local Server) and PPPoE. IpoE(Relay), TWAG, HAG, LAC, LTS, LNS are FFS.

A MS-BNG can support resilient connectivity towards the AN through capabilities such as IEEE 802.3 Link Aggregation, VRRP, MPLS Pseudowires or Ethernet VPN. However, for a traditional (not disaggregated) MS-BNG, there is no standardized method of supporting resilience between multiple MS-BNG instances that allows the subscriber session state to be maintained during the switchover. MR-459.2 describes how the disaggregation of the BNG creates an opportunity to improve subscriber resilience due to the fact that the DBNG-CP contains a centralized state database for all of the DBNG-UP under its control, where the DBNG-CP is the authoritative source of subscriber state. This section details how the DBNG can support switchover of subscriber sessions from one DBNG-UP to another DBNG-UP whilst maintaining their state.

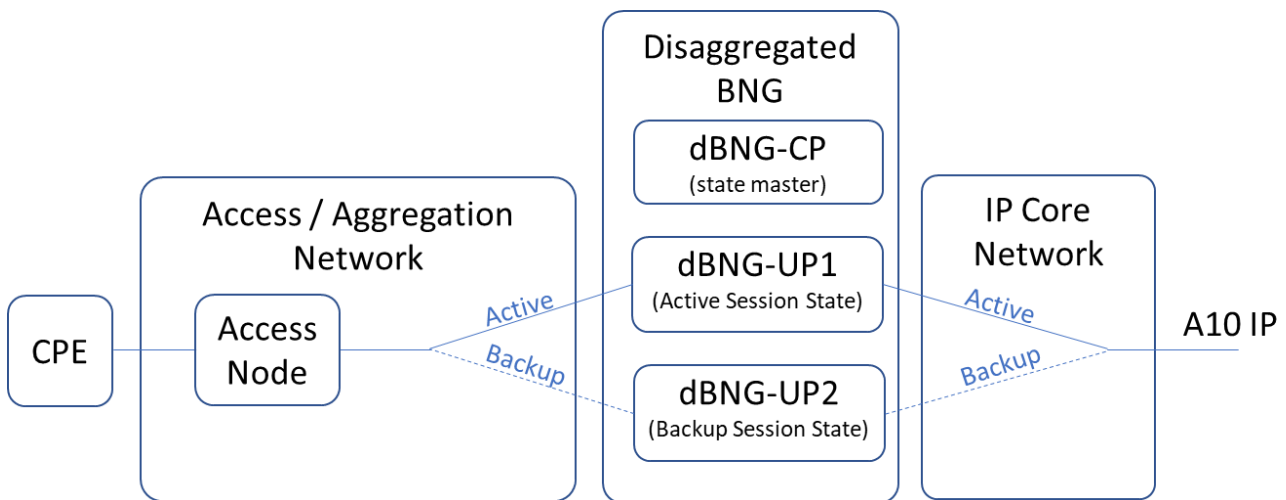


Figure 14: DBNG subscriber session resilience

Figure 14 shows how the DBNG can provide resilience across DBNG-UP where the DBNG-CP holds the authoritative state for any subscriber session, and the DBNG-UP may hold active forwarding state, or backup forwarding state for a particular subscriber session.

A DBNG inherently supports cold-standby resiliency under the condition that the Access Network (in this specification we assume that the Access Network includes the Access Node) has paths from the CPE/RG to more than one of its DBNG-UP. Under this condition, all the CPE/RG with failed sessions on a DBNG-UP would be able to retry a new connection through another DBNG-UP lying under the control of the same DBNG-CP. Notice that with DBNG, it is not necessary to resort to PADO delay mechanism, as there is one control entity that handles the session. Therefore, the CPE/RG will be served by the only DBNG-UP which gets programmed by the DBNG-CP.

It is important to note that the active/backup status applies to a group of subscriber sessions, rather than the DBNG-UP as a whole. This allows one DBNG-UP to be active for a group of subscriber sessions and backup for other sessions simultaneously. This active/backup resiliency mechanism is different from cold-standby resiliency, where the subscriber state on network side is maintained across the failure and therefore the CPE/RGs don't need to have an active role.

There are three key aspects to providing resilience that will allow a session (or a group of sessions) to switchover from one DBNG-UP to another DBNG-UP without requiring the session to reconnect:

1. That the Access Network must have paths from the CPE/RG to more than one DBNG-UP with a way for the traffic to switchover from an active DBNG-UP to a backup DBNG-UP.

2. It must be possible to install the state for the same session onto two or more different DBNG-Ups in an active/backup relationship (where only one DBNG-UP is active for a given session).
3. There must be paths from the core network to the relevant DBNG-Ups, and there must be a way for the routing to be updated to direct traffic to the currently active DBNG-UP. Traffic will be redirected to the current active DBNG-UP during switchover as described in 4.4.10.

The three items above need to be synchronized such that the entire network only considers one DBNG-UP as active for any one subscriber session at any one time.

Note:

1. Active/Active load balancing capability may potentially be feasible but is FFS.
2. The current version of this specification only considers the scenario where the subscriber session is terminated into a L3 routed network by the DBNG-UP.
3. The resiliency of the DBNG-UP, as the role of a LAC, is not defined in this document.

4.4.1 Differences between Active and Backup Session State

The differences in session state between the active and backup DBNG-UP are as follows:

- The detection and forwarding rules will apply to different logical ports and per subscriber session CPR interface tunnels.
- In the active state, the DBNG-UP will generate any offloaded control packets (such as LCP echo requests), if it supports the offloading capability and has been requested to use it.
- In the backup state, the DBNG-UP must not generate offloaded control packets.
- By default, a backup DBNG-UP fully implements all aspects of the subscriber state, other than the offloaded control packet messaging mentioned above. This includes allocation of forwarding, scheduling, and accounting resources exactly as if it were the active DBNG-UP.
- Optionally, the DBNG-CP may allow the backup DBNG-UP to oversubscribe backup session state such that it can have more backup sessions configured than it can have simultaneously active. In the case that the DBNG-CP indicates that this is allowed for a certain Subscriber Group (see section 4.4.10), the DBNG-UP will receive and respond to session establishment messages as normal, but it may optionally choose not to fully allocate or program all resources that would be required to forward traffic for the session. The resources will then only be consumed when the DBNG-UP becomes the active DBNG-UP for the relevant session. This is given the term 'partial state implementation'.

Otherwise, the session state created on a backup DBNG-UP will have the complete set of characteristics that are installed on the active DBNG-UP, for example: Detection and Forwarding rules, QoS, Lawful Intercept and Accounting.

4.4.2 Introduction to Subscriber Groups

The subscriber prefix assignment and the resiliency use case in the previous sections 4.4 and 4.4.1 demonstrate that it is necessary to group subscribers together to share states and resources: active state, backup state, the prefix state for network advertisement. Therefore, the new concept "subscriber group (SGRP)" is introduced in this WT to group subscribers together that share states and resources; for example, subscribers sharing the same back-up state or subscribers sharing the same IP prefix. Currently, in this WT, SGRP is used to group subscribers together to share resiliency state and subscriber prefixes. However, SGRP is not only limited to these use cases and is a general mechanism to group subscribers together. The use of SGRP for additional use cases (such as resilience for PPP sessions over L2TP) is FFS

4.4.3 Use of the Subscriber Group for resilience capability

Subscriber sessions that are subject to the same restoration capability are placed into the same Subscriber Group (SGRP). This helps to minimize the messaging and therefore the elapsed time between the detection of a failure and the restoration of service or the elapsed time between an explicit switchover and its execution.

The active or backup status is set at the SGRP level, and communicated to the relevant DBNG-UP by the DBNG-CP. Subscriber sessions are tagged with the SGRP to which they belong at session establishment. All resilience actions are communicated at the SGRP level rather than at the session level. If all subscriber sessions within the SGRP are terminated on a single logical port, SGRP state installed on each DBNG-UP may also include this logical port to help the DBNG-UP identify relevant failures.

Since resiliency operates at the SGRP level, the switchover of a set of sessions from an active UP can be triggered even without any fault on the active UP (see section 4.4.5 for a list of triggers).

4.4.4 Mapping Subscriber Groups to DBNG-UP

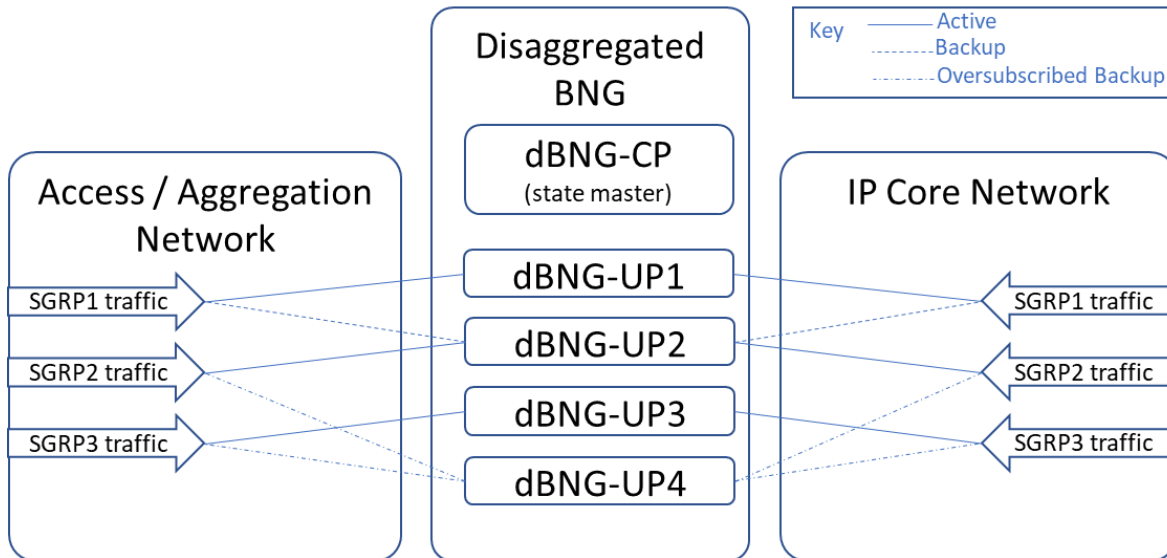


Figure 15: Example SGRP to DBNG UP mapping

Figure 15 illustrates example mapping of SGRP to DBNG-UP for the purposes of delivering resilience. This includes the ability for a DBNG-UP to be simultaneously active for one set of SGRPs and backup for another set of SGRPs (as per DBNG-UP2 in Figure 15) as well as the option for SGRP to be oversubscribed where the DBNG-UP is not capable of simultaneously supporting the activation of all of its Backup SGRPs (as per DBNG-UP4 in Figure 15 assuming that it can only support either SGRP2 or SGRP3 at one time).

4.4.5 Switchover triggers

The following provides a non-exhaustive list of triggers to switchover the active state for a particular SGRP from one DBNG-UP to a different DBNG-UP:

1. Operator Directed: A manual or automated trigger through the management interface of the DBNG-CP.
2. Failure of an entire DBNG-UP.
3. Failure of a component of the DBNG-UP that impacts a set of active subscriber sessions.

4. The failure of a link or interface or virtual circuit directly connected to the DBNG-UP that impacts a logical port and active subscriber sessions.
5. A notification to the DBNG-UP of a failure in the access/aggregation network (for example, this might be communicated through mechanisms such as IGP, MPLS or Pseudowire signaling).
6. A change in the negotiated status of a resilient connection between the DBNG-UP and the Aggregation Network (such as the change in the preferential forwarding bits for a resilient pseudowire, or LACP link selection)
7. A change in the IP core network that isolates a DBNG-UP from the rest of the network. This might be detected by the DBNG-UP but could also be detected by some other entity that notifies the DBNG-CP from a specific network instance.
8. A trigger from the SSS to the DBNG-CP, as a consequence of a logic programmed in the SSS.
9. Long connectivity fault detected by the DBNG-CP

In all of the cases above, the DBNG-CP can take the action to change the active/backup status for a SGRP on the relevant DBNG-UP.

In cases 3 through 7, it is also possible for the DBNG-UP to take some immediate independent action without waiting for an instruction to come from the DBNG-CP, in order to speed up the restoration of service. This is described more in section 4.4.7.

4.4.6 Virtual MAC Address

For a subscriber session to be moved from one DBNG-UP to another DBNG-UP for resilience purposes, the MAC address of the DBNG used to communicate with that subscriber must not change. A Virtual MAC address will therefore be allocated by the Control Plane for all communication between a subscriber and the DBNG, irrespective of the DBNG-UP that is currently active for that subscriber session. This Virtual MAC will be defined at the SGRP level and therefore will be the same for all subscriber sessions within a particular SGRP. The same MAC address may optionally be reused across multiple SGRPs.

4.4.7 DBNG-UP independent actions

In addition to Control-Plane initiated resilience actions, in some cases it is also possible for the DBNG-UP to take independent action without waiting for an instruction to come from the DBNG-CP. If the DBNG-UP becomes aware that its logical-port is now active, it has to advertise the SGRP prefixes associated to that logical-port according to the relevant policy.

For an active DBNG-UP: When the SGRP is tagged with a specific logical port, then a failure of that logical port indicates to the DBNG-UP that it is not able to serve that SGRP, and it can automatically take the relevant actions including stopping sending any offloaded control messages and updating / withdrawing relevant prefixes that might have previously been advertised to the core IP network. Before taking these actions, it will send a notification to the CP.

For a backup DBNG-UP: In the specific case where the status of the logical port is negotiated with the Access/Aggregation Network (such as through the preferential forwarding bits for a resilient pseudowire), then the backup DBNG-UP can independently be aware that it is now required to become active.

In order to support the above, the SGRP state installed on the DBNG-UP includes the actions to be taken when the SGRP is in active and in backup, irrespective of the current status.

4.4.8 SGRP and Logical Port Relationship

As defined in Section 2.3, logical port is an abstract entity. A single logical port can host one or more subscriber session(s). In the case where the logical port hosts more than one subscriber session, it must be possible to group different subscriber sessions into different SGRPs.

4.4.9 CP Logical port resilience groups

Whilst the active logical port for a subscriber can be discovered during session establishment, resilience requires that the DBNG-CP has some advanced knowledge of the logical ports that are resilient for each other across all DBNG-UP under its control. These logical port resilience groups will include two or more logical ports. Without the logical port resilience group information, the CP will not be able to identify the DBNG-UP and logical port against which the backup state should be created.

This information can be provided to the DBNG-CP through local configuration. A mechanism for this information to be automatically reported from the DBNG-UP to the DBNG-CP is for further study.

4.4.10 Resilience Information attributed to the Subscriber Group

The creation of a subscriber group is signaled by the DBNG-CP to the DBNG-UP using the SCi. The information contained in the SGRP to support resilience is as follows:

- The List of IP prefixes that are associated with the SGRP (that will be advertised to the IP Core network)
 - For each prefix: Indication as to whether the prefix is to be advertised within the network instance when the SGRP is backup
 - For each prefix: Information that will influence the routing advertisement attributes within the network instance when the SGRP is active
 - For each prefix: Information that will influence the routing advertisement attributes within the network instance when the SGRP is backup

This information may allow the invocation of a particular routing policy preconfigured on the Ups which can influence the routing attributes (e.g., BGP local preference or other attributes).
- The associated Logical Port (if any) for the DBNG-UP to take independent actions (see section 4.4.7)
- The Virtual MAC: Required to establish resilient sessions. A SGRP contains only one virtual MAC.
- The Requested redundancy state which will be one of:
 - Track Logical Port (if the Logical Port is down, or is not selected as active through signaling in the access network, then the state is Backup, otherwise Active)
 - Active
 - Backup
- The Backup strategy
 - Full state implementation: No oversubscription on DBNG-UP is allowed. The DBNG-UP, even in backup state, must reserve all resources and install subscriber state as soon as it receives the rules from the CP, for all the subscribers in the SGRP.
 - Partial state implementation: Oversubscription on DBNG-UP is allowed. The DBNG-UP, when the SGRP is in backup state, may choose not to reserve all resources to install all state for the

subscribers in the SGRP, but it must keep locally all of the rules received by DBNG-CP and make them effective as soon as it becomes active. See section 4.4.10.1 for further explanations about Partial state.

4.4.10.1 Partial State Properties

As mentioned in section 4.4.3, a DBNG-UP playing the role of backup DBNG-UP for resilient sessions may be oversubscribed. This allows Operators to exploit the spare resources deployed in the network in a more efficient way, shouldering the risk of resource exhaustion in case of multiple simultaneous faults, while guaranteeing a certain degree of protection from non-concurrent faults.

If the DBNG-UP supports “Partial State” and has previously advertised its capability to DBNG-CP, when the DBNG-CP programs or updates the backup session state on a DBNG-UP, this DBNG-UP, at its discretion, may install (or consume local resources for) only a subset of the backup PFCP rules, and store the other rules locally in a way that does not consume as much resource. For example, the DBNG-UP may install the forwarding rules and store the QoS and/or usage reporting rules; it may take this approach for all the sessions or for some sessions only. The subscriber state is then fully installed by the DBNG-UP (from the locally stored information) in case of a switchover, i.e., when an SGRP becomes active on the DBNG-UP. The time of restoration of full subscriber state for sessions of an SGRP depends on how quickly the DBNG-UP can install the remaining rules; however, all the sessions which have undergone the switchover will achieve full state on the DBNG-UP on which they have moved. In case the DBNG-UP, due to the excessive oversubscription, runs out of resources when installing the full state for the sessions of an SGRP, it will notify the DBNG-CP about the failure, for it to take actions.

4.4.11 Use of Subscriber Groups for assignment and management of IP Prefixes

Subscriber Groups can also be used in non-resilient contexts to assign Ipv4 and Ipv6 user prefixes to a DBNG-UP and revoke them when they are no longer needed. This enables the DBNG-CP to dynamically allocate prefixes to DBNG-Ups when it performs the function of IP address assignment. In other words, Subscriber Groups allow to accomplish the function of IP address assignment and management of the DBNG system.

To assign an IP prefix to a DBNG-UP, which in turn will be advertised by the DBNG-UP as owner of the user traffic routing, the DBNG-CP executes two steps:

- First it creates on the DBNG-UP a Subscriber Group (SGRP);
- Then it associates to the SGRP a list of prefixes.

Note that an initial list of IP prefixes can be signaled to the DBNG-UP together with the SGRP creation message.

The prefixes can be pulled-back by the DBNG-CP in case they are not used by any subscriber session terminated on the DBNG-UP, so that an efficient usage of the IP addresses can be administered by the DBNG-CP.

Even in the case resiliency is not required for an SGRP, setting appropriately the “BBF SGRP” IEs and the “BBF UP Subscriber Prefix” IEs (see section 6.7.6 and 6.7.7) allows control of the routing advertisement of IP prefixes assigned to a DBNG-UP. There are a number of IEs that, however, may be not applicable or not needed in case resiliency is not required.

In the following, some guidelines are provided on the use of the IEs when resiliency is not required: refer to sections 6.7.6 and 6.7.7 for the exact meaning of the terms used.

For the following BBF SGRP IEs:

- **SGRP state** could be set to

- “Active” to advertise the prefixes with the routing policy invoked by the “BBF Active Prefix Tag”
- “Backup” to advertise the prefixes with the routing policy invoked by the “BBF Backup Prefix Tag”.
- It is not recommended to set “SGRP state” to “Track Logical Port” as this may consume User Plane resources and trigger routing advertisements without benefit
- **SGRP Virtual MAC** is not required except in the case that resilience may be added at a later time.
- **Route Advertisement State** could be set to control whether or not the DBNG-UP has to advertise the prefixes if the SGRP is set as “Backup”.
- **Logical-Port** is not necessary because “Track-Logical-Port” state is not recommended.
- **Partial State Allowed** is not required as there is no other DBNG-UP where the SGRP is programmed.

For the following BBF UP Subscriber Prefix IEs:

- **BBF Active Prefix Tag** may be set, as well as **BBF Backup Prefix Tag**, to invoke the right routing policy on the UP.
- **Network Instance** may be set to assign the prefix to a network instance other than the default one.

4.5 Call Flows

In this section:

- The call flows are informative and “session forwarding state” is hereinafter referred to as “session”.
- For sessions that utilize immediate session creation, described in section 6.5.1.3, resources are consumed on the DBNG-UP in order to allow individual subscriber control packet management such as blocking, rate limiting, and specific packet filtering.
- For sessions that utilize delayed session creation, described in section 6.5.1.4, additional DBNG-UP resources are not consumed, but individual subscriber control packet management is not possible.

4.5.1 Control Plane and User Plane Association

It is assumed that the two nodes (DBNG-CP and DBNG-UP) are provisioned or otherwise made aware of their partner nodes. E.g., provisioning, configuration, auto-discovery, is outside the scope of the CUPS protocol. DBNG-UP and DBNG-CP must form an association. The reference procedure is defined in 3GPP TS 29.244 [28] clause 6.2.6 and can be initiated by DBNG-CP or DBNG-UP. Afterwards, the DBNG-CP requests to start a generic session with DBNG-UP to program forwarding rules on the DBNG-UP. The forwarding rules instruct the DBNG-UP to redirect control messages over the CPR Interface.

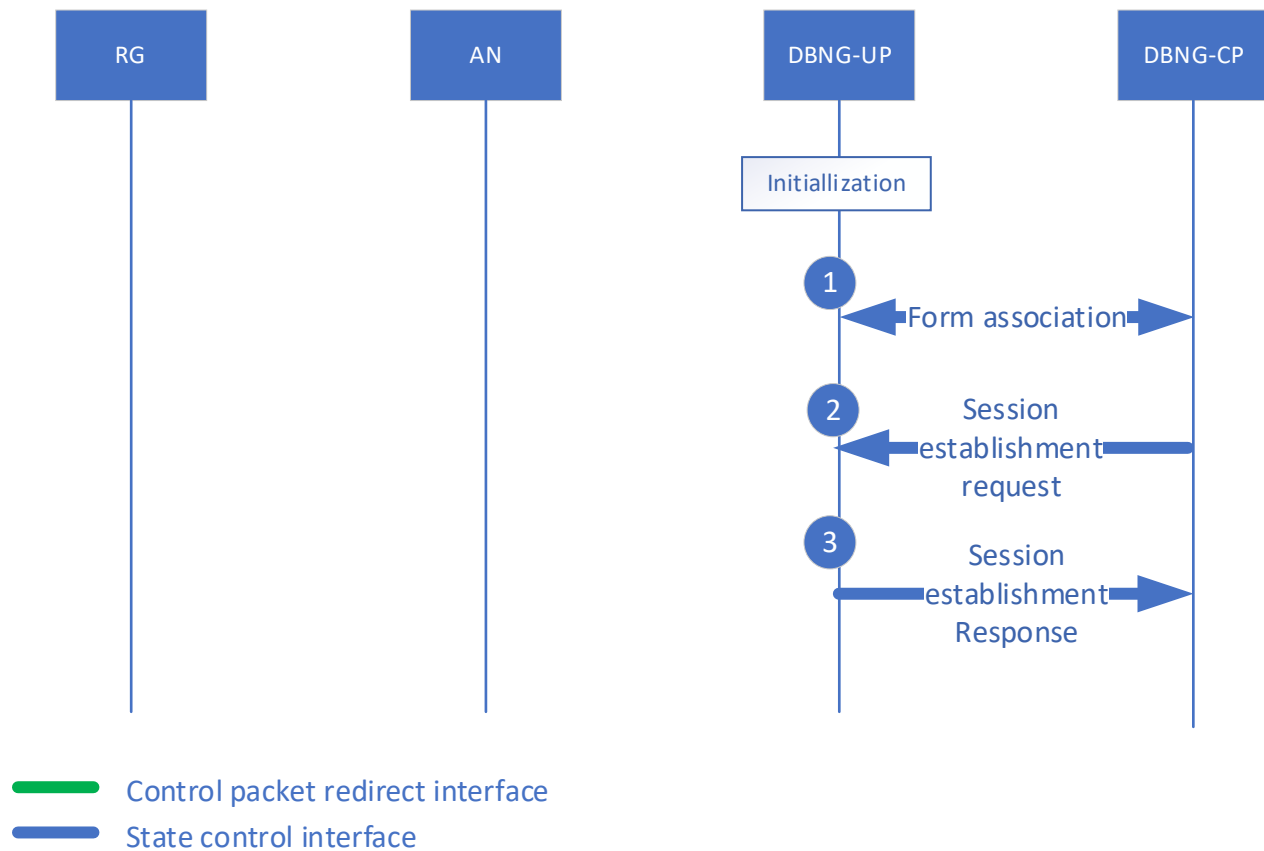


Figure 16: DBNG-CP and DBNG-UP association

1. DBNG-UP and DBNG-CP form an association. The association can be started by either the DBNG-UP or DBNG-CP. During the association process, the DBNG-UP and DBNG-CP will exchange capabilities information, e.g., types of BNG functions.
2. Afterwards, the DBNG-CP requests to start a generic session with DBNG-UP to program forwarding rules, representing the default Control Packet Redirection tunnel, on the DBNG-UP via a PFCP Session Establishment Request. The DBNG-CP may also optionally program additional control packet redirect forwarding rules on a logical port granularity to support delayed session creation. As This is described in section 4.5.2.
3. The DBNG-UP responds to the DBNG-CP session establishment request with either a success or failure.

4.5.2 Initial Control Packet Redirection Rule

After the association between DBNG-UP and DBNG-CP is formed, a PFCP Session Establishment Request is sent by the DBNG-CP to the DBNG-UP to program the default Control Packet Redirection interface (CPRi). The DBNG-CP programs control packet redirection rules to instruct the DBNG-UP to redirect specific types of control packets to the DBNG-CP through the CPR interface for upstream direction.

In addition to the default CPRi, if the DBNG-UP supports “per-logical-port CPR” tunnel, the DBNG-CP may program on DBNG-UP per-logical-port CPR tunnels as well. If programmed, a per-logical-port CPR tunnel must be preferred to the default CPR tunnel by the DBNG-UP when delivering to the DBNG-CP control traffic incoming from that logical-port.

Both the default CPR and per-logical-port CPR interfaces are common to multiple subscriber sessions: therefore, in this section they will be referred as “common CPR interfaces”.

The common CPR interfaces must be programmed with forwarding rules in the upstream direction. For the downstream direction, instead, the DBNG-CP may omit to program forwarding rules.

In the following, the consequences of programming or not programming downstream forwarding rules on common CPR interfaces are explained.

If a common CPRi is not programmed in terms of downstream forwarding rules by the DBNG-CP, upon receiving a connection request from an RG over the common CPRi, the DBNG-CP immediately creates a “dedicated” CPR tunnel to exchange any further control packet with the RG: this option is named “Immediate Session Creation”.

The alternative option is possible in case the DBNG-CP programs not only upstream, but also downstream forwarding rules on the common CPR interface. In this case, the DBNG-CP may adopt the so called “Delayed Session Creation” option, which consists in postponing the creation of the “dedicated” CPR tunnel, making use, during the setup, of the common CPR tunnel to deliver control packets to the RG. The postponement of the dedicated CPRi can last until the moment when the DBNG-CP assigns an address to the RG or, in case the DBNG-CP plays the role of LAC, until it identifies the L2TP tunnel towards the LNS: from that point on, the DBNG-CP uses the dedicated CPR interface to exchange control packets with the RG.

In the following subsections, the ways the DBNG-CP achieves the programming of the dedicated CPRi on the DBNG-UP are described for two different types of access: lpoE and PPPoE.

4.5.2.1 Programming of Control Packet Redirection Interface for IpoE based subscriber sessions

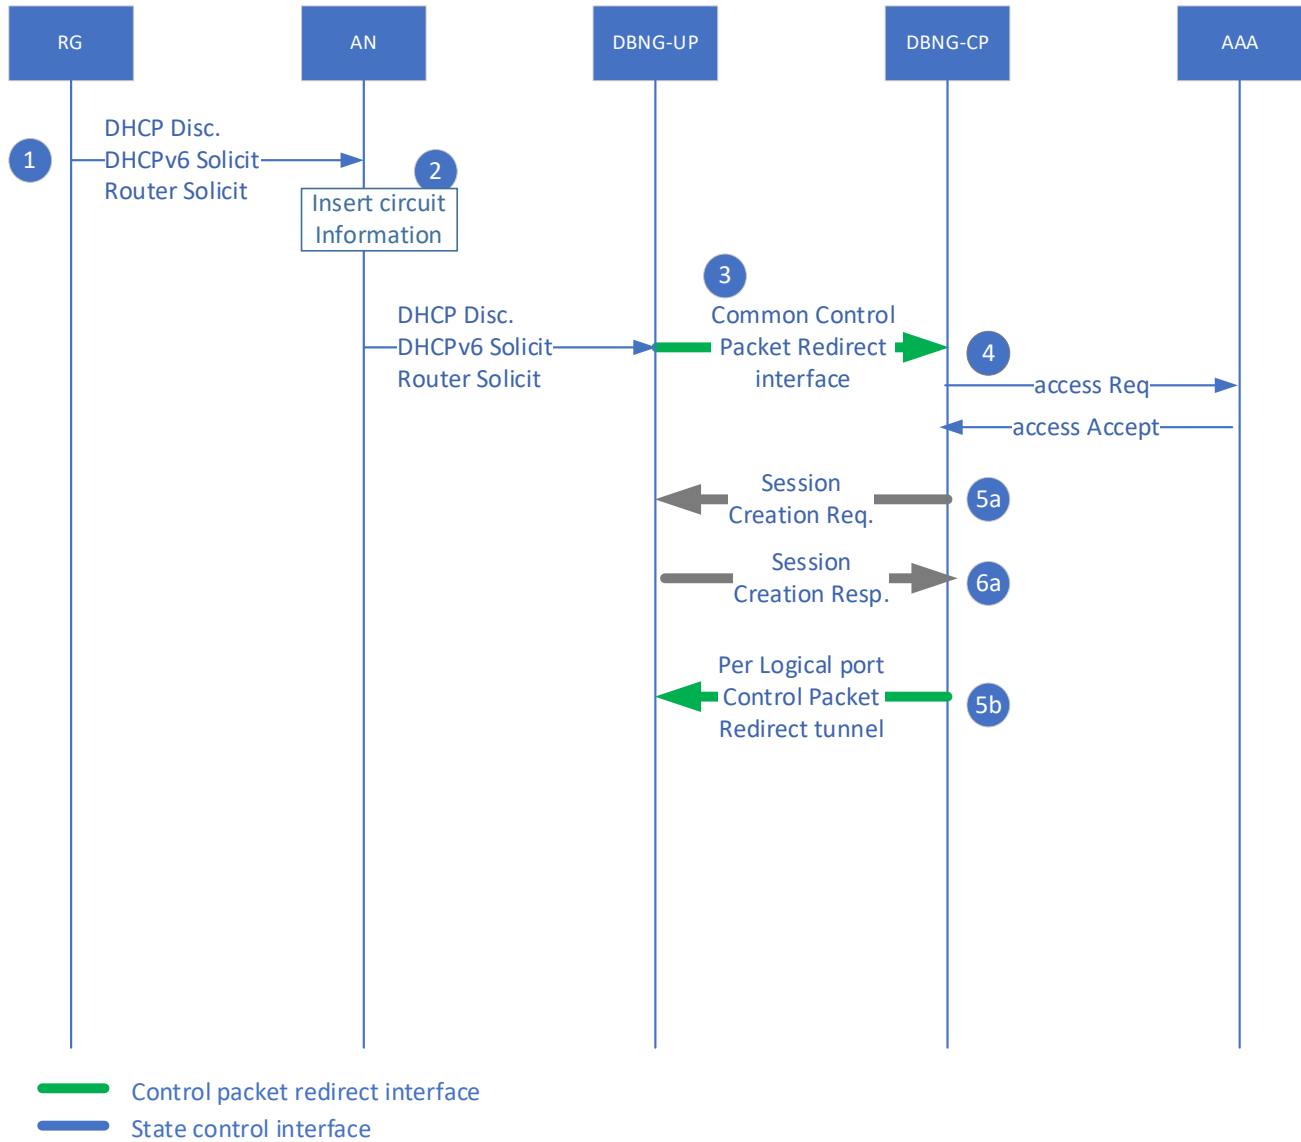


Figure 17: Programming of the Control Packet Redirection interface for IpoE based subscriber sessions

With reference to an IpoE based access (when DBNG-CP is local DHCP server), hereby are the steps to program a dedicated Control Packet Redirection interface on the DBNG-UP. Prior to step 1, the DBNG-CP has programmed at least the default CPR interface; if DBNG-UP supports per-logical-port CPR tunnel, the DBNG-CP may have also programmed a per-logical-port CPR interface. In both cases, it is up to discretion of the DBNG-CP to program on the common CPRi only an upstream PFCP forwarding rule to support Immediate Session Creation or to program both upstream and downstream PFCP forwarding rules to support Delayed Session Creation. In case of bi-directional PFCP rules, the PFCP traffic endpoint will not contain the MAC address of any subscriber.

1. An RG connecting for the first time to the network sends a control packet message. This may be a DHCPv4 discover, a DHCPv6 discover, a Router Solicit, or a data trigger packet.
2. The AN inserts circuit information onto the control message, which for example includes a circuit ID and/or a remote ID.
3. The DBNG-UP detects these as new control messages that are unrelated to currently established subscriber sessions and tunnels these control messages using the PFCP session rules present for default CPR tunnel or per logical port CPR tunnel (If no match is found in the default and/or logical port PFCP sessions, the control messages are dropped).
In the case where the default CPR tunnel is used, the DBNG-UP must also provide access interface information to the DBNG-CP, by encapsulating a NSH header to the control packet (for example the local access port information), because this information is not within the packet itself.
4. The DBNG-CP receives the control message and authenticates the subscriber with the AAA.

At this point, there are two options with regard to the creation of the dedicated CPRi via PFCP: a) immediate subscriber PFCP session establishment b) delayed subscriber PFCP session establishment.

In option a)

- 5a. After authentication succeeds, the DBNG-CP programs via PFCP the dedicated CPRi on the DBNG-UP, in order to exchange control traffic with the RG. The DBNG-CP constructs the PFCP traffic endpoint which will include VLAN tags, logical port, and MAC address.
- 6a. The DBNG-UP informs the DBNG-CP via PFCP that the dedicated CPRi is now established.

In option b), the DBNG-CP delays the establishment of the dedicated CPRi: the DBNG-CP will utilize the default or the per-logical-port bidirectional tunnel to exchange control packets with the RG.

- 5b. After authentication succeeds, control packets between DBNG-CP and RG continue to utilize the per-logical-port CPR tunnel.

Prior to address assignment (i.e., prior to the DHCP ack or the DHCPv6 reply), the DBNG-CP is expected to program the dedicated CPRi on the DBNG-UP.

4.5.2.2 Programming of the Control Packet Redirection Interface for PPPoE based subscriber sessions

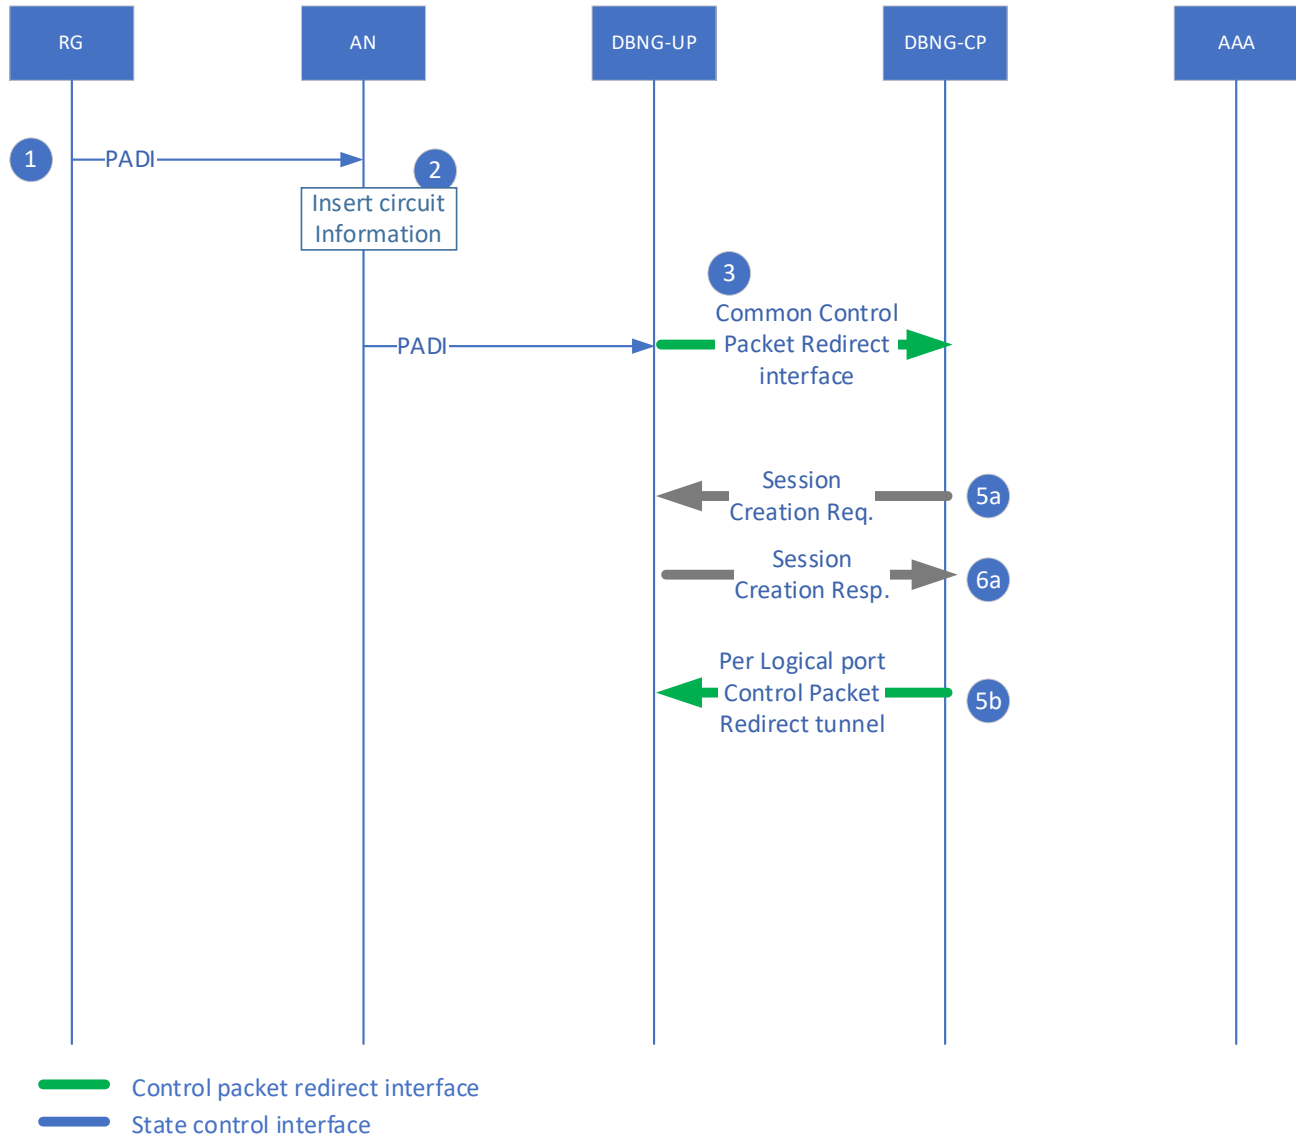


Figure 18: Programming of the Control Packet Redirection Interface for PPPoE based subscriber sessions

With reference to a PPPoE based access, hereby are the steps to program a dedicated Control Packet Redirection interface on the DBNG-UP.

Prior to step 1, the DBNG-CP has programmed at least the default CPR interface; if DBNG-UP supports per-logical-port CPR tunnel, the DBNG-CP may have also programmed a per-logical-port CPR interface. In both cases, it is up to discretion of the DBNG-CP to program on the common CPRi only an upstream PFCP forwarding rule to support Immediate Session Creation or to program both upstream and downstream PFCP forwarding rules to support Delayed Session Creation. In case of bi-directional PFCP rules, the PFCP traffic endpoint will not contain the MAC address of any subscriber.

1. An RG connecting for the first time to the network sends a PPPoE Auto Discovery Initiation (PADI) packet.
2. The AN inserts circuit information onto the control message which for example includes a circuit ID and/or a remote ID.
3. The DBNG-UP detects these as new control messages that are unrelated to currently established subscriber sessions and tunnels these control messages using the PFCP session rules present for default CPR tunnel or per logical port CPR tunnel (If no match is found in the default and/or logical port PFCP sessions, the control messages are dropped). In the case where the default CPR tunnel is used, the DBNG-UP must also provide access interface information to the DBNG-CP by encapsulating a NSH header to the first control packet (for example the local access port information) because this information is not within the packet itself.

At this point, there are two options: a) immediate subscriber PFCP session establishment b) delayed subscriber PFCP session establishment.

In option a)

- 5a. The DBNG-CP programs via PFCP the dedicated CPRi on the DBNG-UP, in order to exchange control traffic with the RG. The DBNG-CP constructs the PFCP traffic endpoint which will include VLAN tags, logical port, PPPoE session ID, and MAC address.
- 6a. The DBNG-UP informs the DBNG-CP via PFCP that the dedicated CPRi is now established.

In option b)

- 5b. Starting from the PADO, control packets between DBNG-CP and RG continue to utilize the per-logical-port bidirectional CPR tunnel.

Prior to address assignment (i.e., prior to IPCP ack or Ipv6CP ack) or prior to the L2TP tunnel identification or creation in case of LAC, the DBNG-CP is expected to program the dedicated CPRi on the DBNG-UP.

4.5.3 External DHCP Server Control Packet Redirection Rule

An external DHCPv4 or DHCPv6 Server can be reachable to the DBNG-CP directly via the cp-d interface or via the DBNG-UP. Configuration will allow the DBNG-CP to determine how the external server is reachable.

When the DBNG is configured to operate as a DHCP relay or relay-proxy or DHCPv6 relay in which the external server is connected to an interface on the UP, after the association between DBNG-UP and DBNG-CP is formed, forwarding rules are programmed from the DBNG-CP to the DBNG-UP for DHCP/DHCPv6 control packet redirection rules. The DBNG-CP programs control packet redirection rules to instruct the DBNG-UP to redirect specific types of DHCP or DHCPv6 control packets received from the external DHCP/DHCPv6 server to the DBNG-CP through the server CPR interface for the upstream direction. The DBNG-CP also programs an additional downstream control packet redirection rule to instruct the DBNG-UP to forward DHCP or DHCPv6 control packets from the DBNG-CP through the server CPR interface to the external DHCP server.

The following diagram depicts the positioning of the server Control Packet Redirect interface within the IpoE DHCPv4 or IpoE DHCPv6 call flows, where the existing common and session Control Packet Redirect interface and common Control Packet Redirect Downstream interface (for delayed session creation) continue to be used exactly as specified throughout this document.

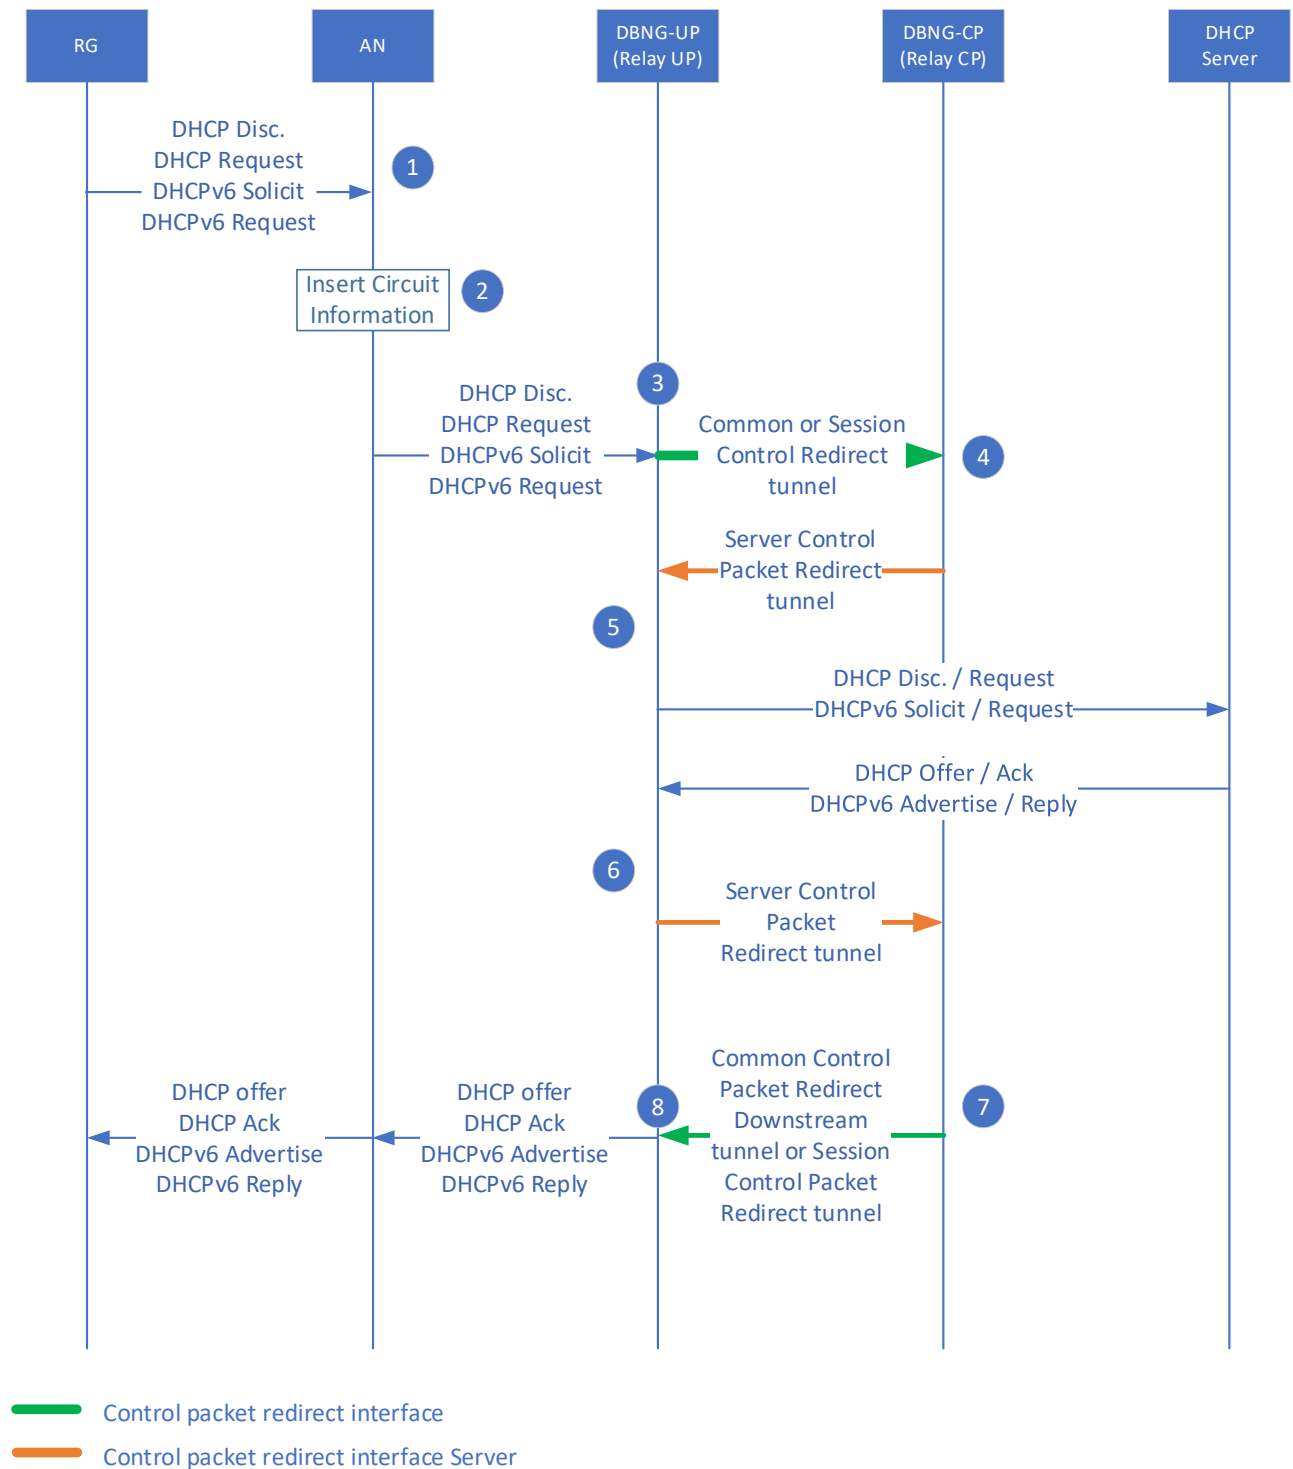


Figure 19: External DHCP server Control Packet Redirection rule

Steps (When the DBNG-CP is configured to work as DHCP v4/6 Relay and the connectivity to the External DHCP Server is via DHCP-UP):

1. All new IpoE DHCP or DHCPv6 RGs first connecting to the network, send a control message packet, specifically DHCPv4 Discover or DHCPv6 Solicit. As part of negotiation, the subsequent upstream control message sent by the IpoE DHCP or DHCPv6 RGs will be a DHCPv4 Request or DHCPv6 Request.
2. The AN inserts circuit information onto the control message which can include: a circuit ID, a remote ID, a Line ID, and/or an interface ID.
3. The DBNG-UP detects these as new control messages that are unrelated to currently established subscriber sessions and tunnels these control messages using the PFCP session rules present for default CPR tunnel or per logical port CPR tunnel (If no match is found in the default and/or logical port PFCP sessions, the control messages are dropped). The DBNG-UP must also provide access interface information as metadata (for example port information) to the DBNG-CP because this information might not be within the packet itself.
4. The DBNG-CP receives the control message and its context as metadata, together providing a session context.
5. After processing the control packet as a stateful relay, DBNG-CP sends the (modified) upstream DHCP control packet to the external DHCP server through DBNG-UP using the server Control Packet Redirect tunnel.
6. The DBNG-UP detects downstream DHCP and DHCPv6 control messages received from the external server in response to the prior upstream control message and redirects the control packets to DBNG-CP through the server Control Packet Redirect tunnel.
7. After processing the control packet as a stateful relay, DBNG-CP sends the (modified) downstream DHCP control packet (responses) to the RG through the DBNG-UP using either the session Control Packet Redirect tunnel or the common downstream control packet redirect tunnel, if the per session redirect tunnel is not setup by this time.
8. DBNG-UP forwards control packets to RG by matching the common downstream redirect tunnel.

4.5.4 IpoE DHCPv4 Immediate Session Creation

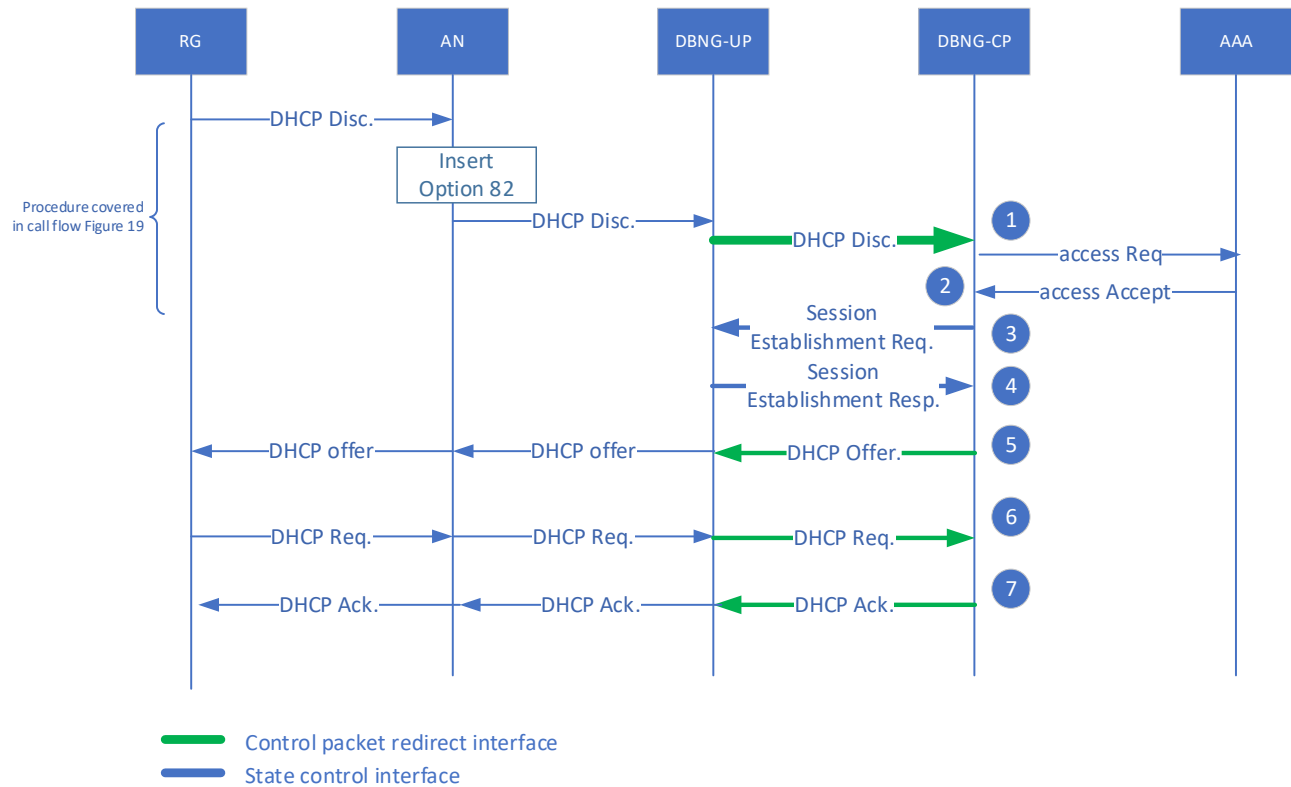


Figure 20: IpoE DHCPv4 immediate session creation call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule.

1. The DBNG-CP triggers an Access-Request to authenticate the RG.
2. The AAA successfully authenticates the RG and replies to the DBNG-CP with an Access-Accept.
3. The DBNG-CP assigns the IP address to the RG which is obtained through either the local address server or returned by AAA as one of the attributes. At this point the DBNG-CP can send a Session Establishment Request to create new packet forwarding states for the data packet. This can update the data plane state.
4. The DBNG-UP sends a Session Establishment Response back to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward the subscriber’s IP data packets.
5. The DHCP Offer from the DBNG-CP is sent back to the RG through the DBNG-UP utilizing the dedicated CPR Interface.
6. The DHCP Request is sent from the RG through the DBNG-UP utilizing a dedicated session control packet redirect tunnel.
7. The DHCP process completes by sending the DHCP acknowledgement. The DBNG-CP forwards the DHCP Ack through the dedicated session control packet redirect tunnel back to the RG though the DBNG-UP.

4.5.5 IpoE DHCPv4 Delayed Session Creation

The DHCPv4 server can either be connected to the DBNG-UP or directly to the DBNG-CP. Note that the call flows apply to both relay and relay-proxy modes.

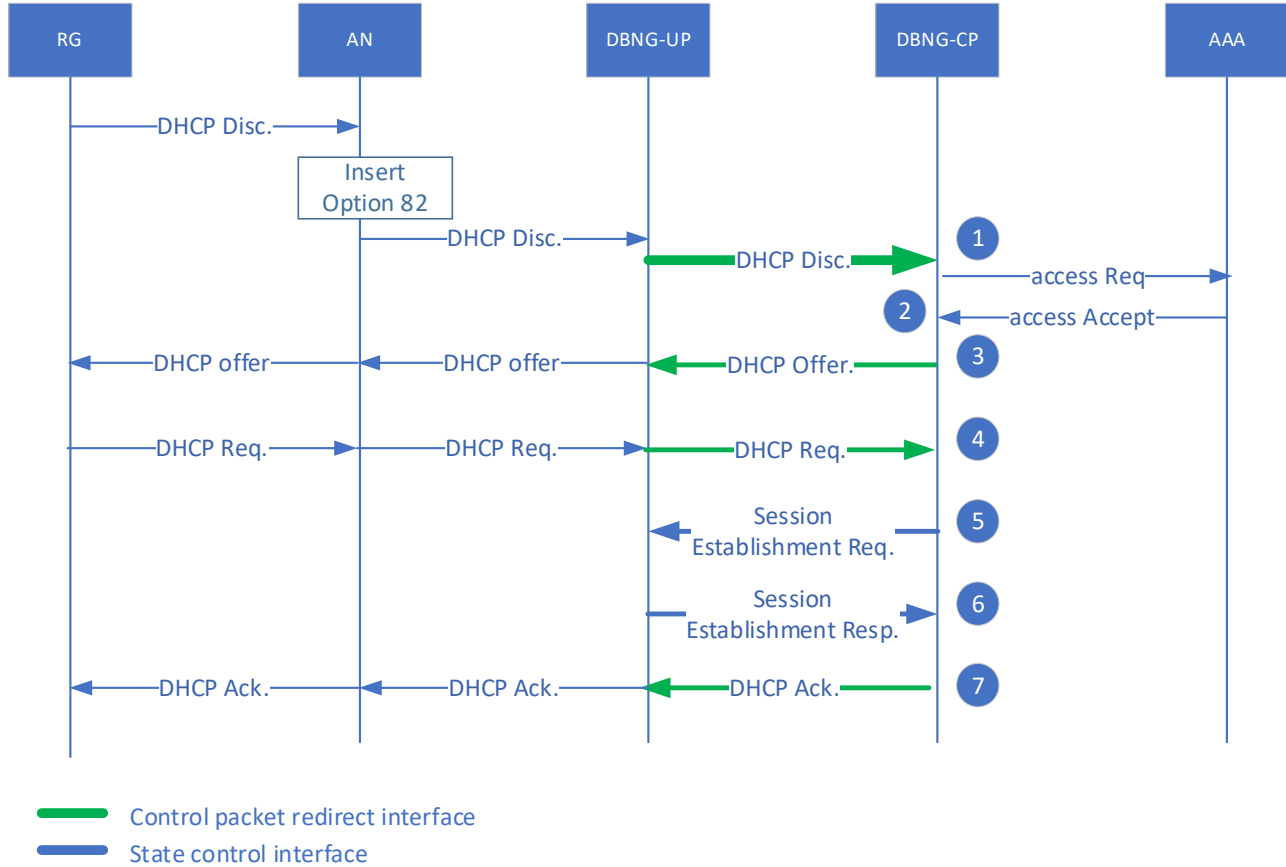


Figure 21: IpoE DHCPv4 delayed session creation call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions.

1. The DBNG-CP triggers an Access-Request to authenticate the RG
2. AAA successfully authenticates the RG and replies with access accept
3. The DBNG-CP assigns the IP address to the RG which is obtained through either the local address server or returned by AAA as one of the attributes. The DHCP offer from the DBNG-CP is sent back to RG through the DBNG-UP utilizing default/common redirect interface tunnel.
4. The DHCP Request is sent from the RG through the DBNG-UP utilizing the default/common control packet redirect tunnel
5. At this stage, DBNG-CP can send subscriber session establishment request containing session specific control packet redirection rules and data packet forwarding rules for this subscriber.
6. The DBNG-UP sends a Session Establishment Response back to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward the subscriber's IP data packets
7. The DHCP process completes by sending the DHCP acknowledgement. The DBNG-CP forwards the DHCP Ack through the dedicated session control packet redirect tunnel back to the RG through the DBNG-UP

4.5.6 IpoE DHCPv4 Relay Call Flows

The DHCPv4 relay server can either be connected to the DBNG-UP or directly to the DBNG-CP via the cp-d interface.

4.5.6.1 IpoE DHCPv4 Relay Immediate Session Creation (via DBNG-UP)

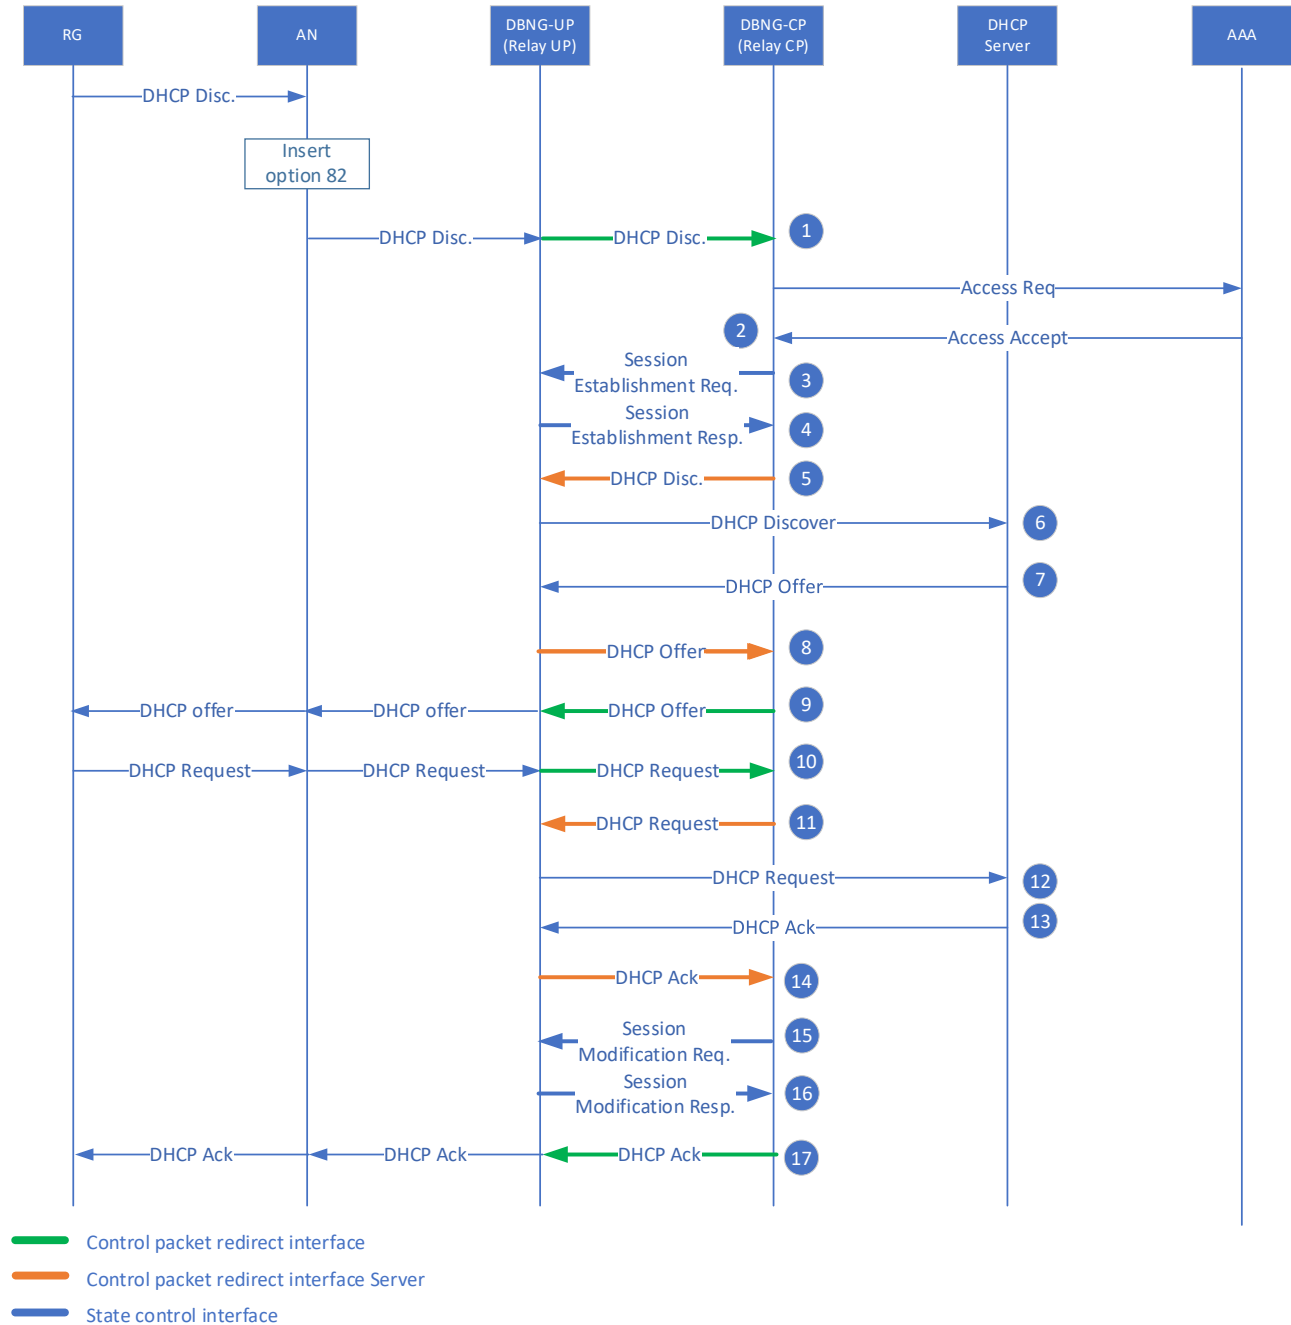


Figure 22: IpoE DHCPv4 Relay Immediate Session Creation (via DBNG-UP) call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions in Figure 18. Also, the call flow in section 4.5.3 and Figure 20 covers the server CPR rule for exchanging DHCP control packets between the DBNG-CP and external DHCP server through the DBNG-UP.

1. The DBNG-CP triggers an Access-Request to authenticate the RG
2. AAA successfully authenticates the RG and replies with Access-Accept
3. At this stage, DBNG-CP sends the subscriber Session Establishment Request, containing session specific control packet redirection rules for this subscriber.
4. The DBNG-UP sends a Session Establishment Response to the DBNG-CP, informing that the states are installed.
5. The DBNG-CP converts the broadcast DHCP Discover to unicast, including setting the relay agent IP address (GIADDR) to the DBNG-UP local Ipv4 address of the logical port (V interface). The DBNG-CP sends the DHCP Discover to the DBNG-UP utilizing the server control packet redirect tunnel. Note that the DBNG-CP may learn of the DBNG-UP local Ipv4 address during UP configuration via Mi.
6. The DBNG-UP forwards the unicast DHCP Discover to the external DHCP server.
7. The external DHCP server responds with a DHCP Offer packet that is directed to the DBNG DHCP relay.
8. The DBNG-UP detects the downstream DHCP Offer packet, containing the assigned IP address, from the external server and redirects the packet to the DBNG-CP through the server Control Packet Redirect tunnel.
9. The DBNG-CP converts the DHCP Offer to a broadcast or unicast packet, per the broadcast bit in the prior Discover packet, and sends it to the RG through the DBNG-UP, utilizing the session control packet redirect tunnel.
10. The DHCP Request is sent from the RG to the DBNG-CP through the DBNG-UP utilizing the session control packet redirect tunnel.
11. The DBNG-CP converts the broadcast DHCP Request to unicast, including setting the relay agent IP address (GIADDR) to the DBNG-UP local Ipv4 address of the subscriber logical port (V interface). The DBNG-CP sends the DHCP Request to the DBNG-UP utilizing the server control packet redirect tunnel.
12. The DBNG-UP forwards the unicast DHCP Request to the external DHCP server.
13. The external DHCP server responds with a DHCP Ack packet that is directed to the DBNG DHCP relay.
14. The DBNG-UP detects the downstream DHCP Ack packet from the external server and redirects the packet to the DBNG-CP through the server Control Packet Redirect tunnel.
15. The DBNG-CP sends a subscriber Session Modification Request to update the traffic forwarding rules to match on the subscriber's Ipv4 address. Note that this step may occur as early as after step 10, but the implications of this approach should be considered.
16. The DBNG-UP sends a Session Modification Response to the DBNG-CP, informing that the forwarding rules have been updated, and the DBNG-UP is ready to forward the subscriber's IP data packets.
17. The DBNG-CP converts the DHCP Ack to a broadcast or unicast packet, per the broadcast bit in the prior Discover packet, and sends it to the RG through the DBNG-UP utilizing the session control packet redirect tunnel.

4.5.6.2 IpoE DHCPv4 Relay Immediate Session Creation (via DBNG-CP)

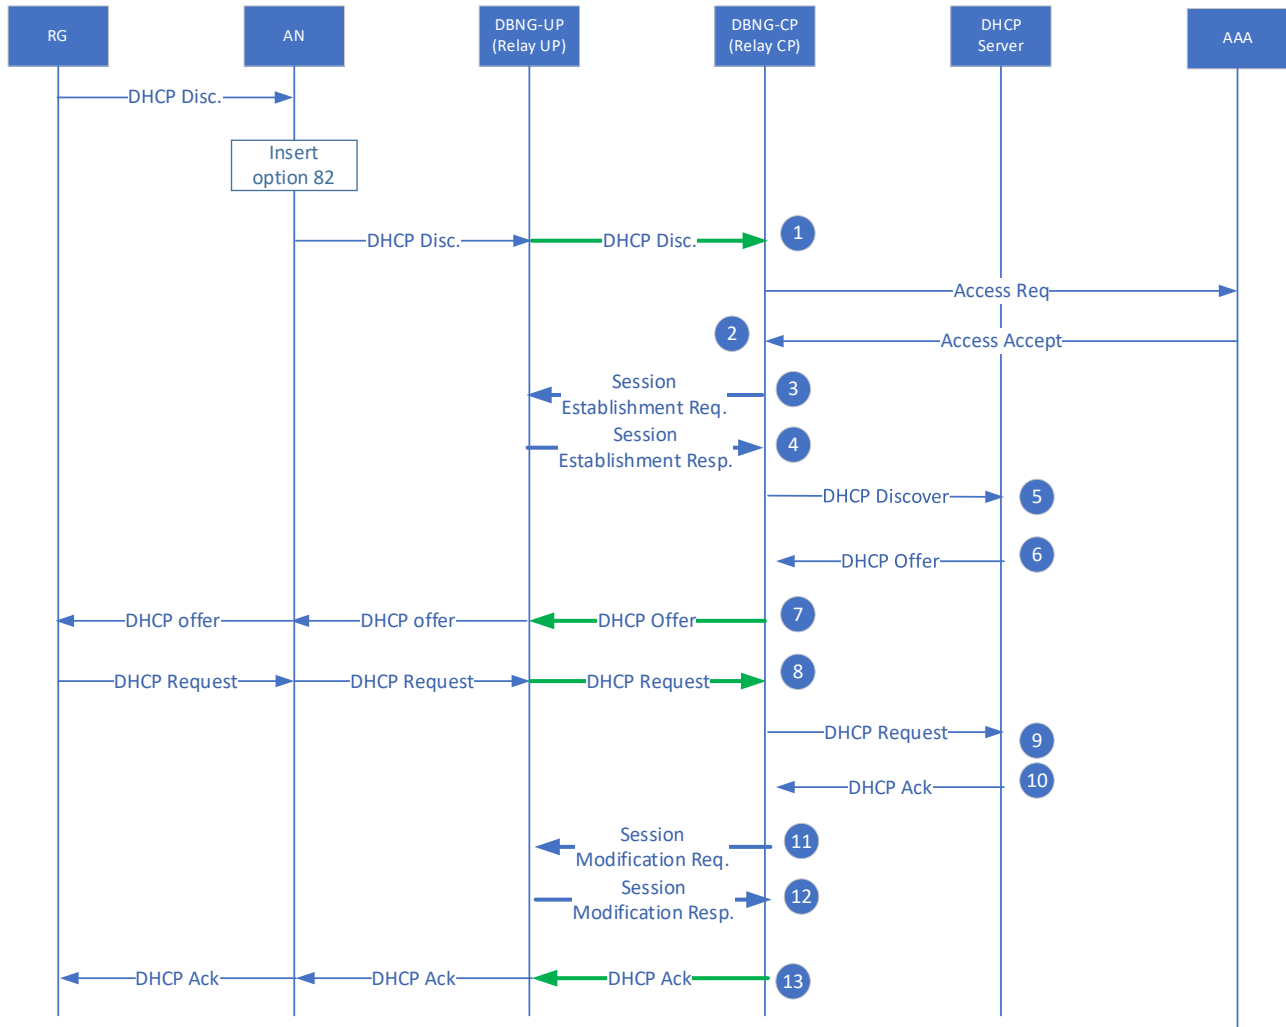


Figure 23: IpoE DHCPv4 Relay Immediate Session Creation (via DBNG-CP) call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions in Figure 18.

1. The DBNG-CP triggers an Access-Request to authenticate the RG
2. AAA successfully authenticates the RG and replies with Access-Accept
3. At this stage, DBNG-CP sends the subscriber Session Establishment Request, containing session specific control packet redirection rules for this subscriber.
4. The DBNG-UP sends a Session Establishment Response to the DBNG-CP, informing that the states are installed.
5. The DBNG-CP converts the broadcast DHCP Discover to unicast, including setting the relay agent IP address (GIADDR) to the DBNG-UP local Ipv4 address of the logical port (V interface). The

- DBNG-CP sends the DHCP Discover to the DHCP server. Note that the DBNG-CP may learn of the DBNG-UP local Ipv4 address during UP configuration via Mi.
6. The external DHCP server responds with a DHCP Offer packet that is directed to the DBNG DHCP relay.
 7. The DBNG-CP converts the DHCP Offer to a broadcast or unicast packet, per the broadcast bit in the prior Discover packet, and sends it to the RG through the DBNG-UP, utilizing the session control packet redirect tunnel.
 8. The DHCP Request is sent from the RG to the DBNG-CP through the DBNG-UP utilizing the session control packet redirect tunnel.
 9. The DBNG-CP converts the broadcast DHCP Request to unicast, including setting the relay agent IP address (GIADDR) to the DBNG-UP local Ipv4 address of the subscriber logical port (V interface). The DBNG-CP sends the DHCP Request to the DHCP server.
 10. The external DHCP server responds with a DHCP Ack packet that is directed to the DBNG DHCP relay.
 11. The DBNG-CP sends a subscriber Session Modification Request to update the traffic forwarding rules to match on the subscriber's Ipv4 address. Note that this step may occur as early as after step 8, but the implications of this approach should be considered.
 12. The DBNG-UP sends a Session Modification Response to the DBNG-CP, informing that the forwarding rules have been updated, and the DBNG-UP is ready to forward the subscriber's IP data packets.
 13. The DBNG-CP converts the DHCP Ack to a broadcast or unicast packet, per the broadcast bit in the prior Discover packet, and sends it to the RG through the DBNG-UP utilizing the session control packet redirect tunnel.

4.5.6.3 IpoE DHCPv4 Relay Delayed Session Creation (via DBNG-UP)

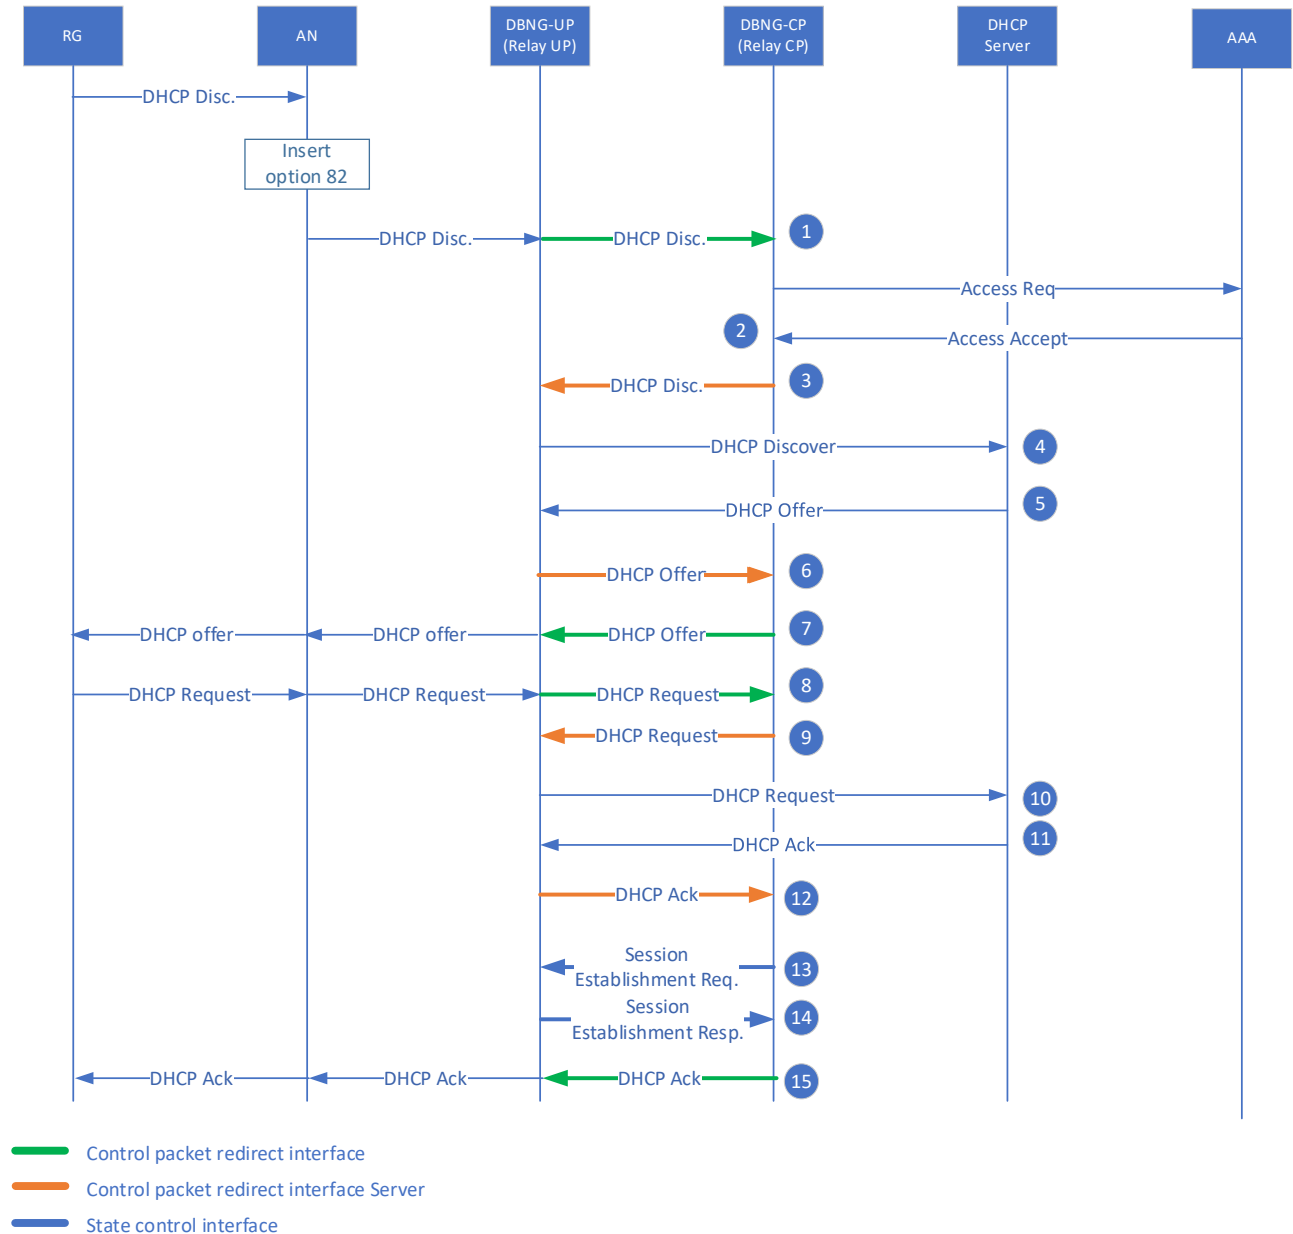


Figure 24: IpoE DHCPv4 Relay Delayed Session Creation (via DBNG-UP) call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions in Figure 17. Also, the call flow in section 4.5.3 and Figure 19 covers the server CPR rule for exchanging DHCP control packets between the DBNG-CP and external DHCP server through the DBNG-UP.

1. The DBNG-CP triggers an Access-Request to authenticate the RG
2. AAA successfully authenticates the RG and replies with Access-Accept
3. The DBNG-CP converts the broadcast DHCP Discover to unicast, including setting the relay agent IP address (GIADDR) to the DBNG-UP local Ipv4 address of the logical port (V interface). The DBNG-CP sends the DHCP Discover to the DBNG-UP utilizing the server control packet redirect

- tunnel. Note that the DBNG-CP may learn of the DBNG-UP local Ipv4 address during UP configuration via Mi.
4. The DBNG-UP forwards the unicast DHCP Discover to the external DHCP server.
 5. The external DHCP server responds with a DHCP Offer packet that is directed to the DBNG DHCP relay.
 6. The DBNG-UP detects the downstream DHCP Offer packet, containing the assigned IP address, from the external server and redirects the packet to the DBNG-CP through the server Control Packet Redirect tunnel.
 7. The DBNG-CP converts the DHCP Offer to a broadcast or unicast packet, per the broadcast bit in the prior Discover packet, and sends it to the RG through the DBNG-UP, utilizing the default/common redirect interface tunnel.
 8. The DHCP Request is sent from the RG to the DBNG-CP through the DBNG-UP utilizing the default/common control packet redirect tunnel.
 9. The DBNG-CP converts the broadcast DHCP Request to unicast, including setting the relay agent IP address (GIADDR) to the DBNG-UP local Ipv4 address of the subscriber logical port (V interface). The DBNG-CP sends the DHCP Request to the DBNG-UP utilizing the server control packet redirect tunnel.
 10. The DBNG-UP forwards the unicast DHCP Request to the external DHCP server.
 11. The external DHCP server responds with a DHCP Ack packet that is directed to the DBNG DHCP relay.
 12. The DBNG-UP detects the downstream DHCP Ack packet from the external server and redirects the packet to the DBNG-CP through the server Control Packet Redirect tunnel.
 13. At this stage, DBNG-CP sends the subscriber Session Establishment Request, containing session specific control packet redirection rules and data packet forwarding rules for this subscriber.
 14. The DBNG-UP sends a Session Establishment Response to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward the subscriber's IP data packets.
 15. The DBNG-CP converts the DHCP Ack to a broadcast or unicast packet, per the broadcast bit in the prior Discover packet, and sends it to the RG through the DBNG-UP utilizing the session control packet redirect tunnel.

Note: Session Establishment Request and Response can occur as soon as step 8 take place, but there are implications of accepting the RG requested IP address without the server validation.

4.5.6.4 IpoE DHCPv4 Relay Delayed Session Creation (via DBNG-CP)

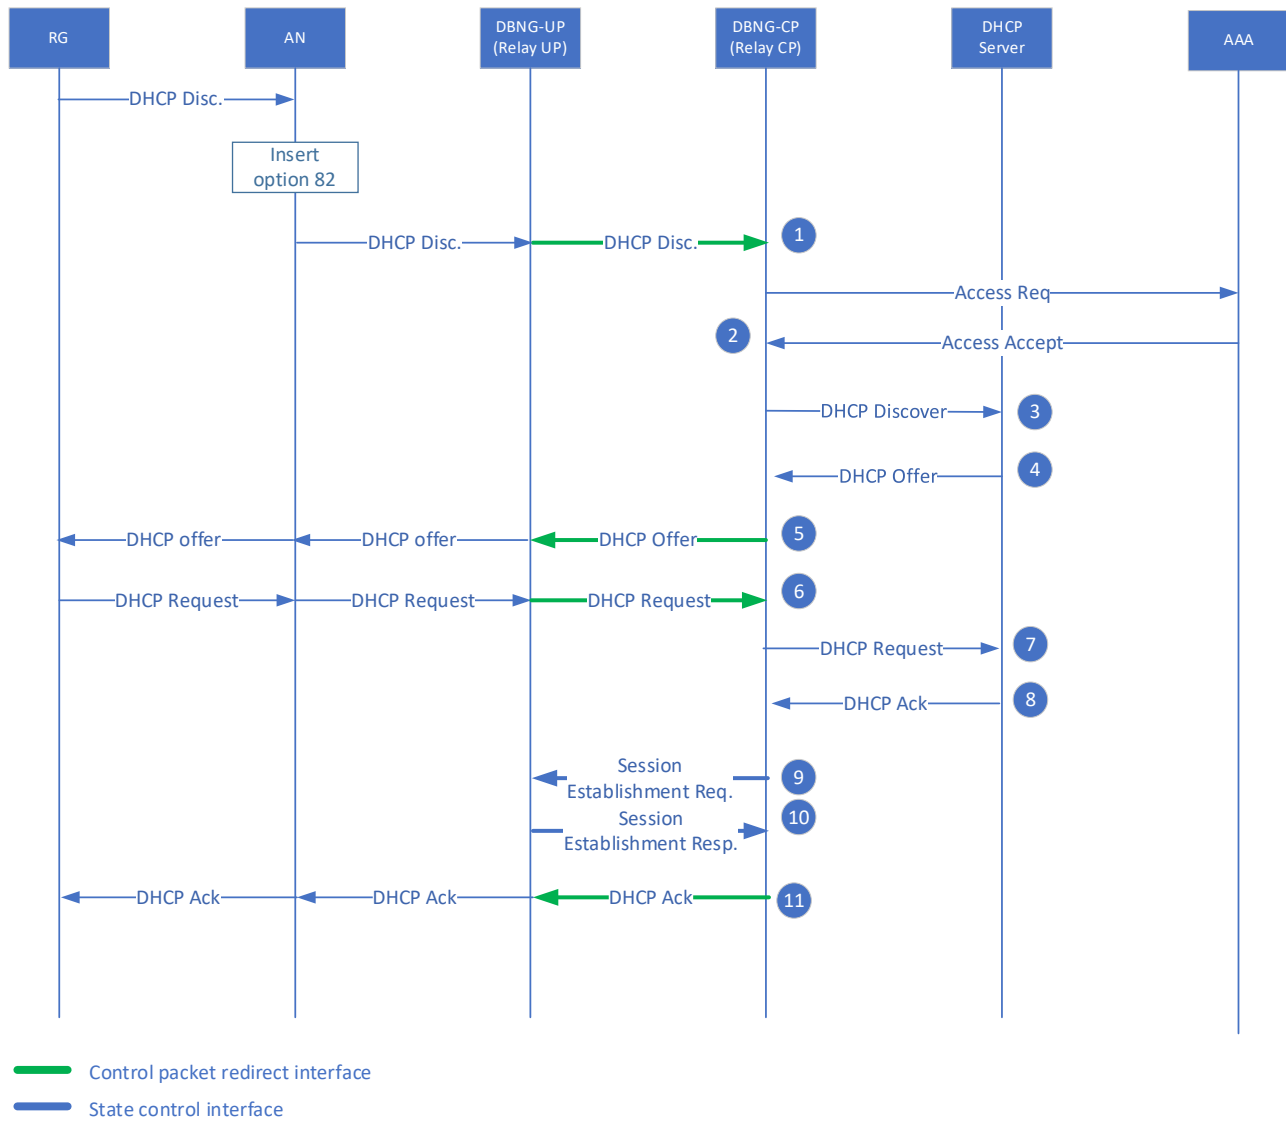


Figure 25: IpoE DHCPv4 Relay Delayed Session Creation (via DBNG-CP) call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions in Figure 17.

1. The DBNG-CP triggers an Access-Request to authenticate the RG
2. AAA successfully authenticates the RG and replies with Access-Accept
3. The DBNG-CP converts the broadcast DHCP Discover to unicast, including setting the relay agent IP address (GIADDR) to the DBNG-UP local Ipv4 address of the logical port (V interface). The DBNG-CP sends the DHCP Discover to the DHCP server via a local interface. Note that the DBNG-CP may learn of the DBNG-UP local Ipv4 address during UP configuration via Mi.
4. The external DHCP server responds with a DHCP Offer packet that is directed to the DBNG DHCP relay.

5. The DBNG-CP converts the DHCP Offer to a broadcast or unicast packet, per the broadcast bit in the prior Discover packet, and sends it to the RG through the DBNG-UP, utilizing the default/common redirect interface tunnel.
6. The DHCP Request is sent from the RG to the DBNG-CP through the DBNG-UP utilizing the default/common control packet redirect tunnel.
7. The DBNG-CP converts the broadcast DHCP Request to unicast, including setting the relay agent IP address (GIADDR) to the DBNG-UP local Ipv4 address of the subscriber logical port (V interface). The DBNG-CP sends the DHCP Request to the DHCP server via a local control/management interface.
8. The external DHCP server responds with a DHCP Ack packet that is directed to the DBNG DHCP relay.
9. At this stage, DBNG-CP sends the subscriber Session Establishment Request, containing session specific control packet redirection rules and data packet forwarding rules for this subscriber.
10. The DBNG-UP sends a Session Establishment Response to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward the subscriber's IP data packets.
11. The DBNG-CP converts the DHCP Ack to a broadcast or unicast packet, per the broadcast bit in the prior Discover packet, and sends it to the RG through the DBNG-UP utilizing the session control packet redirect tunnel.

Note: Session Establishment Request and Response can occur as soon as step 8 take place, but there are implications of accepting the RG requested IP address without the server validation.

4.5.7 IpoE DHCPv6 Immediate Session Creation

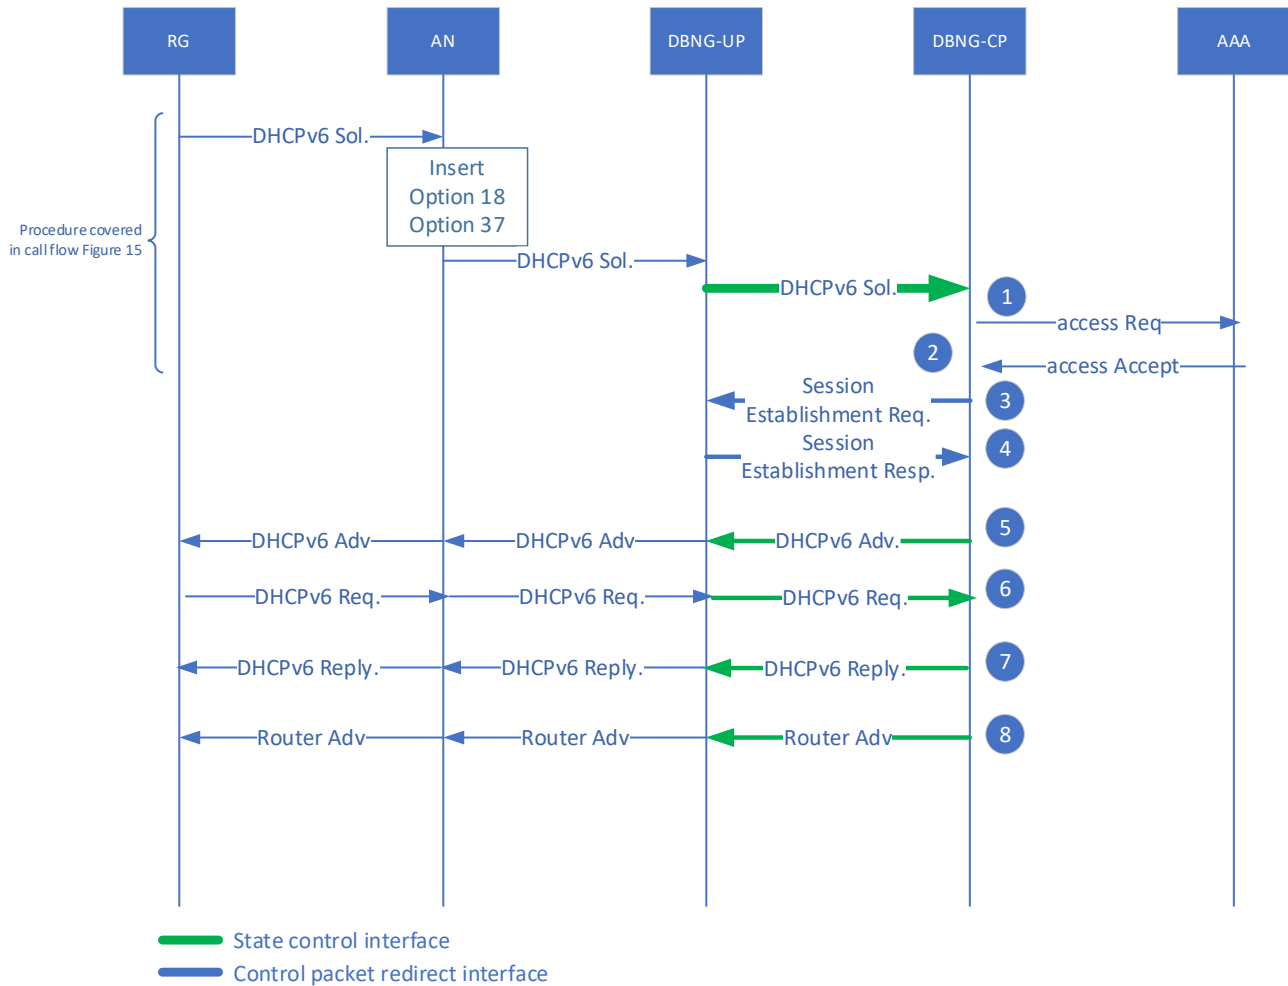


Figure 26: IpoE DHCPv6 immediate session creation call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common control packet redirection rule.

1. The DBNG-CP triggers an Access-Request to authenticate the RG.
2. The AAA successfully authenticates the RG and replies to the DBNG-CP with an Access-Accept.
3. The DBNG-CP assigns the IP address to the RG which is obtained through either the local address server or returned by AAA as one of the attributes. At this point the DBNG-CP can send a Session Establishment Request to create new packet forwarding states for the data packet. This updates the data plane state.
4. The DBNG-UP sends a Session Establishment Response back to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward the subscriber's IP data packets.
5. The DHCPv6 Advertisement from the DBNG-CP is sent back to the RG through the DBNG-UP utilizing the CPR Interface.
6. The DHCPv6 Request is sent from the RG through the DBNG-UP utilizing a dedicated session control packet redirect tunnel.
7. The DHCPv6 process completes by sending the DHCPv6 reply. The DBNG-CP sends the DHCPv6 Reply through the dedicated session control packet redirect tunnel back to the RG though the DBNG-UP.
8. DBNG-CP informs the RG the default gateway (Link Local address).

4.5.8 IpoE DHCPv6 Delayed Session Creation

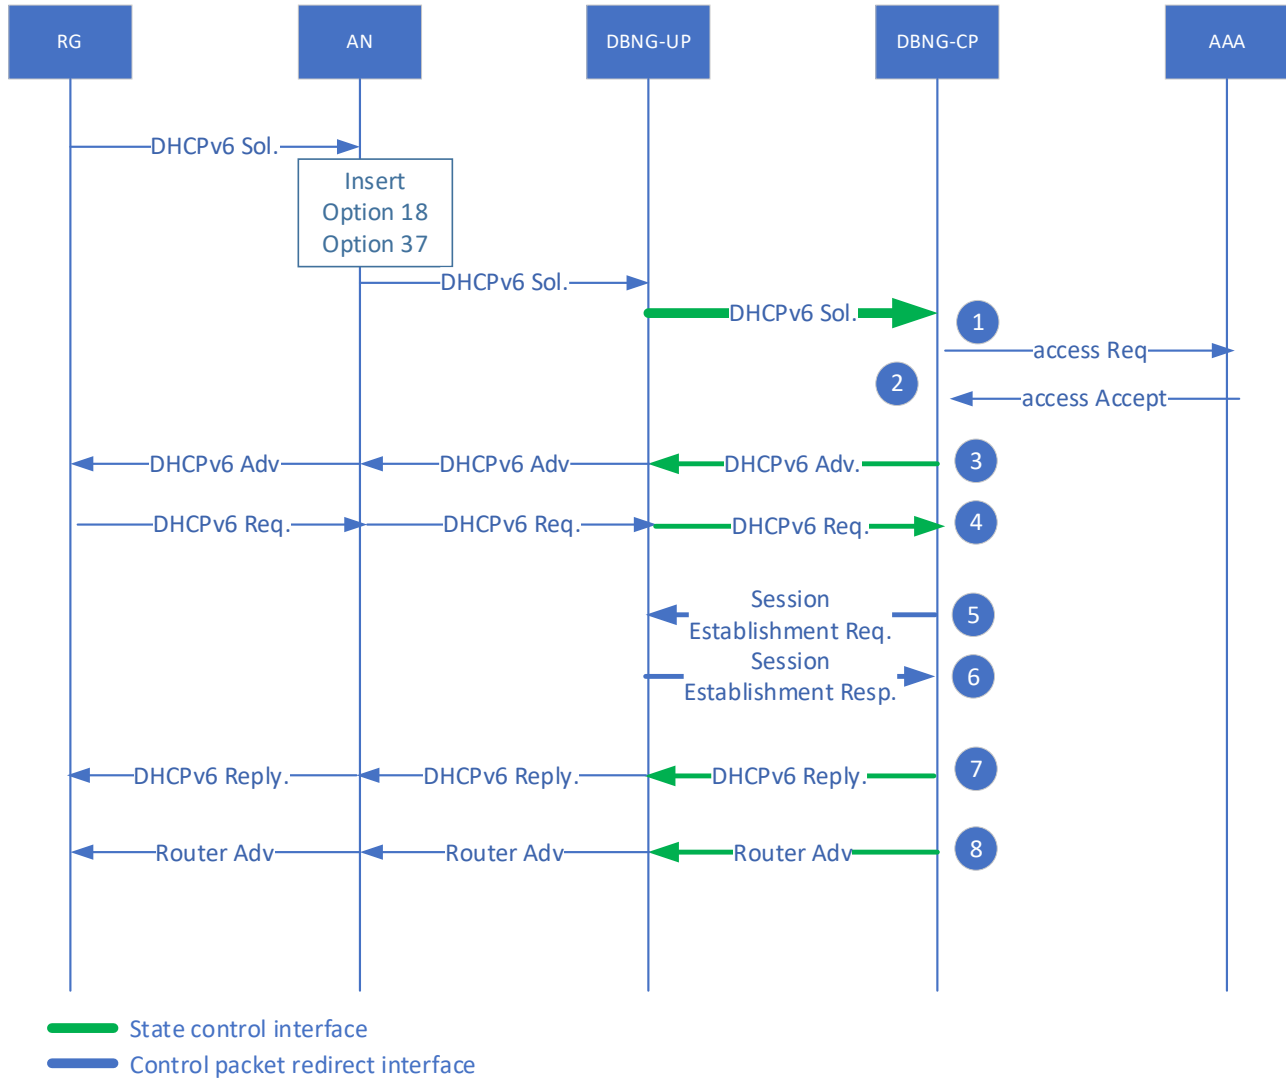


Figure 27: IpoE DHCPv6 delayed session creation call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions.

1. The DBNG-CP triggers an Access-Request to authenticate the RG.
2. The AAA successfully authenticates the RG and replies to the DBNG-CP with an Access-Accept.
3. The DBNG-CP assigns the IP address to the RG which is obtained through either the local address server or returned by AAA as one of the attributes. DHCPv6 Adv. From DBNG-CP is sent back to RG utilizing the downstream default redirect tunnel interface.
4. The DHCPv6 Request is sent from the RG through the DBNG-UP utilizing default CPRi.
5. At this point the DBNG-CP send a Session Establishment Request to create new packet forwarding states for the data packet. This updates the data plane state.

6. The DBNG-UP sends a Session Establishment Response back to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward the subscriber's IP data packets.
7. The DHCPv6 process completes by sending the DHCPv6 reply. The DBNG-CP sends the DHCPv6 Reply through the dedicated session control packet redirect tunnel back to the RG through the DBNG-UP.
8. DBNG-CP informs the RG the default gateway (Link Local Address).

4.5.9 IpoE DHCPv6 Relay Call Flows

The DHCPv6 server can either be connected to the DBNG-UP via the A10 interface or directly to the DBNG-CP via the B interface.

4.5.9.1 IpoE DHCPv6 Relay Immediate Session Creation (via DBNG-UP)

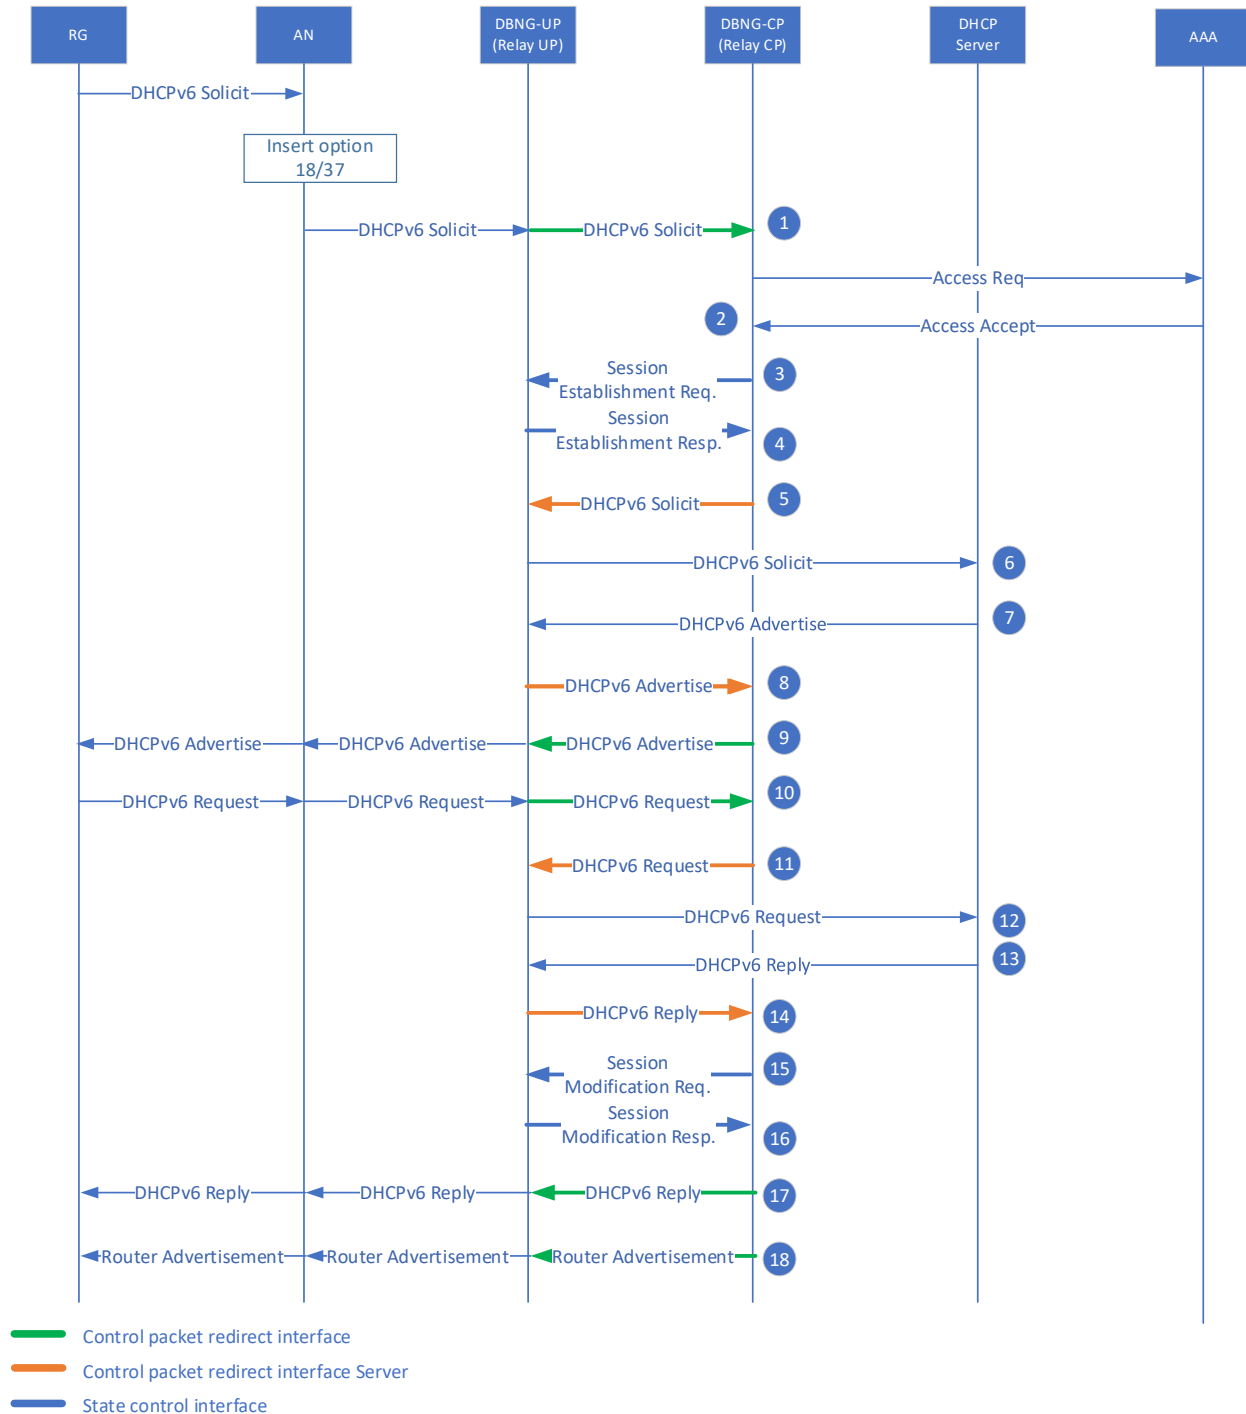


Figure 28: IpoE DHCPv6 Relay Immediate Session Creation (via DBNG-UP) call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions in Figure 18. Also, the call flow in section 4.5.3 and Figure 20 covers the server CPR rule for exchanging DHCP control packets between the DBNG-CP and external DHCPv6 server through the DBNG-UP.

1. The DBNG-CP triggers an Access-Request to authenticate the RG
2. AAA successfully authenticates the RG and replies with Access-Accept
3. At this stage, DBNG-CP sends the subscriber Session Establishment Request, containing session specific control packet redirection rules for this subscriber.
4. The DBNG-UP sends a Session Establishment Response to the DBNG-CP, informing that the states are installed.
5. The DBNG-CP encapsulates the DHCPv6 Solicit with a RELAY-FORW header, setting the link address to the DBNG-UP local Ipv6 address of the logical port (V interface), and converts the multicast DHCPv6 Solicit to unicast. The DBNG-CP sends the DHCPv6 Solicit to the DBNG-UP utilizing the server control packet redirect tunnel. Note that the DBNG-CP may learn of the DBNG-UP local Ipv6 address during UP configuration via Mi.
6. The DBNG-UP forwards the DHCPv6 Solicit to the external DHCPv6 server.
7. The external DHCPv6 server responds with a DHCPv6 Advertise packet that is directed to the DBNG DHCPv6 relay.
8. The DBNG-UP detects the downstream DHCPv6 Advertise packet, containing the assigned Ipv6 address(es), from the external server and redirects the packet to the DBNG-CP through the server Control Packet Redirect tunnel.
9. The DBNG-CP decapsulates the RELAY-REPLY header from the DHCPv6 Advertise packet before sending it to the RG through the DBNG-UP, utilizing the session control packet redirect tunnel.
10. The DHCPv6 Request is sent from the RG to the DBNG-CP through the DBNG-UP utilizing the session control packet redirect tunnel.
11. The DBNG-CP encapsulates the DHCPv6 Request with a RELAY-FORW header, setting the link address to the DBNG-UP local Ipv6 address of the subscriber logical port (V interface), and converts the multicast DHCPv6 Request to unicast. The DBNG-CP sends the DHCPv6 Request to the DBNG-UP utilizing the server control packet redirect tunnel.
12. The DBNG-UP forwards the DHCPv6 Request to the external DHCPv6 server.
13. The external DHCPv6 server responds with a DHCPv6 Reply packet that is directed to the DBNG DHCPv6 relay.
14. The DBNG-UP detects the downstream DHCPv6 Reply packet from the external server and redirects the packet to the DBNG-CP through the server Control Packet Redirect tunnel.
15. The DBNG-CP sends a subscriber Session Modification Request to update the traffic forwarding rules to match on the subscriber's Ipv6 address and/or prefixes. Note that this step may occur as early as after step 10, but the implications of this approach should be considered.
16. The DBNG-UP sends a Session Modification Response to the DBNG-CP, informing that the forwarding rules have been updated, and the DBNG-UP is ready to forward the subscriber's Ipv6 data packets.
17. The DBNG-CP decapsulates the RELAY-REPLY header from the DHCPv6 Reply packet before sending it to the RG through the DBNG-UP, utilizing the session control packet redirect tunnel.
18. DBNG-CP informs the RG of the default gateway Link Local Address (LLA).

4.5.9.2 IpoE DHCPv6 Relay Immediate Session Creation (via DBNG-CP)

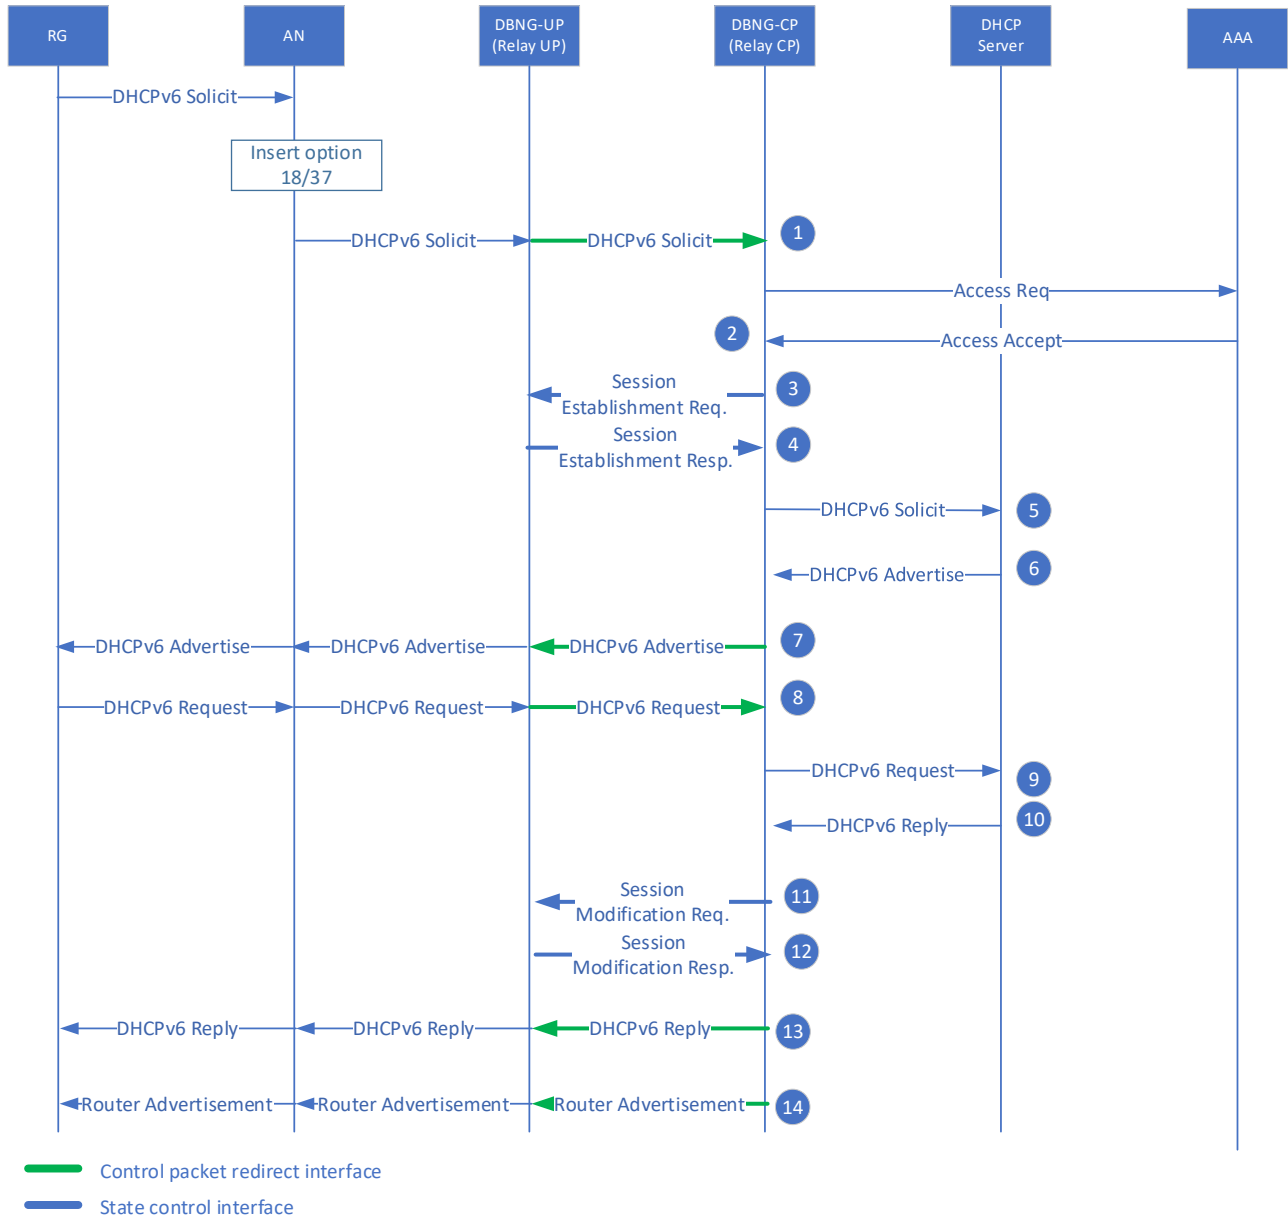


Figure 29: IpoE DHCPv6 Relay Immediate Session Creation (via DBNG-CP) call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions in Figure 18.

1. The DBNG-CP triggers an Access-Request to authenticate the RG
2. AAA successfully authenticates the RG and replies with Access-Accept
3. At this stage, DBNG-CP sends the subscriber Session Establishment Request, containing session specific control packet redirection rules for this subscriber.

4. The DBNG-UP sends a Session Establishment Response to the DBNG-CP, informing that the states are installed.
5. The DBNG-CP encapsulates the DHCPv6 Solicit with a RELAY-FORW header, setting the link address to the DBNG-UP local Ipv6 address of the logical port (V interface), and converts the multicast DHCPv6 Solicit to unicast. The DBNG-CP sends the DHCPv6 Solicit to the DHCPv6 server. Note that the DBNG-CP may learn of the DBNG-UP local Ipv6 address during UP configuration via Mi.
6. The external DHCPv6 server responds with a DHCPv6 Advertise packet that is directed to the DBNG DHCPv6 relay.
7. The DBNG-CP decapsulates the RELAY-REPLY header from the DHCPv6 Advertise packet before sending it to the RG through the DBNG-UP, utilizing the session control packet redirect tunnel.
8. The DHCPv6 Request is sent from the RG to the DBNG-CP through the DBNG-UP utilizing the session control packet redirect tunnel.
9. The DBNG-CP encapsulates the DHCPv6 Request with a RELAY-FORW header, setting the link address to the DBNG-UP local Ipv6 address of the subscriber logical port (V interface), and converts the multicast DHCPv6 Request to unicast. The DBNG-CP sends the DHCPv6 Request to the DHCPv6 server.
10. The external DHCPv6 server responds with a DHCPv6 Reply packet that is directed to the DBNG DHCPv6 relay.
11. The DBNG-CP sends a subscriber Session Modification Request to update the traffic forwarding rules to match on the subscriber's Ipv6 address and/or prefixes. Note that this step may occur as early as after step 8, but the implications of this approach should be considered.
12. The DBNG-UP sends a Session Modification Response to the DBNG-CP, informing that the forwarding rules have been updated, and the DBNG-UP is ready to forward the subscriber's Ipv6 data packets.
13. The DBNG-CP decapsulates the RELAY-REPLY header from the DHCPv6 Reply packet before sending it to the RG through the DBNG-UP, utilizing the session control packet redirect tunnel.
14. DBNG-CP informs the RG of the default gateway Link Local Address (LLA).

4.5.9.3 IpoE DHCPv6 Relay Delayed Session Creation (via DBNG-UP)

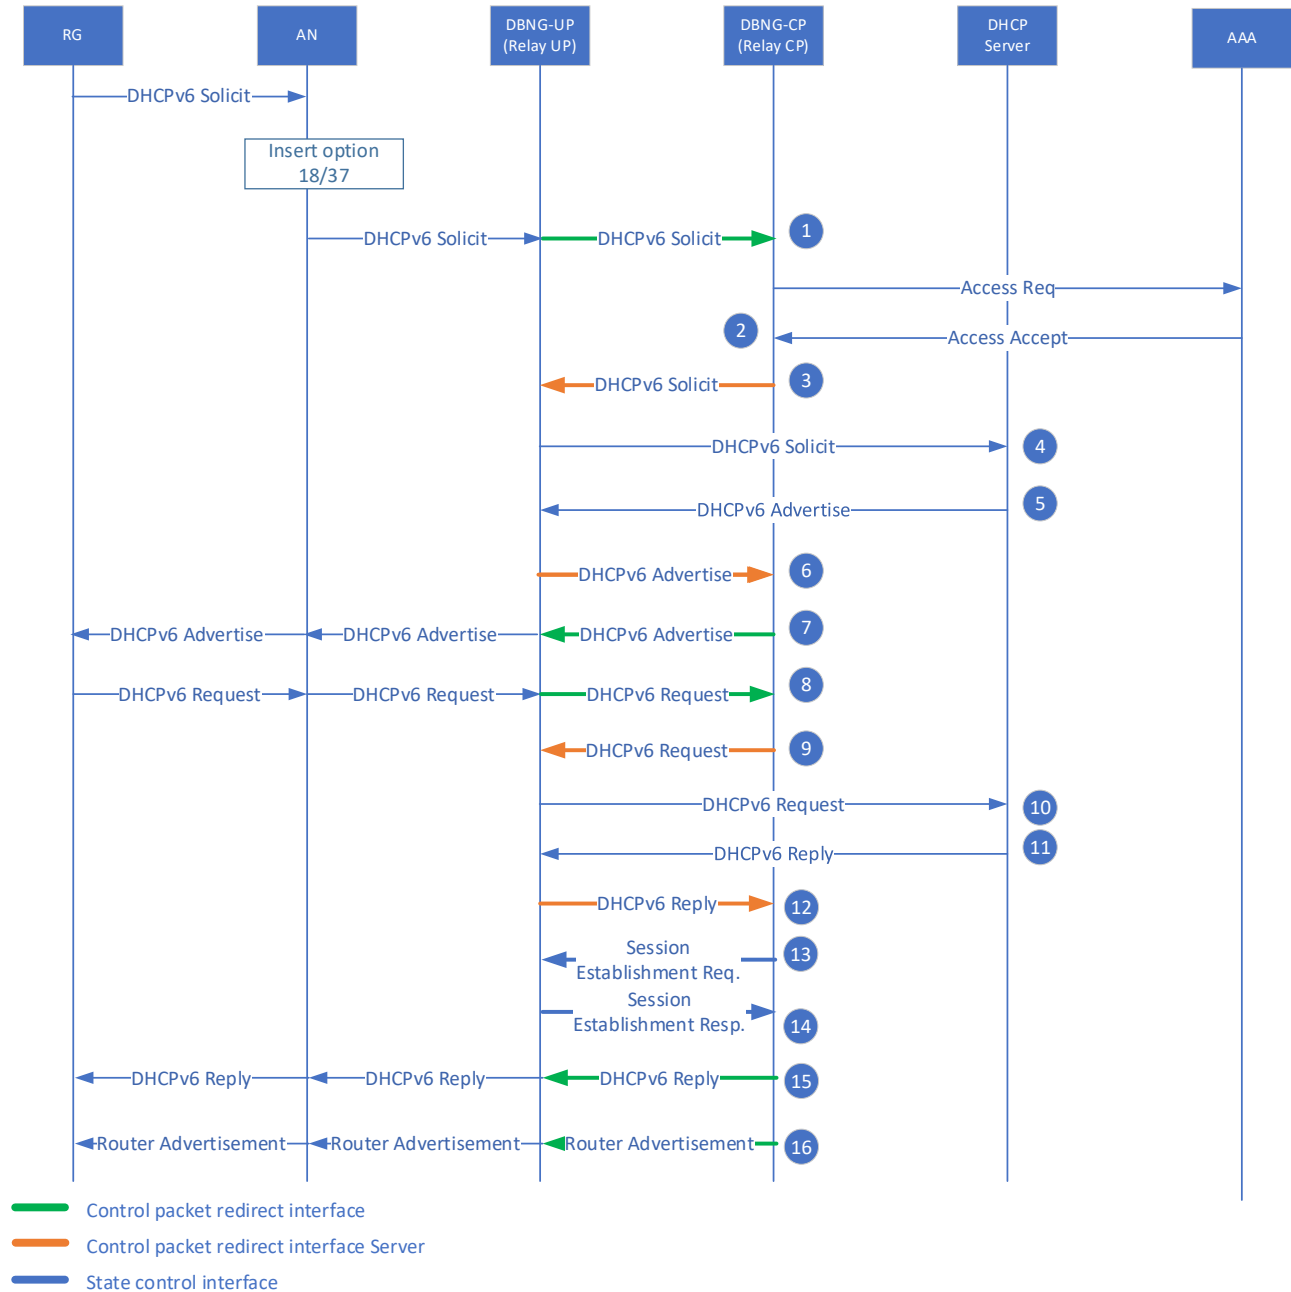


Figure 30: IpoE DHCPv6 Relay Delayed Session Creation (via DBNG-UP) call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions in Figure 17. Also, the call flow in section 4.5.3 and Figure 19 covers the server CPR rule for exchanging DHCP control packets between the DBNG-CP and external DHCPv6 server through the DBNG-UP.

1. The DBNG-CP triggers an Access-Request to authenticate the RG
2. AAA successfully authenticates the RG and replies with Access-Accept
3. The DBNG-CP encapsulates the DHCPv6 Solicit with a RELAY-FORW header, setting the link address to the DBNG-UP local Ipv6 address of the logical port (V interface), and converts the multicast DHCPv6 Solicit to unicast. The DBNG-CP sends the DHCPv6 Solicit to the DBNG-UP utilizing the server control packet redirect tunnel. Note that the DBNG-CP may learn of the DBNG-UP local Ipv6 address during UP configuration via Mi.
4. The DBNG-UP forwards the DHCPv6 Solicit to the external DHCPv6 server.
5. The external DHCPv6 server responds with a DHCPv6 Advertise packet that is directed to the DBNG DHCPv6 relay.
6. The DBNG-UP detects the downstream DHCPv6 Advertise packet, containing the assigned Ipv6 address(es), from the external server and redirects the packet to the DBNG-CP through the server Control Packet Redirect tunnel.
7. The DBNG-CP decapsulates the RELAY-REPLY header from the DHCPv6 Advertise packet before sending it to the RG through the DBNG-UP, utilizing the default/common redirect interface tunnel.
8. The DHCPv6 Request is sent from the RG to the DBNG-CP through the DBNG-UP utilizing the default/common control packet redirect tunnel.
9. The DBNG-CP encapsulates the DHCPv6 Request with a RELAY-FORW header, setting the link address to the DBNG-UP local Ipv6 address of the subscriber logical port (V interface), and converts the multicast DHCPv6 Request to unicast. The DBNG-CP sends the DHCPv6 Request to the DBNG-UP utilizing the server control packet redirect tunnel.
10. The DBNG-UP forwards the DHCPv6 Request to the external DHCPv6 server.
11. The external DHCPv6 server responds with a DHCPv6 Reply packet that is directed to the DBNG DHCPv6 relay.
12. The DBNG-UP detects the downstream DHCPv6 Reply packet from the external server and redirects the packet to the DBNG-CP through the server Control Packet Redirect tunnel.
13. At this stage, DBNG-CP sends the subscriber Session Establishment Request, containing session specific control packet redirection rules and data packet forwarding rules for this subscriber.
14. The DBNG-UP sends a Session Establishment Response to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward the subscriber's Ipv6 data packets.
15. The DBNG-CP decapsulates the RELAY-REPLY header from the DHCPv6 Reply packet before sending it to the RG through the DBNG-UP, utilizing the session control packet redirect tunnel.
16. DBNG-CP informs the RG of the default gateway (Link Local Address).

Note: Session Establishment Request and Response can occur as soon as step 9 take place, but there are implications of accepting the RG requested IP address without the server validation.

4.5.9.4 IpoE DHCPv6 Relay Delayed Session Creation (via DBNG-CP)

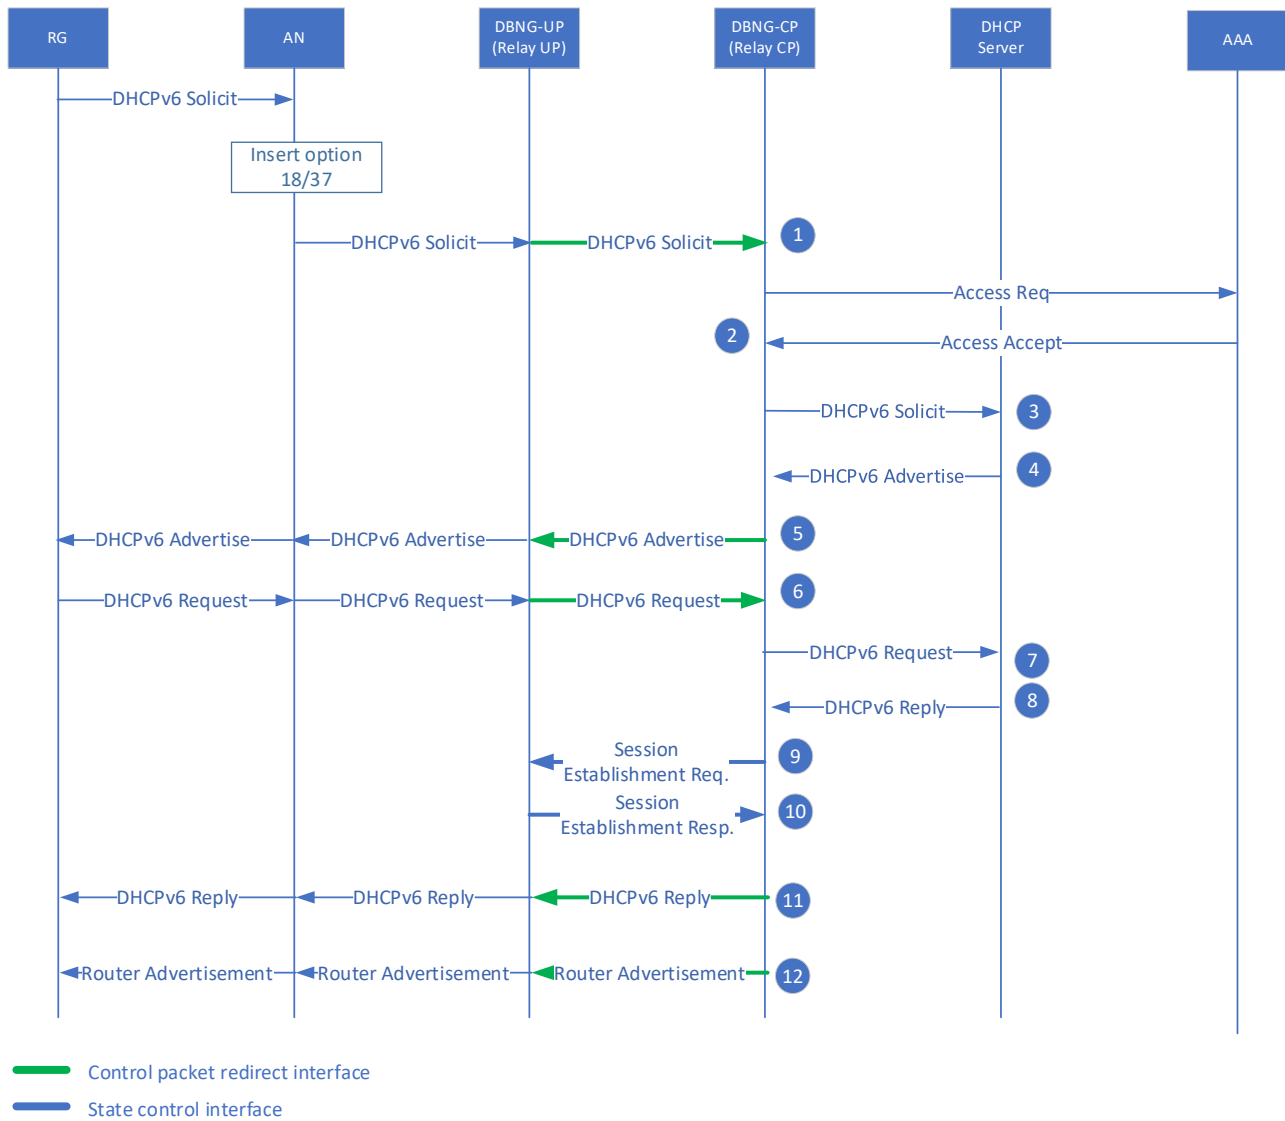


Figure 31: IpoE DHCPv6 Relay Delayed Session Creation (via DBNG-CP) call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions in Figure 17.

- 15. The DBNG-CP triggers an Access-Request to authenticate the RG
- 16. AAA successfully authenticates the RG and replies with Access-Accept
- 17. The DBNG-CP encapsulates the DHCPv6 Solicit with a RELAY-FORW header, setting the link address to the DBNG-UP local Ipv6 address of the logical port (V interface), and converts the multicast DHCPv6 Solicit to unicast. The DBNG-CP sends the DHCPv6 Solicit to the DHCP server via a local interface. Note that the DBNG-CP may learn of the DBNG-UP local Ipv6 address during UP configuration via Mi.

18. The external DHCPv6 server responds with a DHCPv6 Advertise packet that is directed to the DBNG DHCPv6 relay.
19. The DBNG-CP decapsulates the RELAY-REPLY header from the DHCPv6 Advertise packet before sending it to the RG through the DBNG-UP, utilizing the default/common redirect interface tunnel.
20. The DHCPv6 Request is sent from the RG to the DBNG-CP through the DBNG-UP utilizing the default/common control packet redirect tunnel.
21. The DBNG-CP encapsulates the DHCPv6 Request with a RELAY-FORW header, setting the link address to the DBNG-UP local Ipv6 address of the subscriber logical port (V interface), and converts the multicast DHCPv6 Request to unicast. The DBNG-CP sends the DHCPv6 Request to the DHCPv6 server via a local control/management interface.
22. The external DHCPv6 server responds with a DHCPv6 Reply packet that is directed to the DBNG DHCPv6 relay.
23. At this stage, DBNG-CP sends the subscriber Session Establishment Request, containing session specific control packet redirection rules and data packet forwarding rules for this subscriber.
24. The DBNG-UP sends a Session Establishment Response to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward the subscriber's Ipv6 data packets.
25. The DBNG-CP decapsulates the RELAY-REPLY header from the DHCPv6 Reply packet before sending it to the RG through the DBNG-UP, utilizing the session control packet redirect tunnel.
26. DBNG-CP informs the RG of the default gateway (Link Local Address).

Note: Session Establishment Request and Response can occur as soon as step 7 take place, but there are implications of accepting the RG requested IP address without the server validation.

4.5.10 IpoE SLAAC

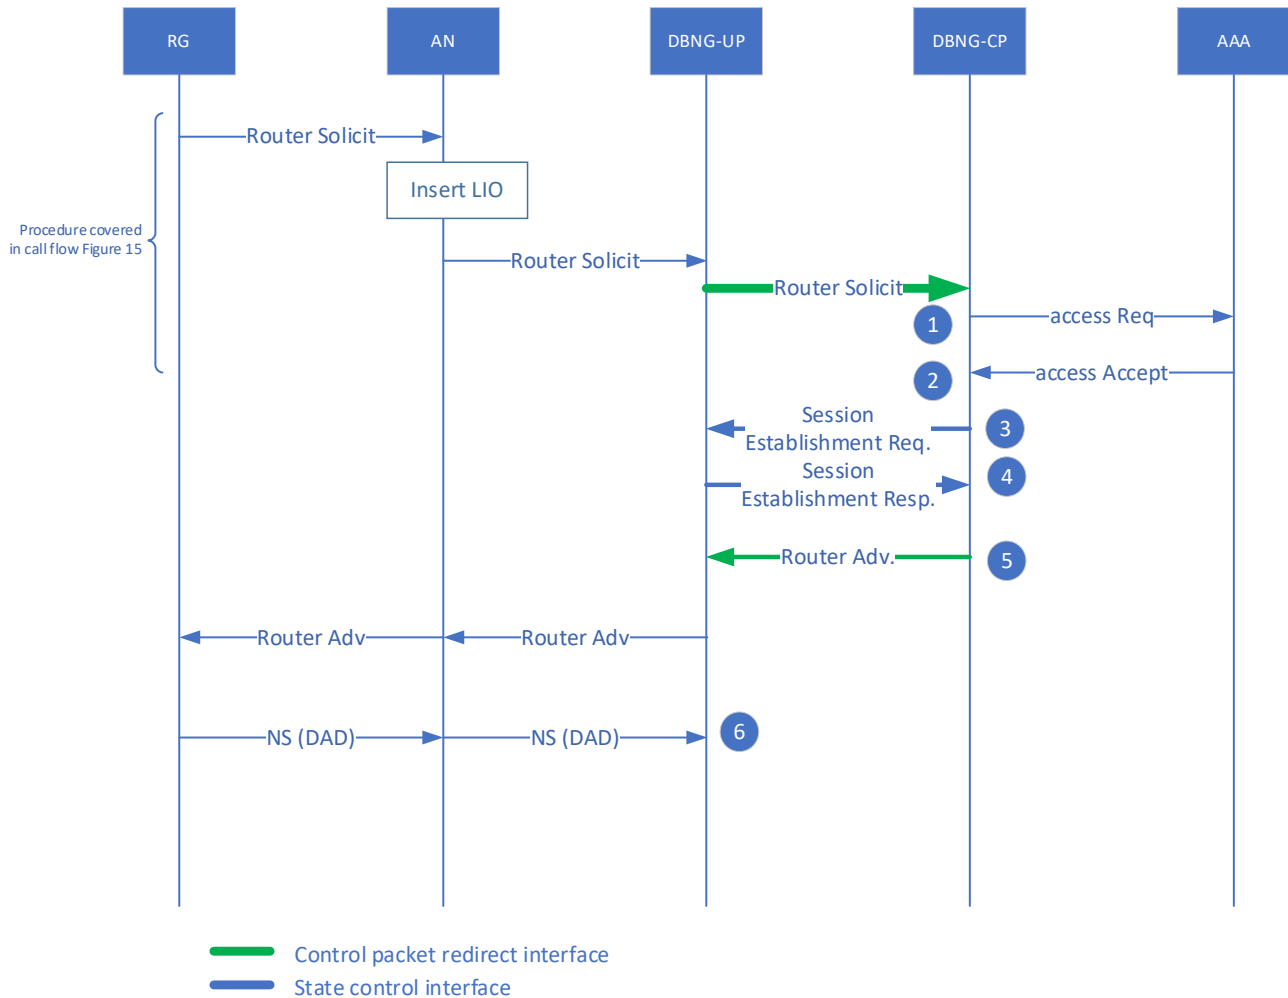


Figure 32: IpoE SLAAC call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common control packet redirection rule.

1. The DBNG-CP triggers an Access-Request to authenticate the RG.
2. The AAA successfully authenticates the RG and replies to the DBNG-CP with an Access-Accept.
3. The DBNG-CP assigns the IP prefix to the RG which is obtained through either the local address server or returned by AAA as one of the attributes. The DBNG-CP sends a Session Establishment Request to create new packet forwarding states for the data packet. This updates the data plane state.
4. The DBNG-UP sends a Session Establishment Response back to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward the subscriber's IP data packets.
5. The SLAAC process completes by sending the Router Advertisement back to the RG. The DBNG-CP forwards the RA through the dedicated session control packet redirect tunnel back to the RG though the DBNG-UP.
6. RG sends Neighbor Solicit for Duplicate Address Detection (DAD)

4.5.11 IpoE Data Trigger

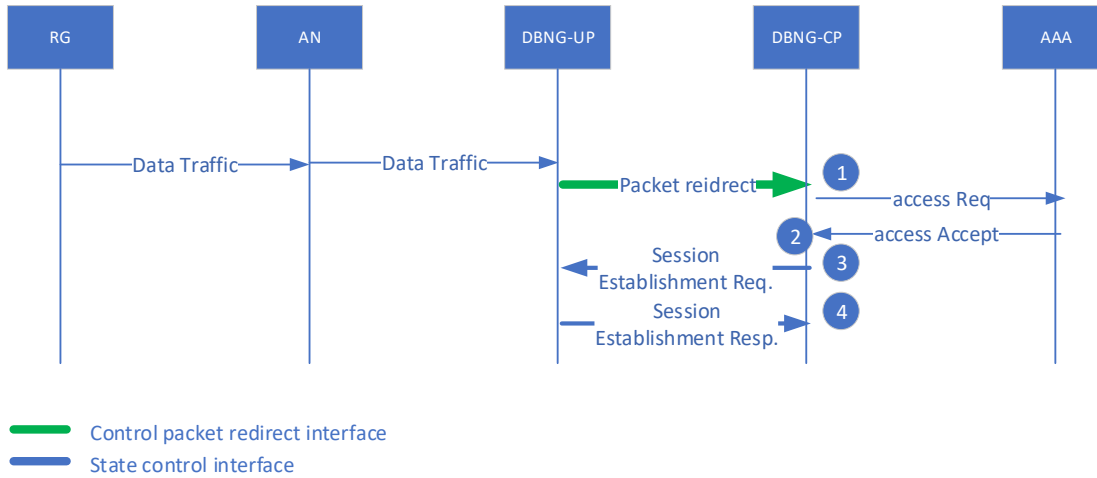


Figure 33: IpoE Data Trigger call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common control packet redirection rule.

1. A data packet is received from the RG is redirected to the DBNG-CP from the DBNG-UP for authentication.
2. The AAA successfully authenticates the RG based on parameters including IP and MAC, then replies to the DBNG-CP with an Access-Accept.
3. At this point the DBNG-CP can send a session establishment request to create new packet forwarding states for the data packet. This updates the data plane forwarding state.
4. The DBNG-UP sends a response back to the DBNG-CP, informing that the forwarding states are installed, and the DBNG-UP is ready to forward the subscriber's IP data packets.

4.5.12 IpoE Dual Stack immediate session creation

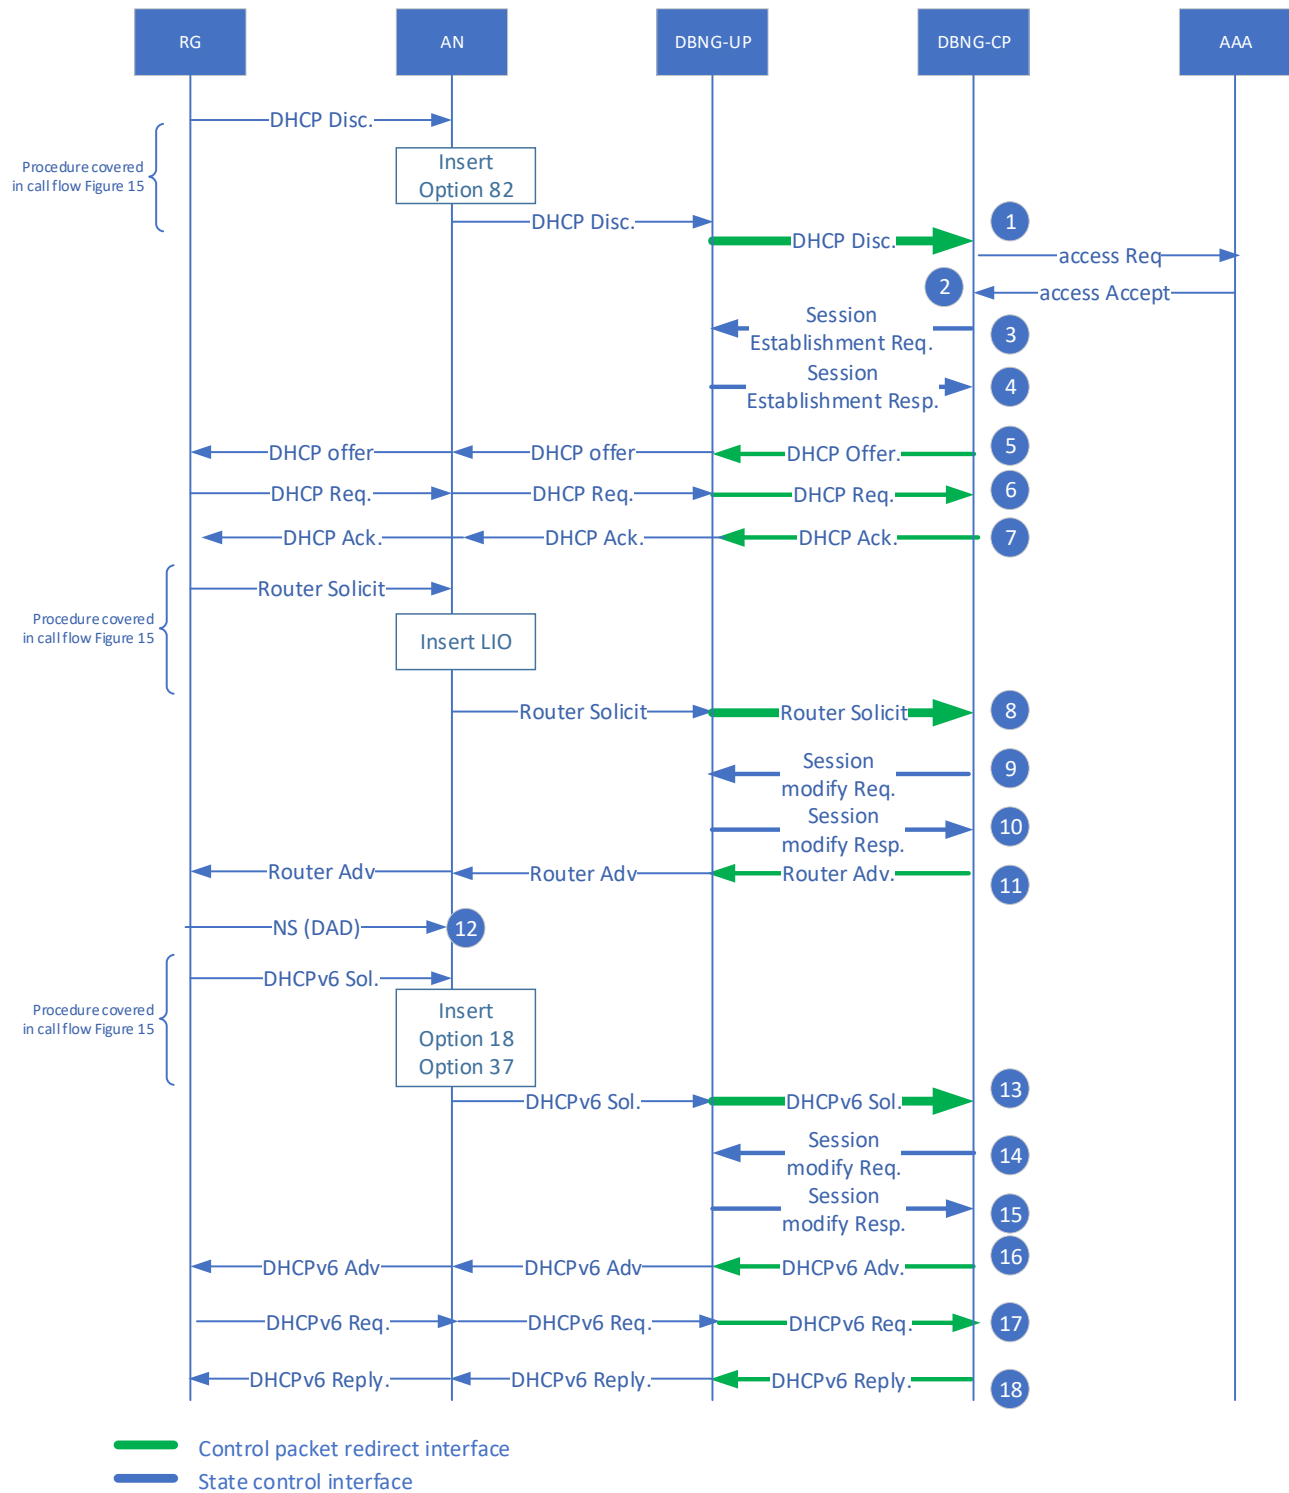


Figure 34: IpoE Dual Stack immediate session creation call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common control packet redirection rule.

1-7. Follows the same procedure as section 4.5.4 step 1-7.

Ipv6 SLAAC configuration:

8. RG initiates a Router Solicit which is redirected through the CPR interface from the DBNG-UP to the DBNG-CP.

9-12: Follows the same procedure as section 4.5.10 step 3-6. However, instead of a Session establishment request and response between the DBNG-CP and DBNG-UP, it is a session modification request and respond.

DHCPv6 negotiation for user's address (PD) configuration;

13. RG initiates a DHCPv6 solicit which is redirected through the CPR interface from the DBNG-UP to the DBNG-CP.

14-18: Follows the same procedure as section 4.5.7 step 3-7. However, instead of a Session establishment request and response between the DBNG-CP and DBNG-UP, it is a session modification request and response.

Note: The dual-stack access processes could be concurrent, one before another, or one after another.

4.5.13 IpoE Dual Stack delayed session creation

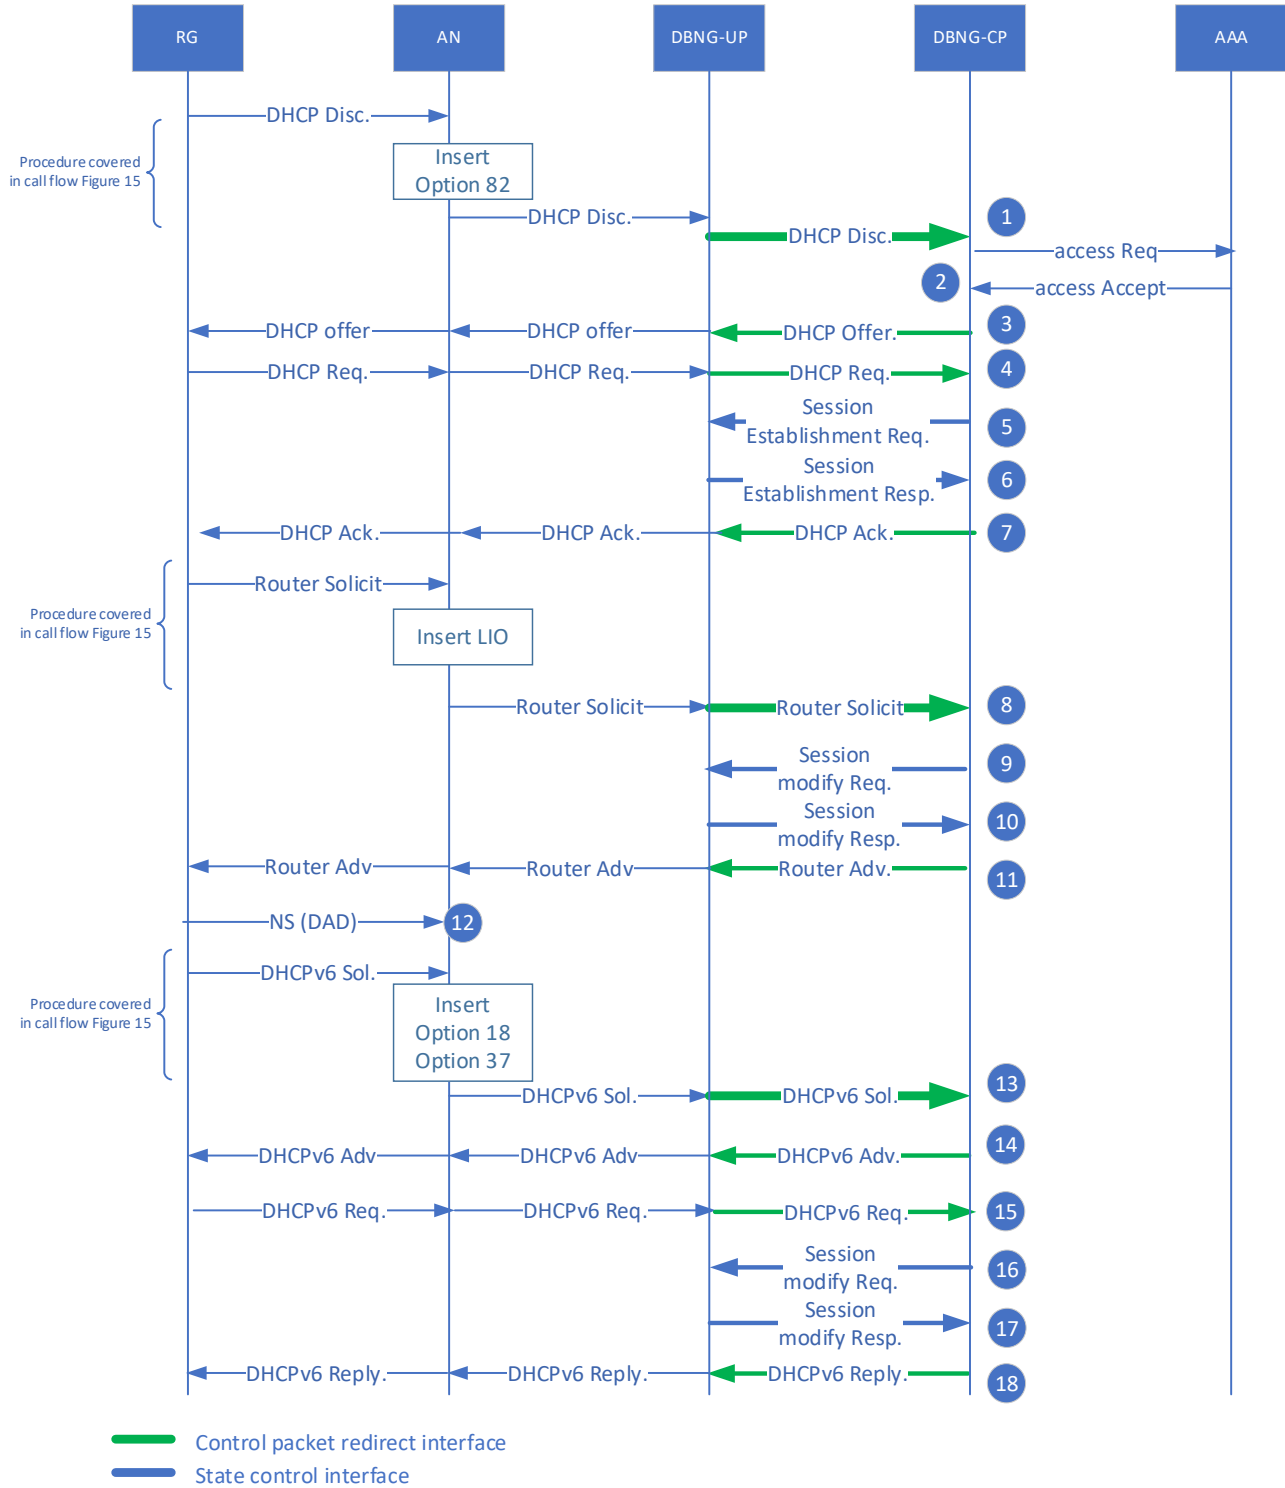


Figure 35: IpoE Dual Stack delayed session creation call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions.

1-7. Follows the same procedure as section 4.5.4 step 1-7 for delayed session creation model.

Ipv6 SLAAC configuration:

8. RG initiates a Router Solicit which is redirected through the upstream default redirect tunnel interface from the DBNG-UP to the DBNG-CP.

9-12: Follows the same procedure as section 4.5.10 ,step 3-6. However, instead of a Session Create request and response between the DBNG-CP and DBNG-UP, it is a session modification request and respond.

DHCPv6 negotiation for user's address (PD) configuration;

13. RG initiates a DHCPv6 solicit which is redirected through the upstream default redirect tunnel interface from the DBNG-UP to the DBNG-CP.

14-18: Follows the same procedure as section 4.5.7 delayed session creation model step 3-7. However, instead of a Session Create request and response between the DBNG-CP and DBNG-UP, it is a session modification request and response.

Note: The dual-stack access processes could be concurrent, one before another, or one after another.

4.5.14 IpoE SLAAC and DHCPv6 PD immediate session creation

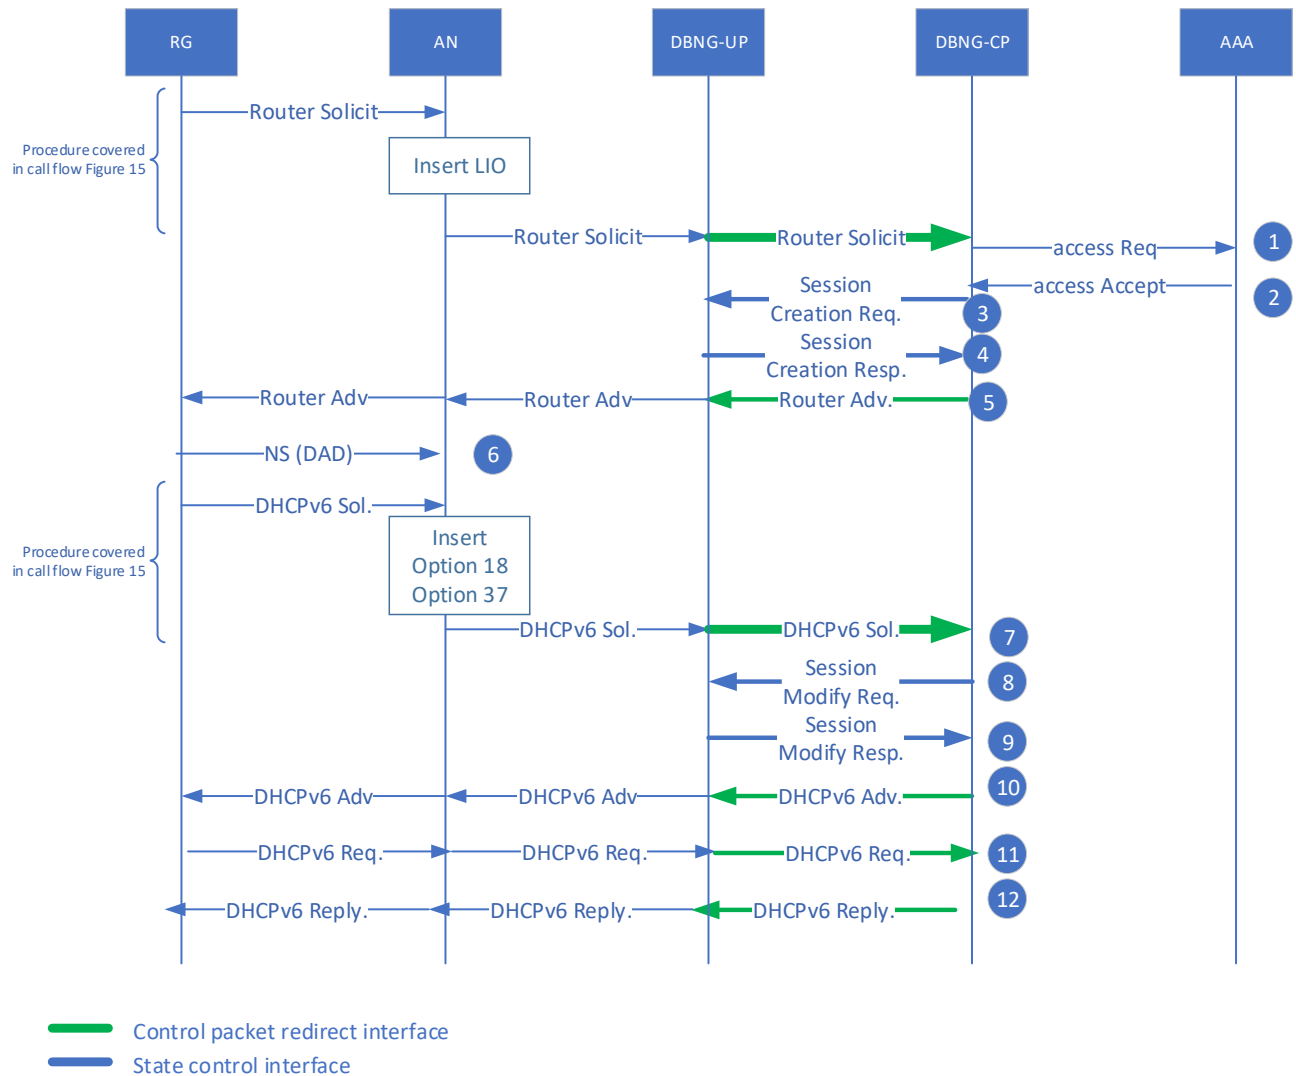


Figure 36: IpoE SLAAC and DHCPv6 PD immediate session creation call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common control packet redirection rule.

1-6: Follows the same procedure as section 4.5.10 step 1-6.

7. DHCPv6 solicit is initiated from RG, and is sent through AN, DBNG-UP redirects it to DBNG-CP where it responds with DHCPv6 Advertise from the DBNG-CP or DHCPv6 server as a neighboring system;

8-12 Follows the same procedure as section 4.5.7 step 3-7.

4.5.15 IpoE SLAAC and DHCPv6 PD delayed session creation

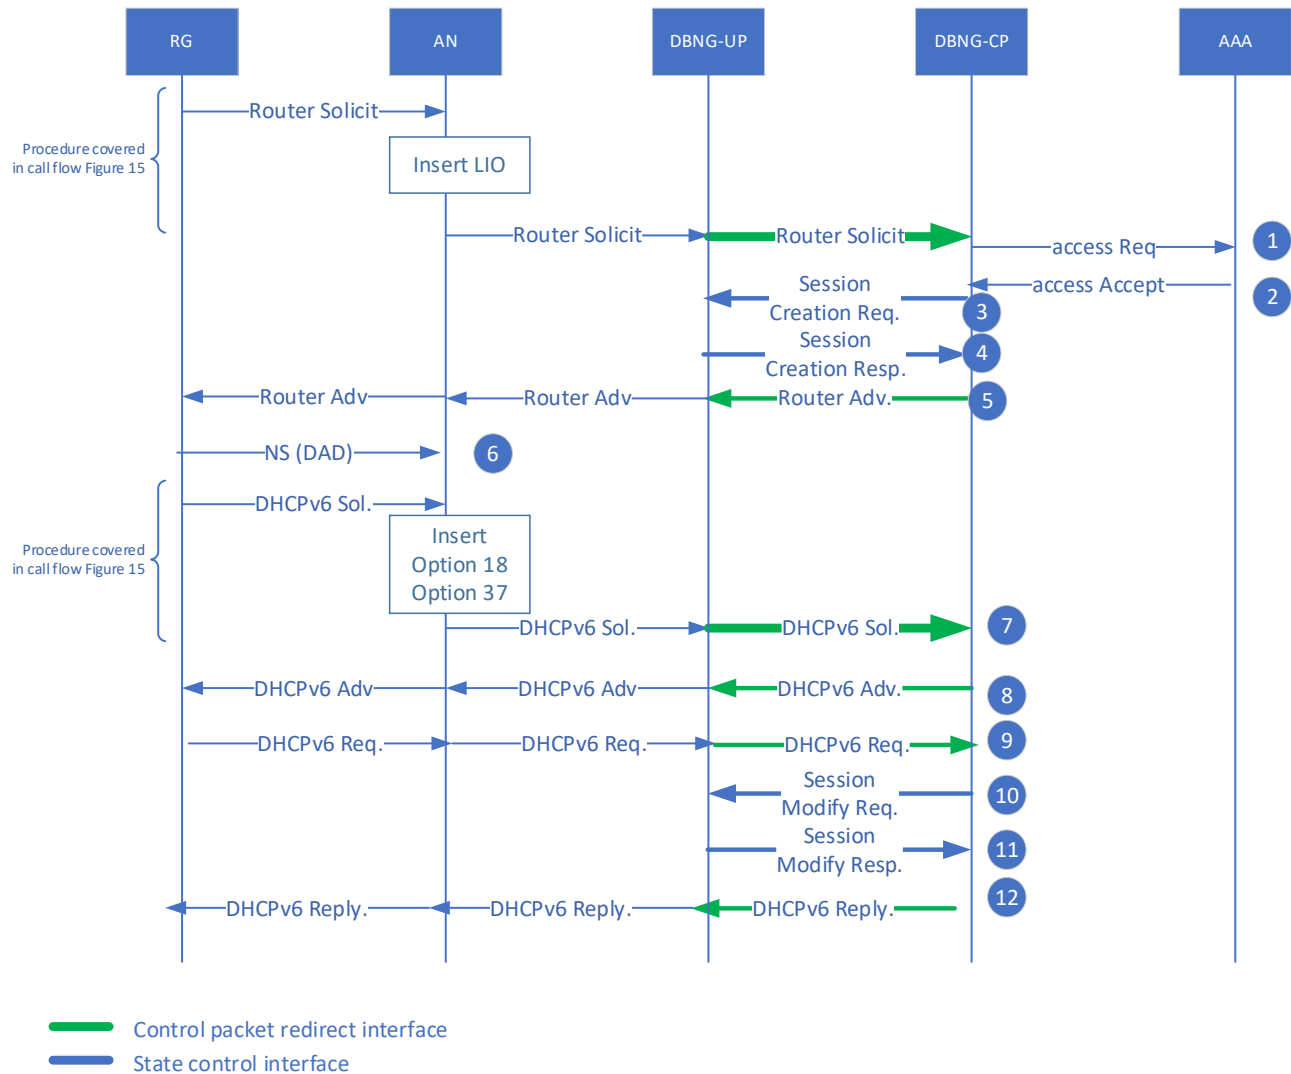


Figure 37: IpoE SLAAC and DHCPv6 PD delayed session creation call flow

P Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions.

1-6: Follows the same procedure as section 4.5.10 step 1-6.

7. DHCPv6 solicit is initiated from RG, and is sent through AN, DBNG-UP redirects through the upstream default redirect tunnel interface from the DBNG-UP to the DBNG-CP where it responds with DHCPv6 Advertise utilizing the downstream default redirect tunnel interface from the DBNG-CP or DHCPv6 server as a neighboring system;

8-12 Follows the same procedure as section 4.5.7 step 3-7.

4.5.16 PPPoE immediate session creation

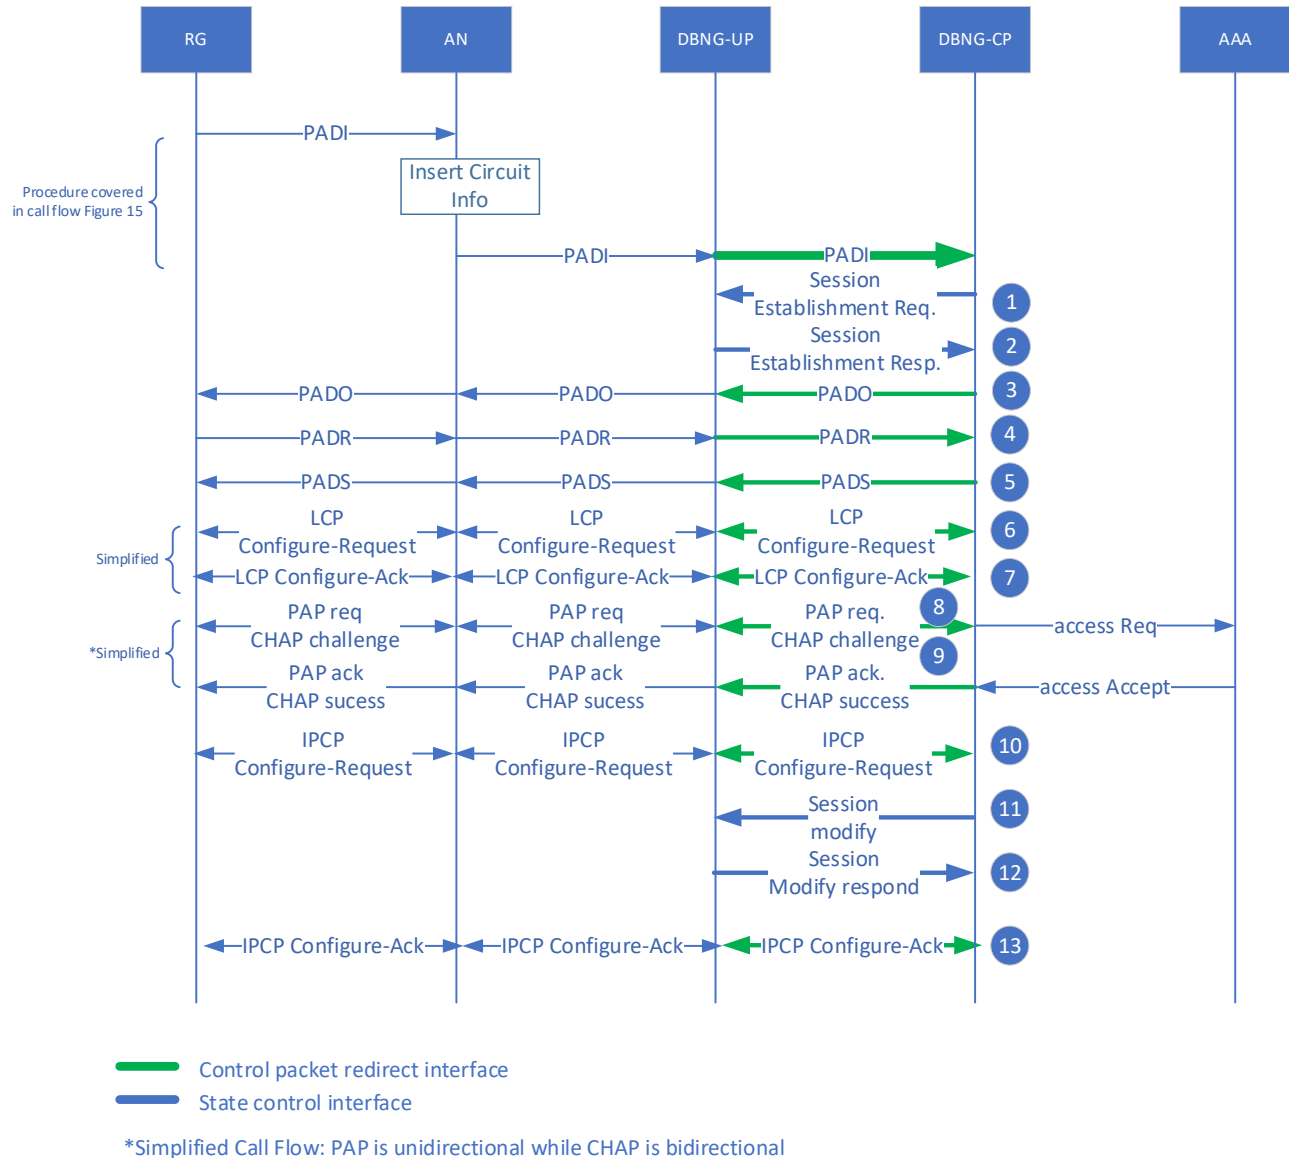


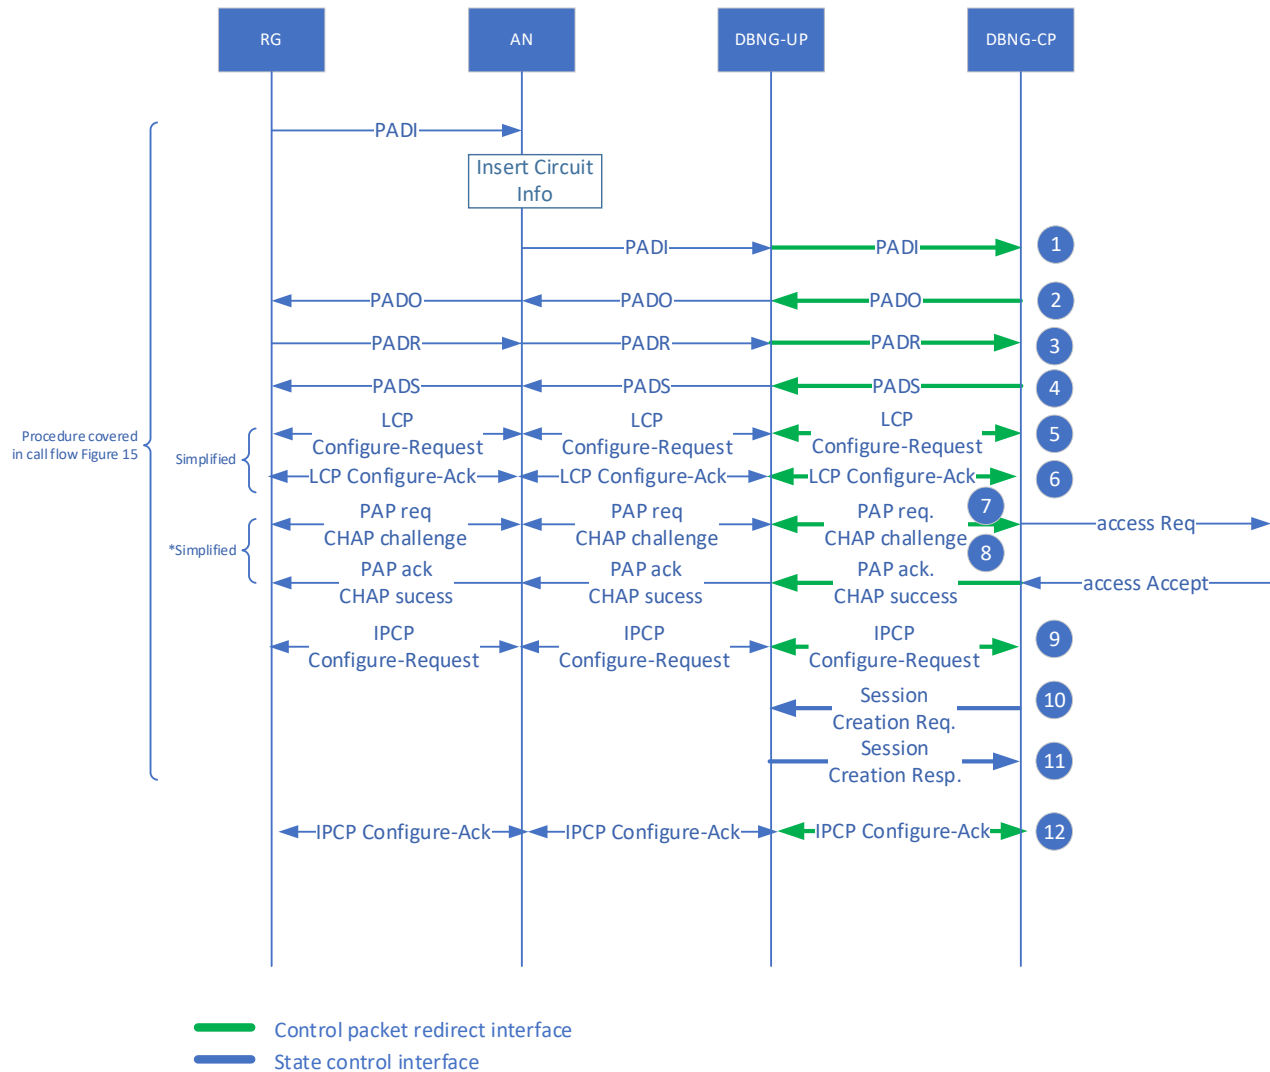
Figure 38: PPPoE immediate session creation call flow

Section 4.5.2 covers the generic common CPRi creation, which is carried out prior to step 1.

1. Upon receiving the first control packet through the common CPRi, the DBNG-CP sends a session establishment request to DBNG-UP to create forwarding rules for the control packets of the subscriber session. This creates a CPRi dedicated to the subscriber session.
2. The DBNG-UP sends a response back to the DBNG-CP, informing that the rules are installed, and the DBNG-UP is ready to forward the subscriber's PPPoE and PPP control packets.
3. The DBNG-CP sends the PADO back to the RG through the DBNG-UP utilizing the per session CPR interface.
4. The PADR is sent from the RG through the DBNG-UP utilizing the CPR interface.
5. The DBNG-CP sends the PADS back to the RG through the DBNG-UP utilizing the per session CPR interface.
6. Both DBNG-CP and RG send LCP requests to the peer utilizing the per session CPRi. In this step the DBNG-CP communicates the accepted authentication method (PAP or CHAP) to the RG.

7. Both DBNG-CP and RG send LCP Config-acks to the peer utilizing the per session CPRI.
8. Depending on the authentication method previously negotiated between RG and DBNG-CP, the RG might send a PAP Config-Req to the DBNG-CP or the DBNG-CP might send a CHAP Authentication Challenge to the RG.:
 - Option 1: If PAP was negotiated in the previous steps, the RG sends a PAP request to the DBNG-CP through the DBNG-UP utilizing the per session CPR Interface. The credentials are sent in the Access-Request to the AAA server.
 - Option 2: If CHAP was negotiated in the previous steps, the DBNG-CP initiates a challenge to the RG though the DBNG-UP utilizing the per session CPR Interface. The RG responds back to the challenge to the DBNG-CP. The challenge response is sent to the AAA server.
9. The AAA successfully authenticates the RG and replies to the DBNG-CP with an Access Accept, which in turn results in a PAP/CHAP success message sent to the RG by the DBNG-CP through the per session CPR interface.
10. Both DBNG-CP and RG send IPCP Configure-Request for parameter negotiation, utilizing the per session CPR interface. The RG is assigned an Ipv4 address. This address could have been assigned either by the AAA or chosen by the DBNG-CP as local PPP server.
11. Once the RG is assigned an IP address, the DBNG-CP sends a session modification request to the DBNG-UP in order to update the packet forwarding rules, adding those for data packets.
12. The DBNG-UP sends a response back to the DBNG-CP, informing that the rules are installed, and the DBNG-UP is ready to forward the subscribers IP data packets.
13. The DBNG-CP and RG exchange the IPCP Configure-Acks through the DBNG-UP utilizing the per session CPR interface.

4.5.17 PPPoE delayed session creation



*Simplified Call Flow: PAP is unidirectional while CHAP is bidirectional

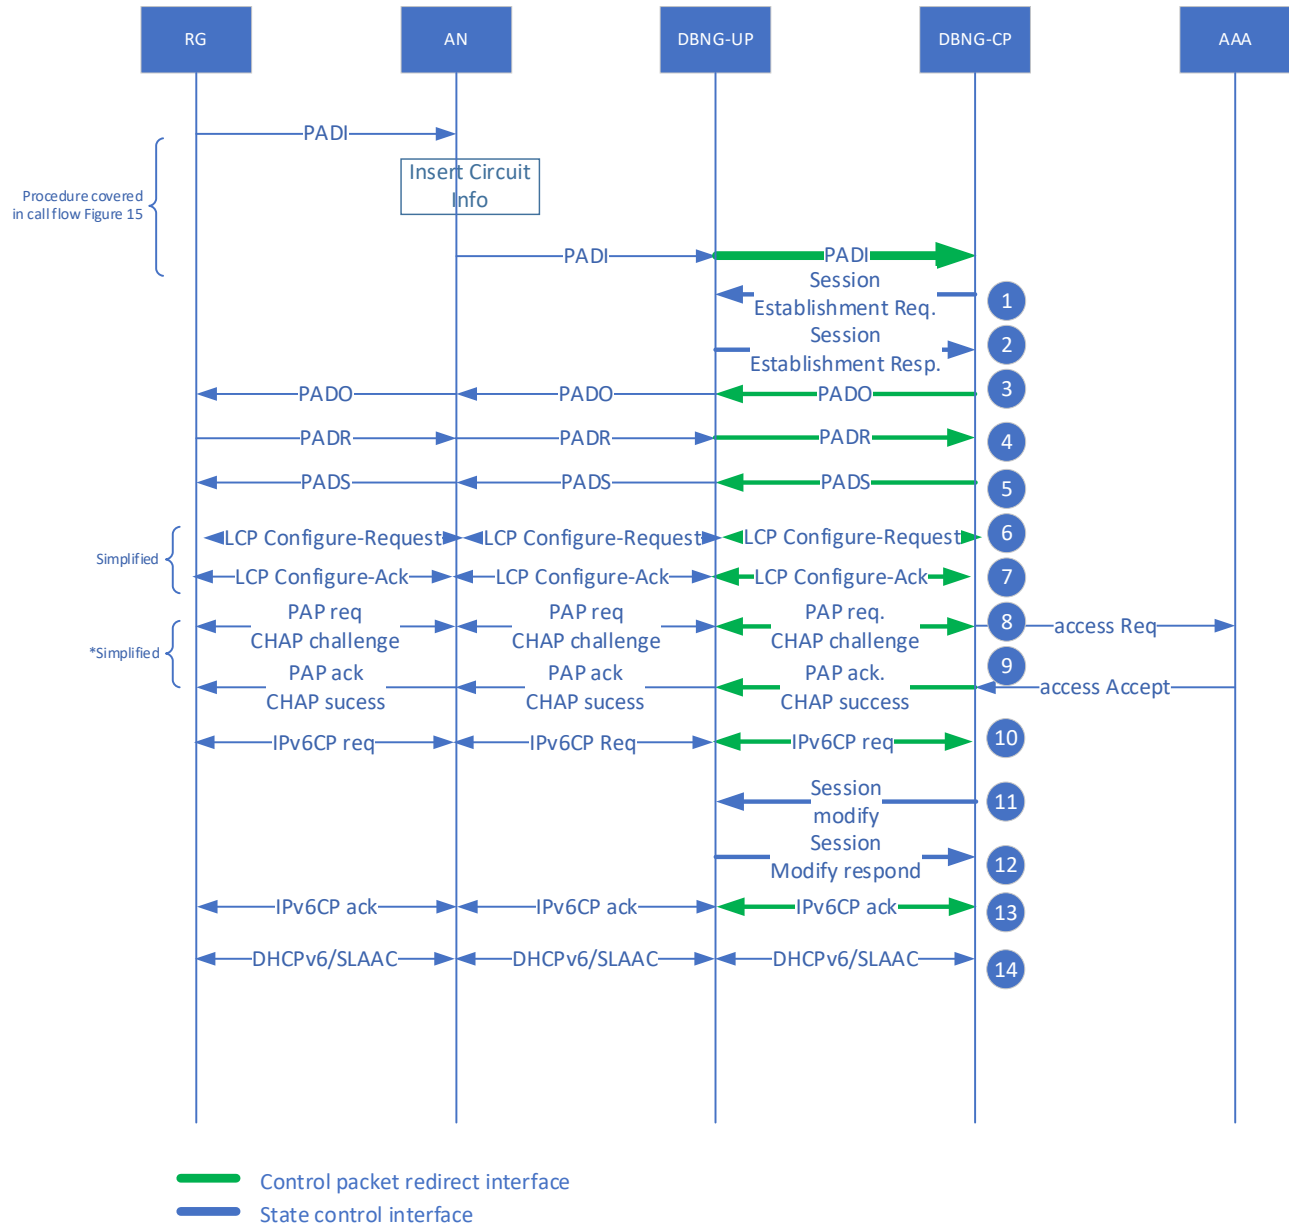
Figure 39: PPPoE delayed session creation call flow

Section 4.5.2 covers the generic CPRi creation, which is carried out prior to step 1. Delayed Session Creation requires that the DBNG-CP has programmed bidirectionally a common CPR interface.

1. DBNG-CP receives the PADI enhanced by the Line Id inserted by the AN through the common CPR interface.
2. The DBNG-CP sends the PADO back to the RG through the DBNG-UP utilizing the common CPR interface.
3. The PADR is sent from the RG through the DBNG-UP utilizing the common CPR interface.
4. The DBNG-CP sends the PADS back to the RG through the DBNG-UP utilizing the common CPR interface.
5. The LCP Configure-Request is sent from the RG through the DBNG-UP utilizing the common CPR interface.

6. Both DBNG-CP and RG send LCP Config-acks to the peer utilizing the common CPR interface.
7. Depending on the authentication method previously negotiated between RG and DBNG-CP, the RG might send a PAP Config-Req to the DBNG-CP or the DBNG-CP might send a CHAP Authentication Challenge to the RG.:
 - Option 1: If PAP was negotiated in the previous steps, the RG sends a PAP request to the DBNG-CP through the DBNG-UP utilizing the per session CPR Interface. The credentials are sent in the Access-Request to the AAA server.
 - Option 2: If CHAP was negotiated in the previous steps, the DBNG-CP initiates a challenge to the RG through the DBNG-UP utilizing the per session CPR Interface. The RG responds back to the challenge to the DBNG-CP. The challenge response is sent to the AAA server.
8. The AAA successfully authenticates the RG and replies to the DBNG-CP with an Access Accept, which in turn results in a PAP/CHAP success message sent to the RG by the DBNG-CP through the common CPR interface.
9. Both DBNG-CP and RG send IPCP Configure-Request for parameter negotiation, utilizing the common CPRi. The RG is assigned an Ipv4 address. This address could have been assigned either by the AAA or chosen by the DBNG-CP as local PPP server.
10. If a session has not been established yet on the DBNG-UP, the session must be programmed at this step. The DBNG-CP sends a session creation request to create new packet forwarding rules for both the control packets and data packets of the subscriber session.
11. The DBNG-UP sends a response back to the DBNG-CP, informing that the rules are installed, and the DBNG-UP is ready for forwarding both control and data packets.
12. The DBNG-CP and RG exchange the IPCP Configure-Acks through the DBNG-UP utilizing the dedicated CPR interface.

4.5.18 PPPoEv6 immediate session creation



*Simplified Call Flow: PAP is unidirectional while CHAP is bidirectional

Figure 40: PPPoEv6 immediate session creation call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common control packet redirection rule.

1. Upon receiving the first control packet, the DBNG-CP at this point can send a session establishment request to create new packet forwarding states for the data packet. This updates the data plane state.
2. The DBNG-UP sends a response back to the DBNG-CP, informing that the rules are installed, and the DBNG-UP is ready to forward the subscriber’s PPP and PPPoE control packets.
3. After the Session establishment request and response, the DBNG-CP sends the PADO back to the RG through the DBNG-UP utilizing the per session CPR interface.
4. The PADR is sent from the RG through the DBNG-UP utilizing the CPR interface.

5. The DBNG-CP sends the PADS back to the RG through the DBNG-UP utilizing the per session CPR interface.
6. Both DBNG-CP and RG send LCP requests to the peer utilizing the per session CPRi. In this step the DBNG-CP communicates the accepted authentication method (PAP or CHAP) to the RG.
7. Both DBNG-CP and RG send LCP Config-acks to the peer utilizing the per session CPR.
8. Depending on the authentication method previously negotiated between RG and DBNG-CP, the RG might send a PAP Config-Req to the DBNG-CP or the DBNG-CP might send a CHAP Authentication Challenge to the RG:
 - Option 1: If PAP was negotiated in the previous steps, the RG sends a PAP request to the DBNG-CP through the DBNG-UP utilizing the per session CPR Interface. The credentials are sent in an Access-Request to the AAA server.
 - Option 2: If CHAP was negotiated in the previous steps, the DBNG-CP initiates a challenge to the RG though the DBNG-UP utilizing the per session CPR Interface and the RG responds back to the challenge to the DBNG-CP. The challenge response is sent to the AAA server.
9. The AAA successfully authenticates the RG and replies to the DBNG-CP with an Access Accept, which on turn becomes a PAP/CHAP success message sent to the RG by the DBNG-CP through the per session CPR interface.
10. The Ipv6CP Configure-Request is sent from the RG through the DBNG-UP utilizing the per session CPR interface.
11. At this point, the DBNG-CP had obtained the Ipv6 addresses and prefixes for the RG either from the local address server or from AAA returned VSAs. The DBNG-CP sends a session modification request to create new packet forwarding rules for both control and data packet.
 - Packet handling rules for control packets: redirect DHCPv6 and SLAAC request to the DBNG-CP
 - Packet handling rule for data packets: match data packet and perform forwarding action.
12. The DBNG-UP sends a response back to the DBNG-CP, informing that the rules are installed and the DBNG-UP is ready to forward the subscriber's IP data packets.
13. The DBNG-CP and RG exchange Ipv6CP Configure-Ack through the DBNG-UP utilizing the per session CPR interface.
14. In the case of SLAAC prefix assignment, the DBNG-CP sends an RA to the RG informing the Link Local Address which is described in detail in section 4.5.10. In the case of DHCPv6, refer to section 4.5.7 with "Session Establishment Request/Respond" steps to be replaced by "Session Modification Request/Response"

4.5.19 PPPoEv6 delayed session creation

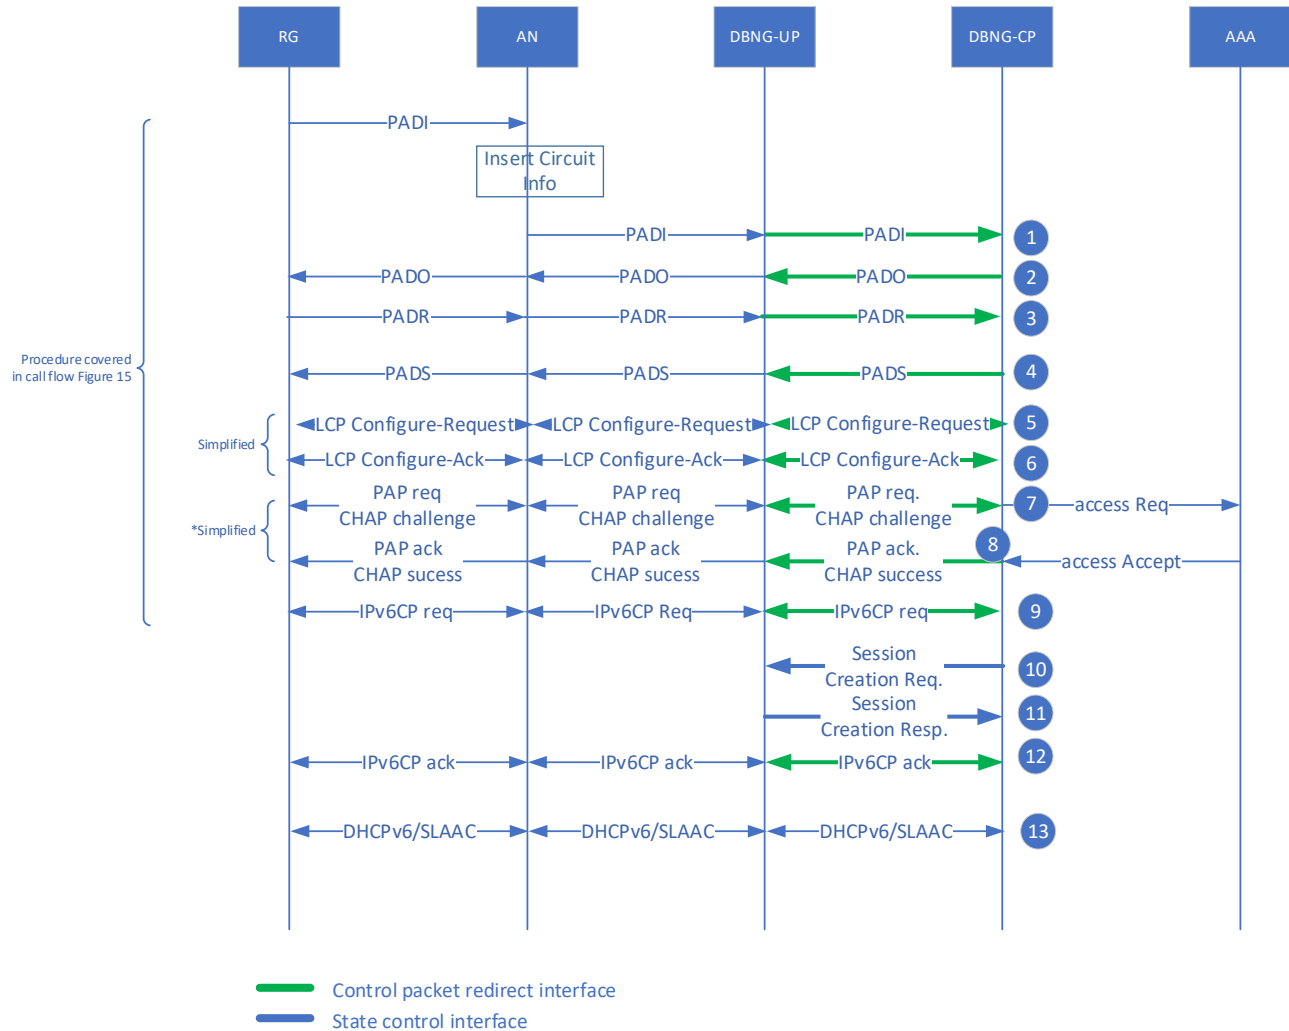


Figure 41: PPPoEv6 delayed session creation call flow

Section 4.5.2 covers the generic CPR interface creation, which is carried out prior to step 1. Delayed Session Creation requires that the DBNG-CP has programmed bidirectionally a common CPR interface.

1. DBNG-CP receives the PADI enhanced by the Line Id inserted by the AN though the common CPR interface.
2. The DBNG-CP sends the PADO back to the RG through the DBNG-UP utilizing the common CPR interface.
3. The PADR is sent from the RG through the DBNG-UP utilizing the common CPR interface.
4. The DBNG-CP sends the PADS back to the RG through the DBNG-UP utilizing the common CPR interface.
5. The LCP Configure-Request is sent from the RG through the DBNG-UP utilizing the common CPR interface.

6. Both DBNG-CP and RG send LCP Config-acks to the peer utilizing the common CPR interface.
7. Depending on the authentication method previously negotiated between RG and DBNG-CP, the RG might send a PAP Config-Req to the DBNG-CP or the DBNG-CP might send a CHAP Authentication Challenge to the RG.:
 - Option 1: If PAP was negotiated in the previous steps, the RG sends a PAP request to the DBNG-CP through the DBNG-UP utilizing the per session CPR Interface. The credentials are sent in the Access-Request to the AAA server.
 - Option 2: If CHAP was negotiated in the previous steps, the DBNG-CP initiates a challenge to the RG through the DBNG-UP utilizing the per session CPR Interface. The RG responds back to the challenge to the DBNG-CP. The challenge response is sent to the AAA server.
8. The AAA successfully authenticates the RG and replies to the DBNG-CP with an Access Accept, which in turn becomes a PAP/CHAP success message sent to the RG by the DBNG-CP through the common CPR interface.
9. The Ipv6CP Configure-Request is sent from the RG through the DBNG-UP utilizing the common CPR interface.
10. At this point, the DBNG-CP has obtained the Ipv6 addresses and prefixes for the RG either from the local address server or from AAA returned VSAs. The DBNG-CP sends a session creation request to create session specific packet forwarding states for both control and data packet. The data plane state is updated.
 - Traffic management rules for control packets: redirect DHCPv6 and SLAAC request to the DBNG-CP
 - Traffic management rule for data packets: match data packet and perform forwarding action.
11. The DBNG-UP sends a response back to the DBNG-CP, informing that the rules are installed and the DBNG-UP is ready to forward the subscriber's IP data packets.
12. The DBNG-CP sends the Ipv6CP Configure-Ack to the RG through the DBNG-UP utilizing the per session CPR interface.
13. In the case of SLAAC prefix assignment, the DBNG-CP sends an RA to the RG informing the Link Local Address which is described in detail in section 4.5.10. In the case of DHCPv6 assignment, please refer to section 4.5.7 delayed session creation mode.

4.5.20 PPPoE Dual Stack immediate session creation

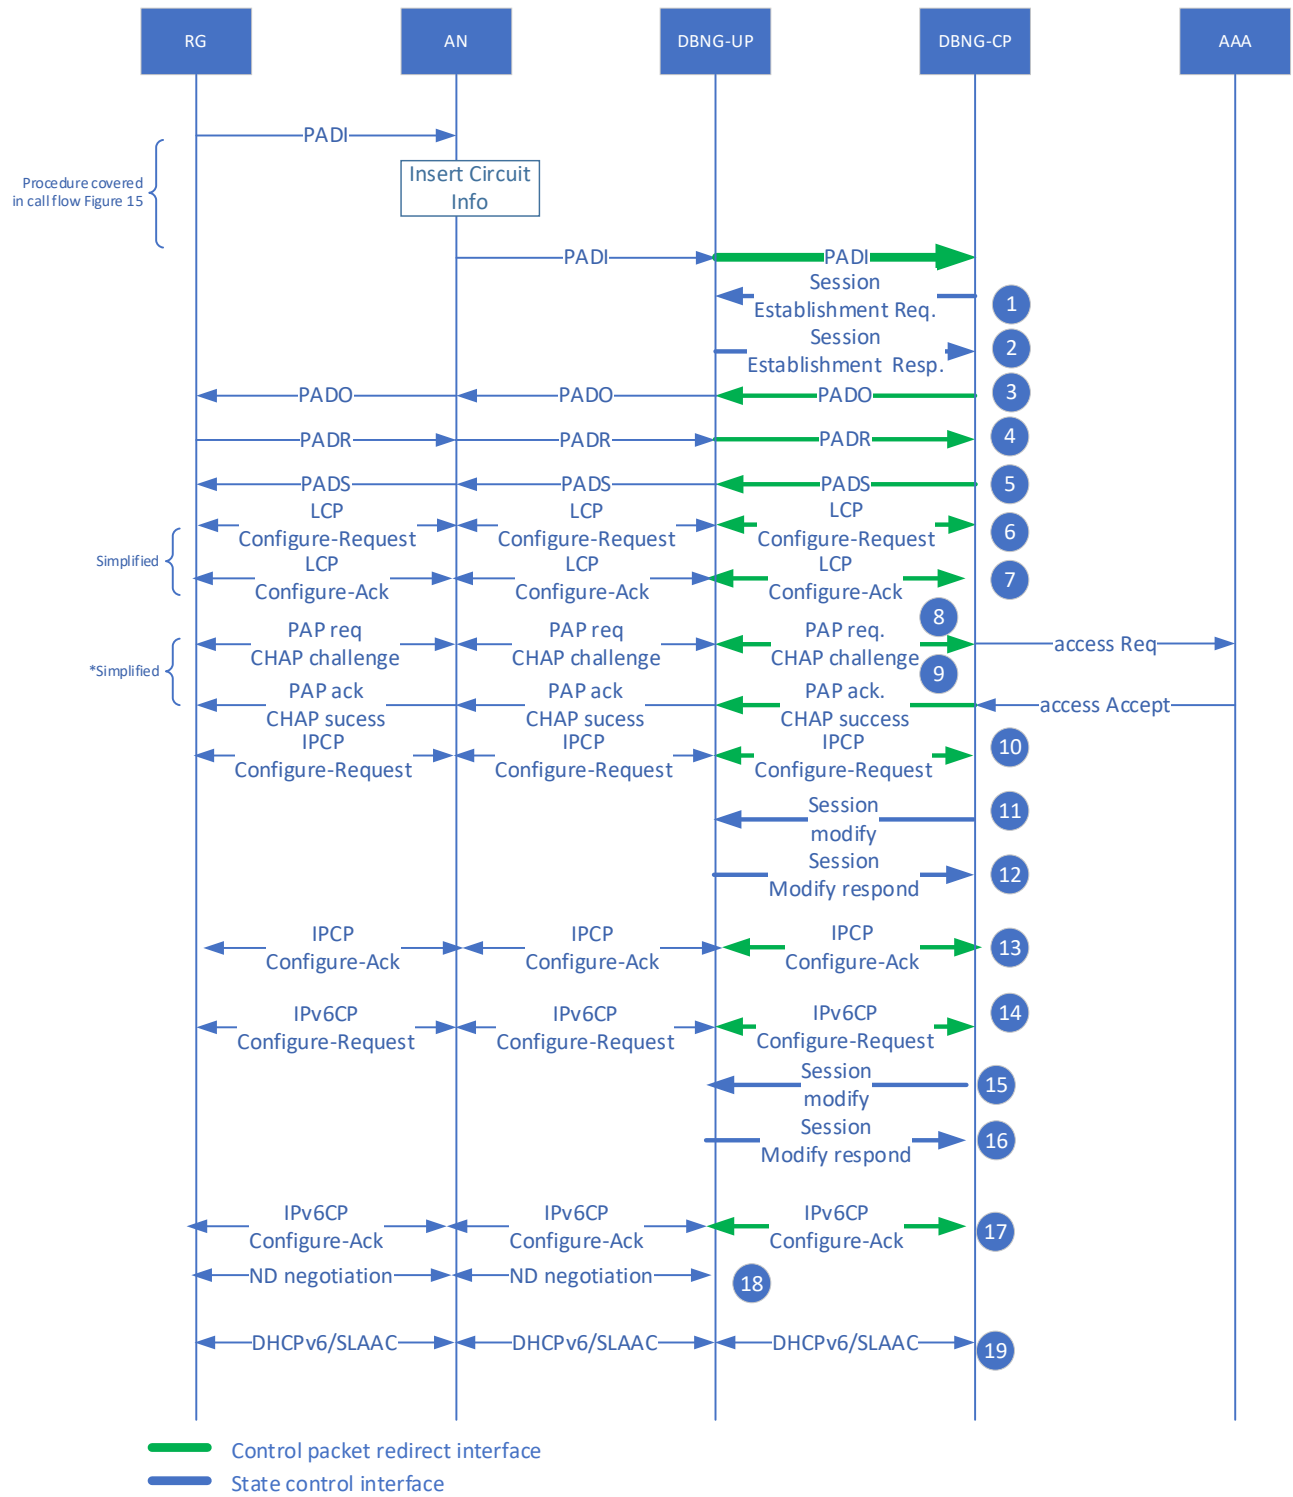


Figure 42: PPPoE Dual Stack immediate session creation call flow

Section 4.5.2 covers the generic CPR interface creation, which is carried out prior to Step 1.

1-13. Follows the same procedure as section 4.5.16 step 1-13

14-19. Follows the same procedure as section 4.5.18 step 10-14

Note: The dual-stack access processes could be concurrent, one before the other, or one after the other.

4.5.21 PPPoE Dual Stack delayed session creation

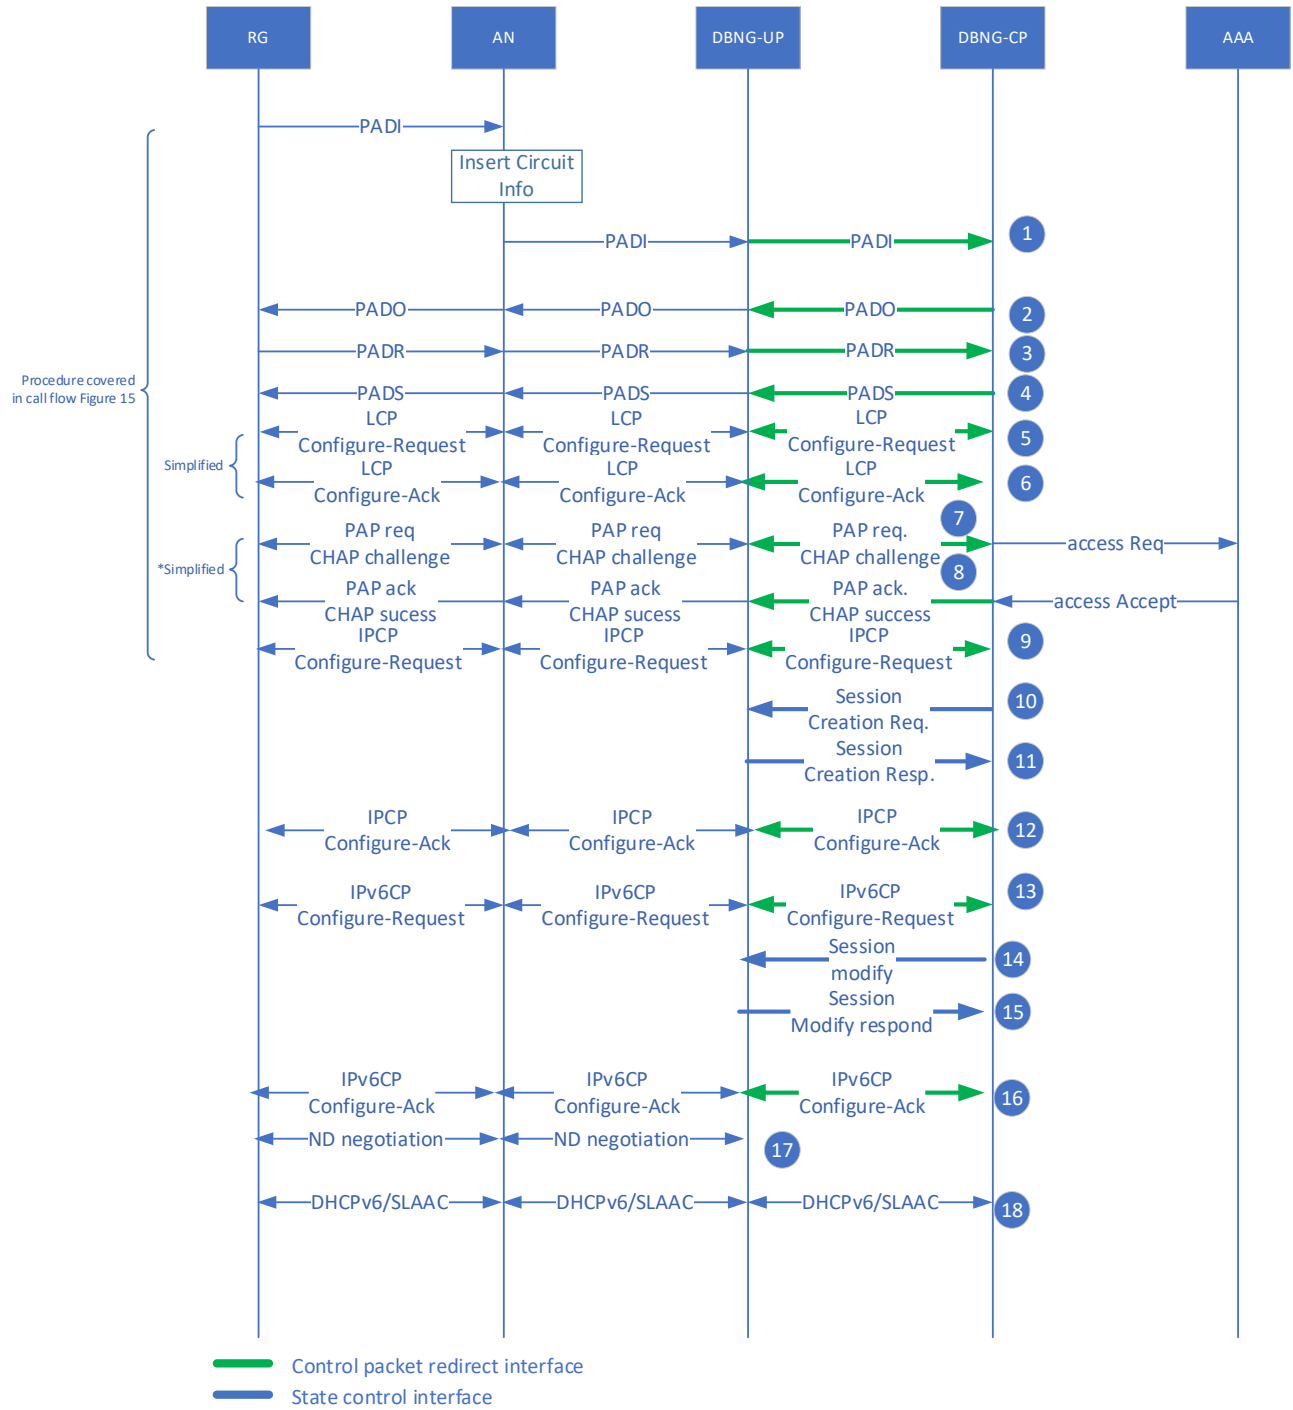


Figure 43: PPPoE Dual Stack delayed session creation call flow

Section 4.5.2 covers the generic CPRi creation, which is carried out prior to step 1. Delayed Session Creation requires that the DBNG-CP has programmed bidirectionally a common CPR interface.

1-13. Follows the same procedure as section 4.5.16 delayed session creation model step 1-13

14-19. Follows the same procedure as section 4.5.18 delayed session creation model step 10-14

Note: The dual-stack access processes could be concurrent, one before the other, or one after the other.

4.5.22 LAC immediate session creation

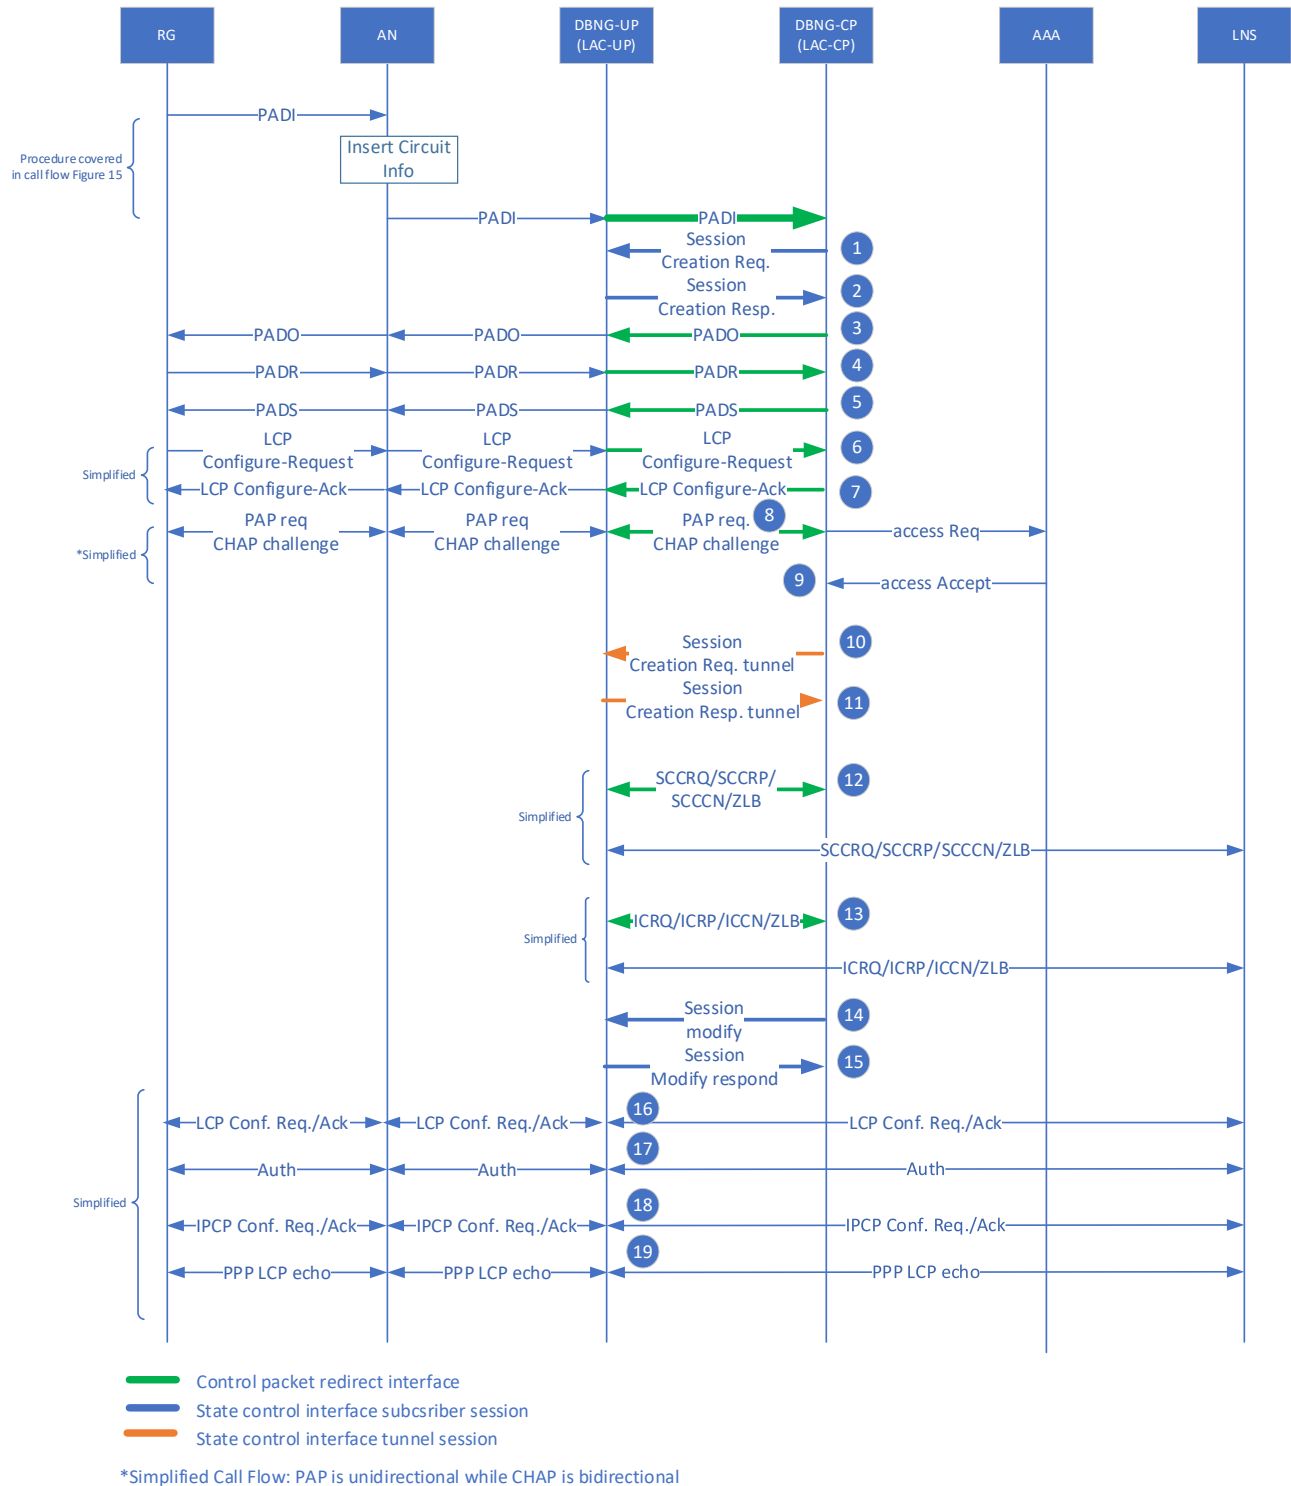


Figure 44: LAC immediate session creation call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common control packet redirection rule.

For Step 1-8, please refer to Section 4.5.16 Steps 1-8.

9. The AAA successfully authenticates the RG and determines that subscriber must be tunneled to an LNS.

Please note step 10-12 would be omitted if a prior L2TP tunnel exists between the LAC and the LNS.

10. The DBNG-CP programs the redirection of the L2TP control packets coming from the LNS towards the DBNG-CP itself over a per-tunnel CPR interface. As explained in section 6.5.4, this per-tunnel CPR interface is programmed bidirectionally by the DBNG-CP on the DBNG-UP.
11. The DBNG-UP sends a response back to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward the L2TP control packets.
12. The DBNG-CP exchanges Start-Control-Connection-Request (SCCRQ), Start-Control-Connection-Reply (SCCRP), Start-Control-Connection-Connected (SCCCN), and Zero-Length Body (ZLB) with the LNS via the DBNG-UP through the CPR interface.
13. The DBNG-CP sends Incoming-Call-Request (ICRQ), Incoming-Call-Reply (ICRP), Incoming-Call-Connected (ICCN), and ZLB to the LNS via the DBNG-UP through the CPR interface.
14. The DBNG-CP programs the DBNG-UP with a Session Modification Request, so that the PPP data packet of the session are forwarded to the LNS over the L2TP tunnel; regarding the PPP control packets, they are also encapsulated over L2TP and sent towards the LNS, except the PADT packets, which are redirected to the DBNG-CP.
15. The DBNG-UP sends a response back to the DBNG-CP, informing that the rules are installed, and the DBNG-UP is ready to forward subscriber's PPP control and data packets.
16. If the LNS cached the LCP Configure-Request and there is no negotiation disagreement, this step can be skipped. If LNS has not cached LCP Configure-Request or the session requires renegotiation, then LCP negotiation takes place.
17. If the LNS had cached authentication information and there is no disagreement on authentication, this step can be skipped. If LCP had not cached authentication information or authentication failed, then a session re-authentication is required.
18. IPCP takes place between the RG and the LNS through the DBNG-UP.
19. PPP LCP echo requests/replies are exchanged between the RG and the LNS through the DBNG-UP.

Note that LAC-UP does not handle any PPP LCP Keepalives at any point.

4.5.23 LAC delayed session creation

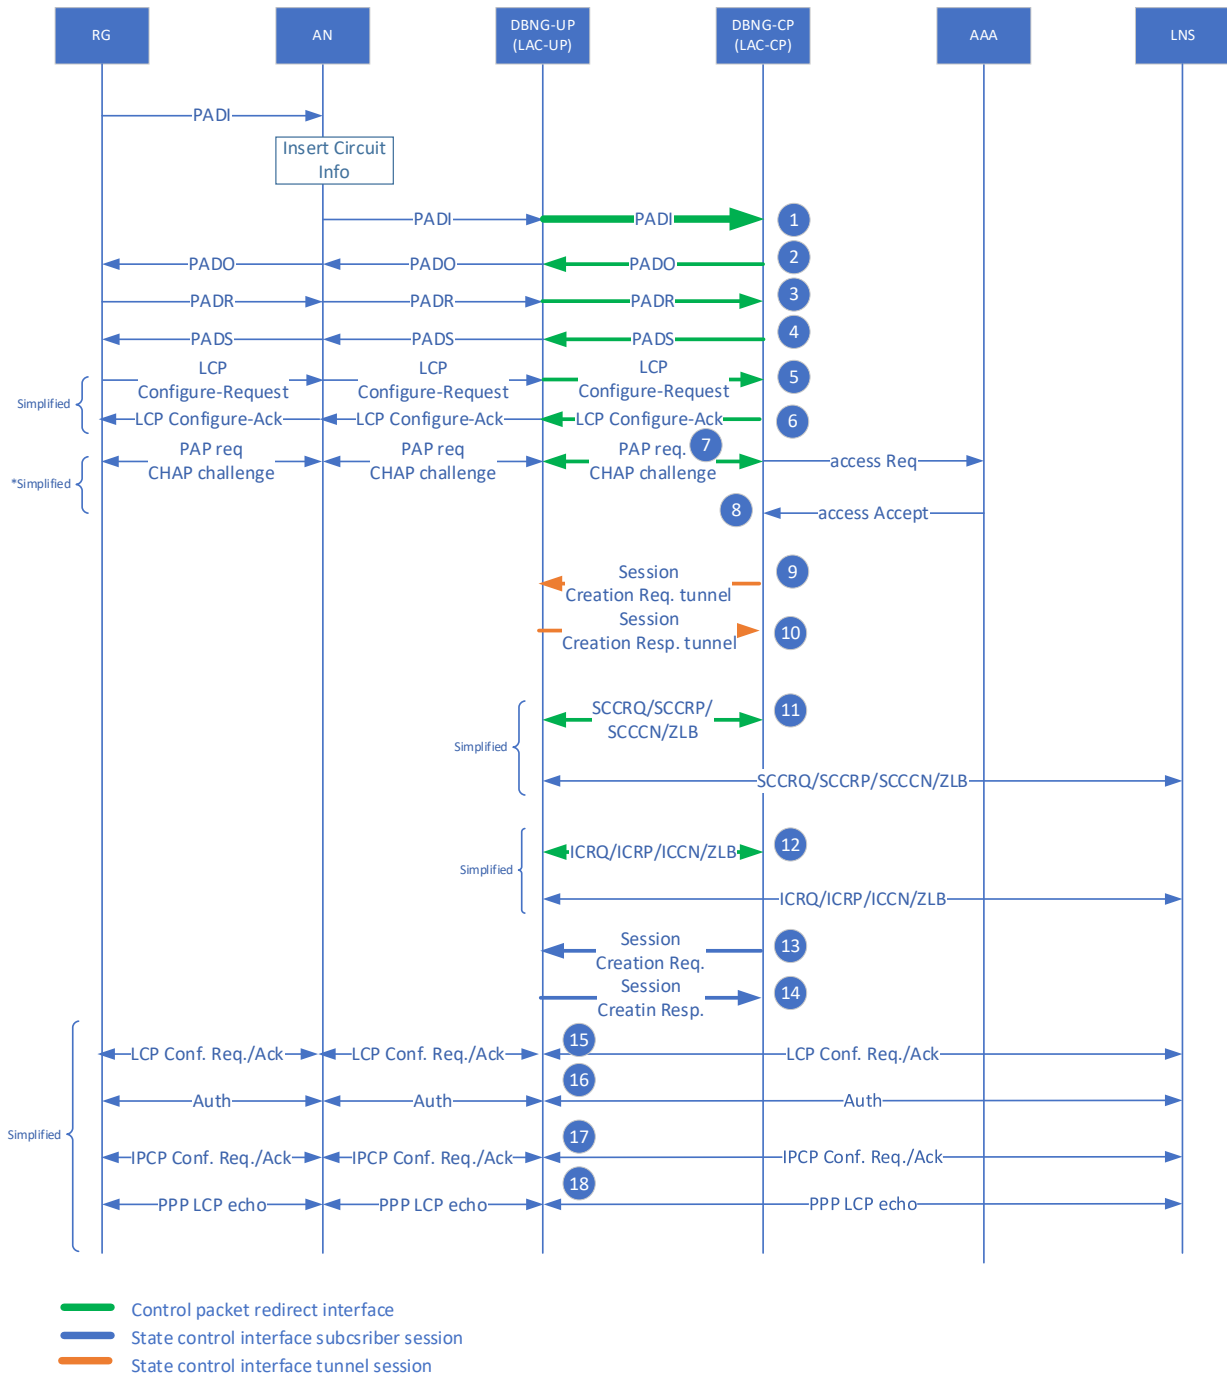


Figure 45: LAC Delayed Session Creation call flow

When the DBNG-CP takes the role of a LAC-CP, it can delay the session creation on DBNG-UP until the moment when the tunnel is chosen or created. This differs from the other cases, where the assignment of the

address to the RG is the last moment until the session creation can be delayed. This is due to the fact that a LAC does not assign the final user address.

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions.

This call flow starts with steps 1-8 of section 4.5.17, followed by L2TP tunnel creation steps 10-13 of section 4.5.22 after which the subscriber session is created. Until the session is created, all the control packets will use default redirection rules for upstream and downstream.

Note: If a prior tunnel is already established between the LAC and the LNS, steps 9-11 can be omitted.

4.5.24 LNS – PPPoEv4 immediate session creation

The LAC DBNG-CP and DBNG-UP split is not relevant in the LNS call flow

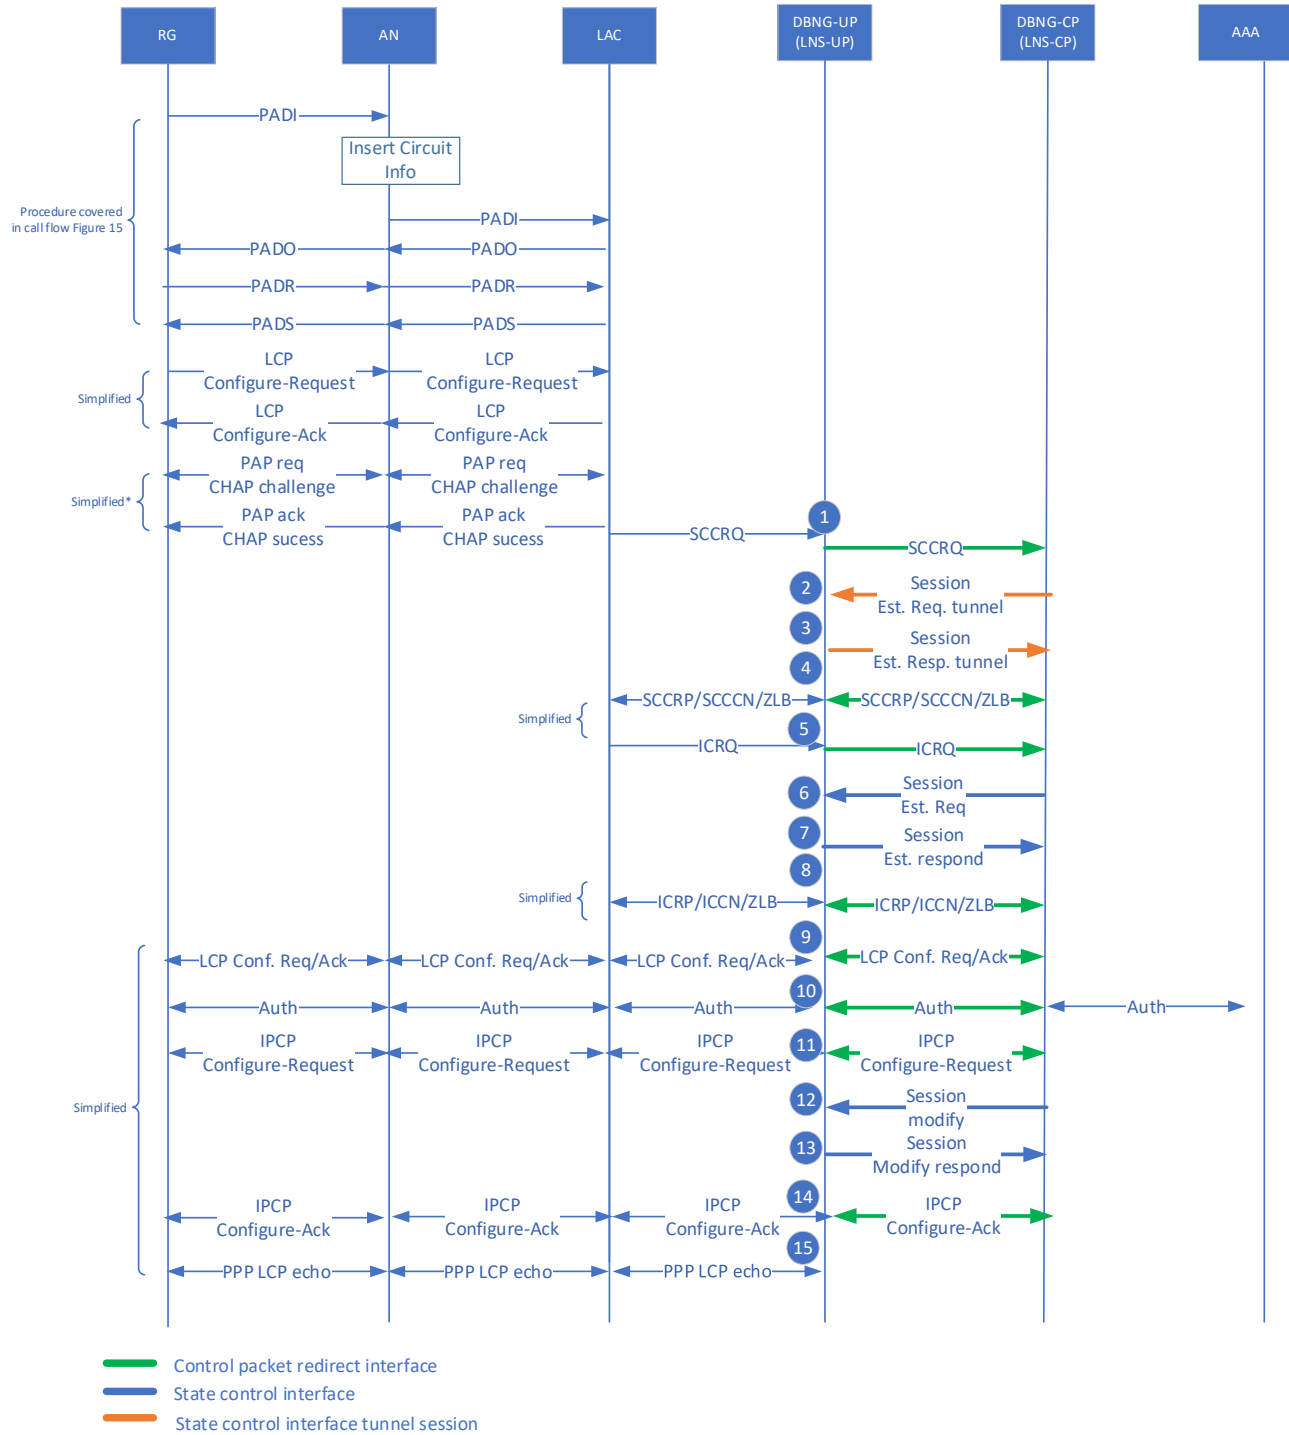


Figure 46: LNS PPPoEv4 immediate session creation call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common control packet redirection rule.

1. The SCCRQ message is received through the CPR interface following the common packet redirect rule.
2. DBNG-CP sends a session establishment request message to the DBNG-UP. The DBNG-CP programs the DBNG-UP control packet redirect rules to send L2TP control message towards the DBNG-CP to only accept particular tunnels.
3. The DBNG-UP sends a response back to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward the L2TP control packets.
4. The DBNG-CP exchanges SCCRP, SCCCN, and ZLB with the LAC by utilizing the CPR interface
5. The DBNG-CP receives the ICRQ message (including the AVP defined in RFC 5515 [42])
6. Upon receiving the ICRQ message, the DBNG-CP has the L2TP session ID information. The DBNG-CP can send a session establishment request to the DBNG-UP to ensure only known L2TP session are accepted.
7. The DBNG-UP sends a response back to the DBNG-CP, informing that the states are installed, and the DBNG-UP only accepts L2TP control packet from known sessions.
8. The DBNG-CP exchanges ICRP, ICCN, and ZLB with the LAC by utilizing the CPR interface
9. If the LNS had cached LCP Configure-Request and there is no negotiation disagreement, this step can be skipped. If LCP had not cached LCP Configure-Request or the session requires renegotiation, then LCP negotiation takes place.
10. If the LNS had cached authentication information and there is no disagreement on authentication, this step can be skipped. If LCP had not cached authentication information or authentication failed, then a session renegotiation takes place.
11. Both DBNG and RG sends IPCP Configure-Request for parameter negotiation, utilizing a dedicated session control packet redirect tunnel. The RG is assigned an Ipv4 address. Address could be assigned to the RG either through the AAA reply or through a local address server.
12. The DBNG-CP sends a session modify request to update the data plane state.
13. The DBNG-UP sends a response back to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward subscriber's PPP control and data packets.
14. The IPCP Configure-Ack is sent from the DBNG-CP to the RG through the DBNG-UP utilizing a dedicated session control packet redirect tunnel.
15. PPP LCP echo requests/replies are exchanged between the RG and the LNS through the DBNG-UP

4.5.25 LNS – PPPoEv4 delayed session creation

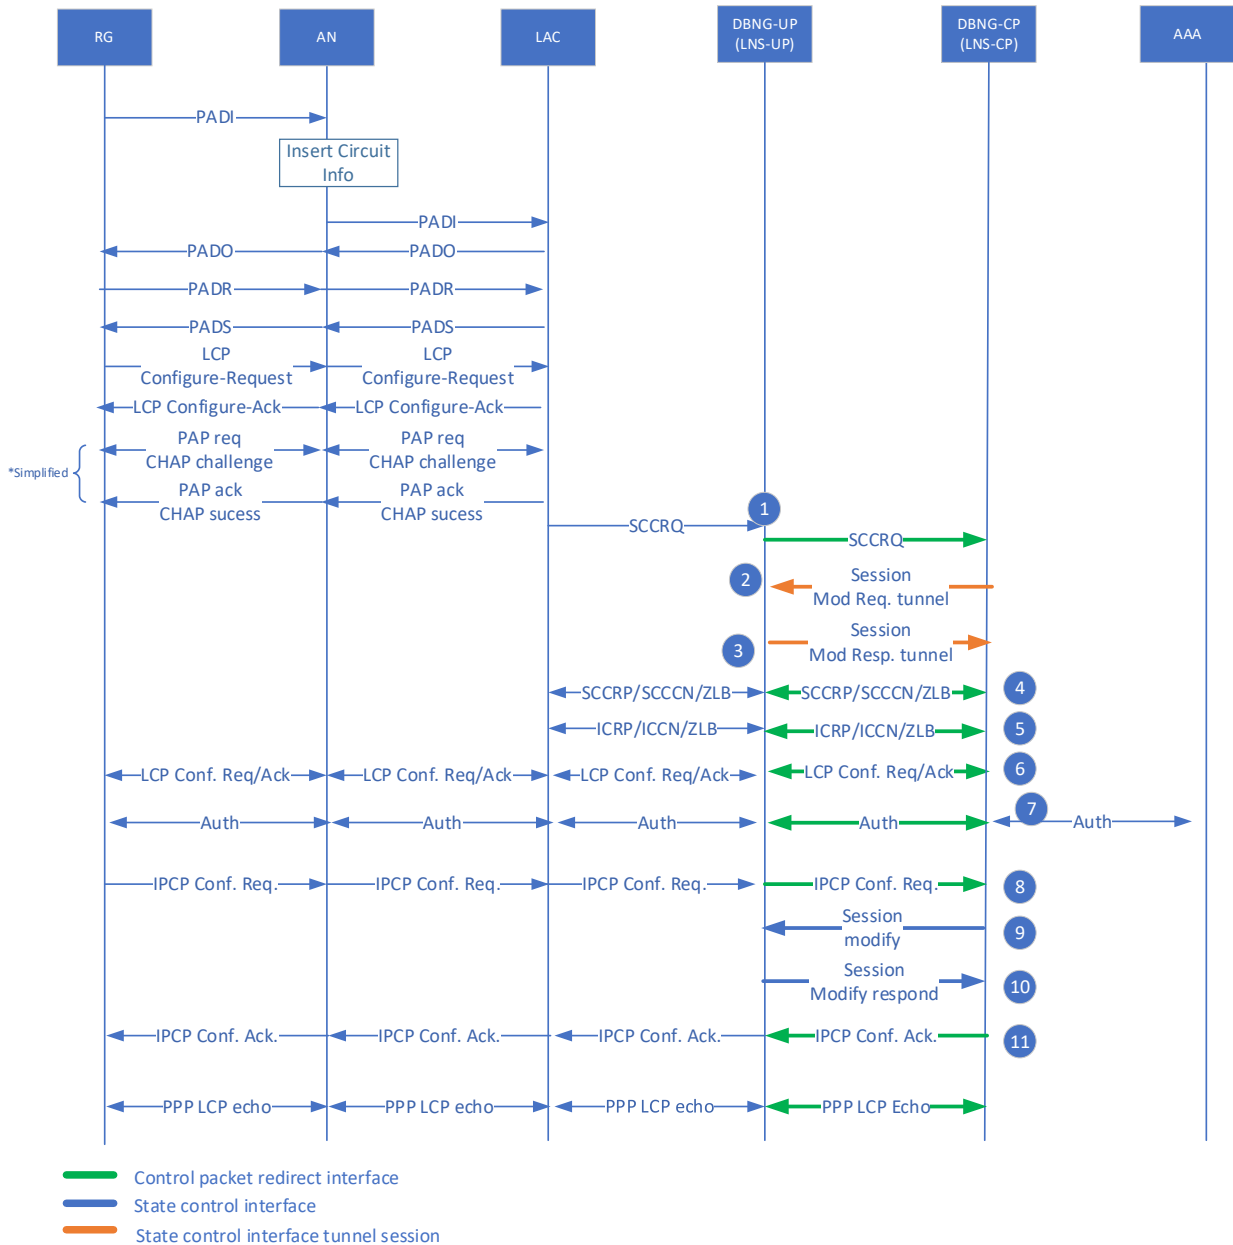
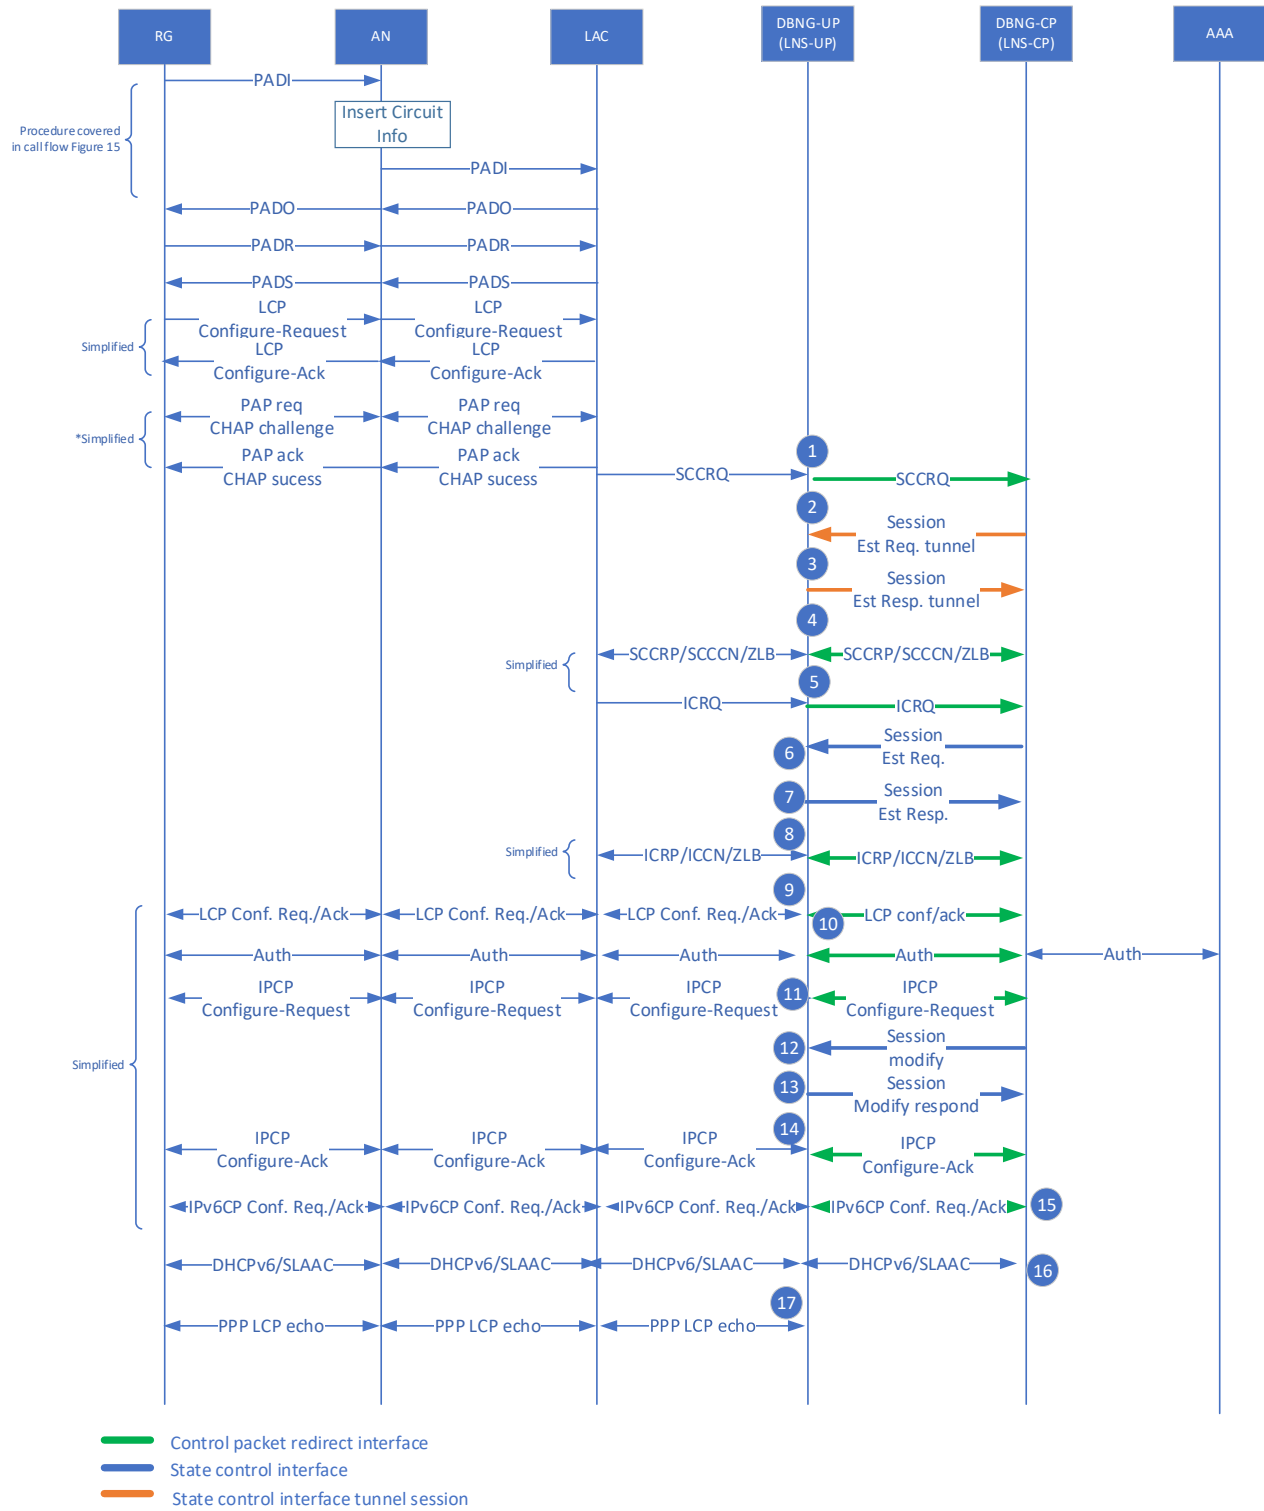


Figure 47: LNS Ipv4 Delayed Session Creation call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions.

This call flow describes the case where the steps 6 & 7 of section 4.5.24 are delayed until step 8 to ensure the L2TP session is created only after IP address is allocated for the RG.

4.5.26 LNS – Dual Stack immediate session creation



*Simplified Call Flow: PAP is unidirectional while CHAP is bidirectional

Figure 48: LNS Dual Stack immediate session creation call flow

Procedure 1 to 14 would follow the same LNS procedure outlined in section 4.5.24:

15. The Ipv6CP Configure-Request is sent from the RG through the DBNG-UP utilizing the CPR interface. The DBNG-CP sends the Ipv6CP Configure-Ack to the RG through the DBNG-UP utilizing the CPR interface.
16. In the case of SLAAC prefix assignment, the DBNG-CP sends an RA to the RG informing the Link Local Address, for more detail please refer to section 4.5.10. In the case of DHCPv6 assignment, please refer section 4.5.7.
17. PPP LCP echo requests/replies are exchanged between the RG and the LNS through the DBNG-UP

4.5.27 LNS – Dual Stack delayed session creation

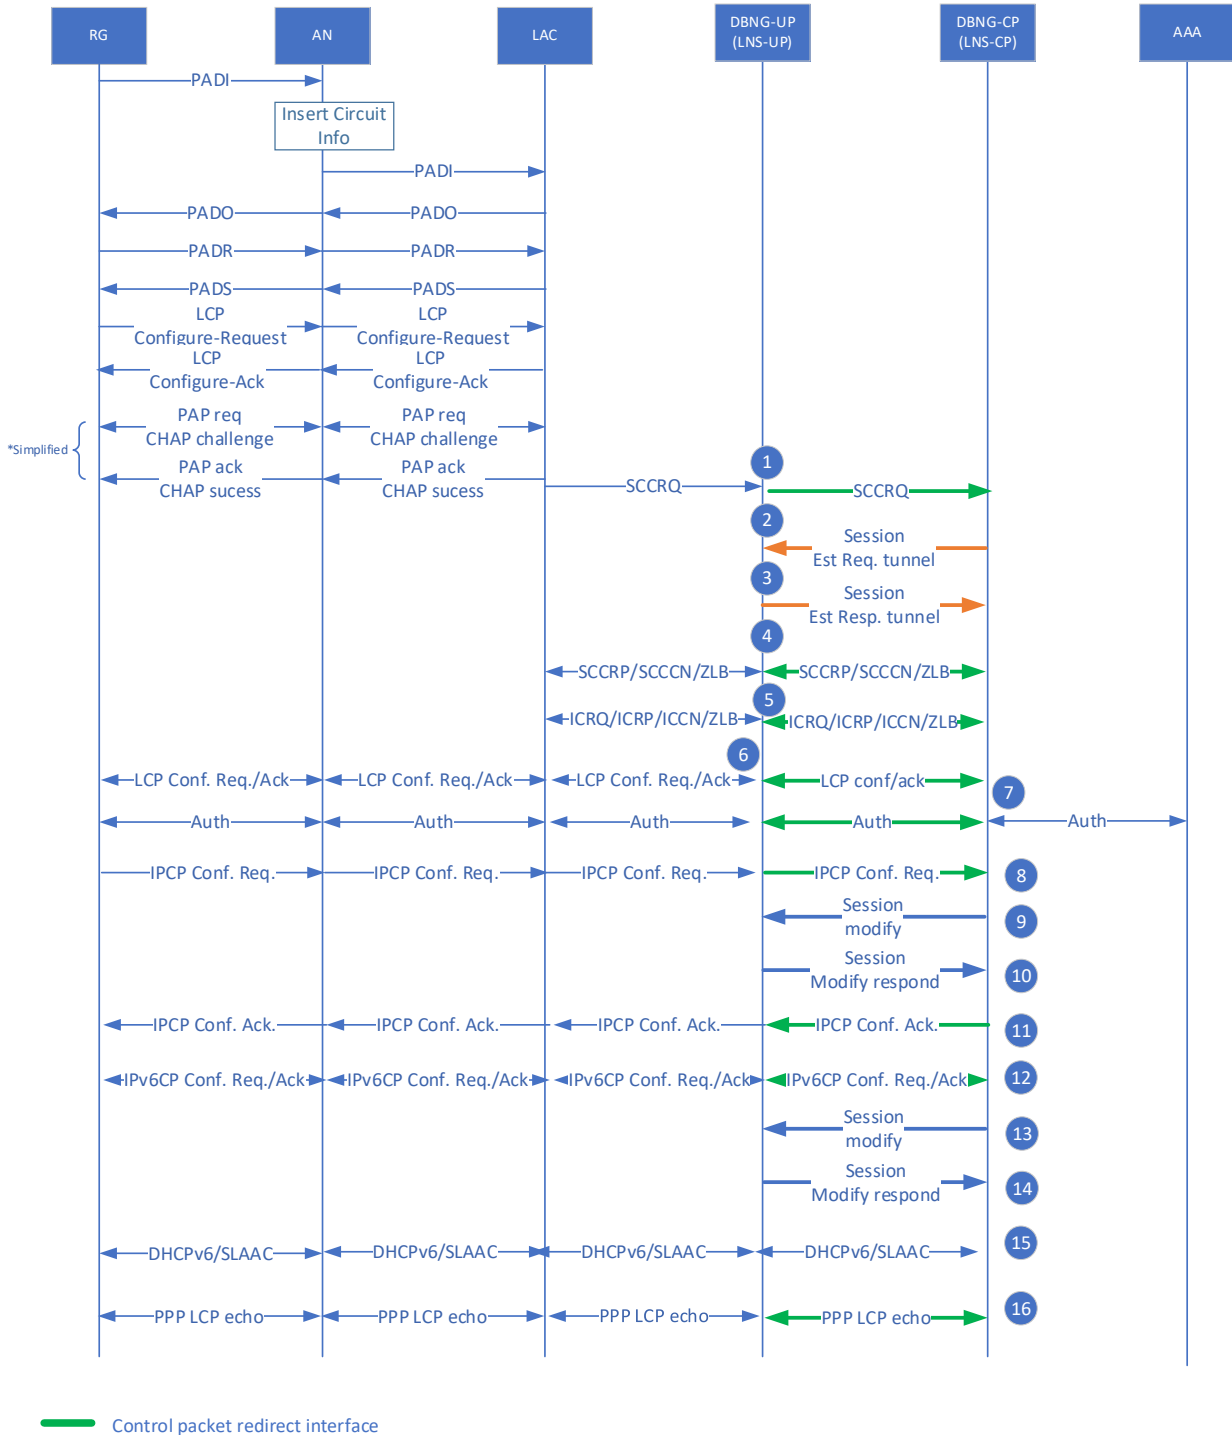


Figure 49: LNS Dual Stack Delayed Session Creation call flow

This case is similar to section 4.5.25, but covers the dual stack case. Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions.

4.5.28 Public Wi-Fi Access

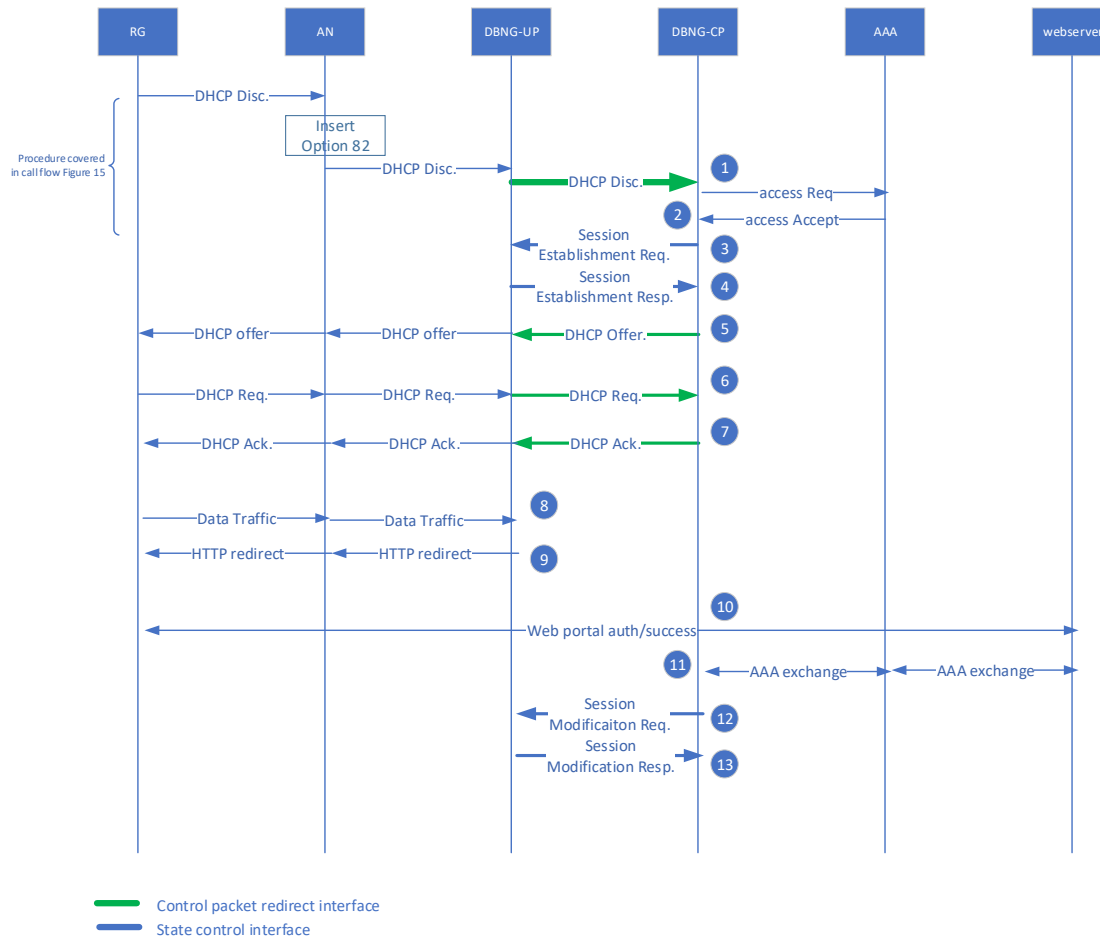


Figure 50: Public Wi-Fi Access call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common control packet redirection rule.

1-7: Follows the same procedure as section 4.5.4 step 1-7.

Note in step 3, the DBNG-CP also installs forwarding rules on the DBNG-UP to redirect http data packets to a web server

8. As user traffic arrives at the DBNG-UP, the packet is redirected for web authentication.

9. The DBNG must inform RG of HTTP redirection, this can be done by the DBNG-UP or the DBNG-CP. RG then sends the traffic directly to the designated web server for authentication.

10. Subscriber successfully authenticates

11. AAA updates DBNG-CP to allow subscriber internet access and removes the http redirection rule for the subscriber.

12. The DBNG-CP sends a session modification message to allow the subscriber internet access and removes the HTTP rule from the DBNG-UP.

13. DBNG-UP sends a session modification response to the DBNG-CP

4.5.29 Public Wi-Fi Layer 3 Access

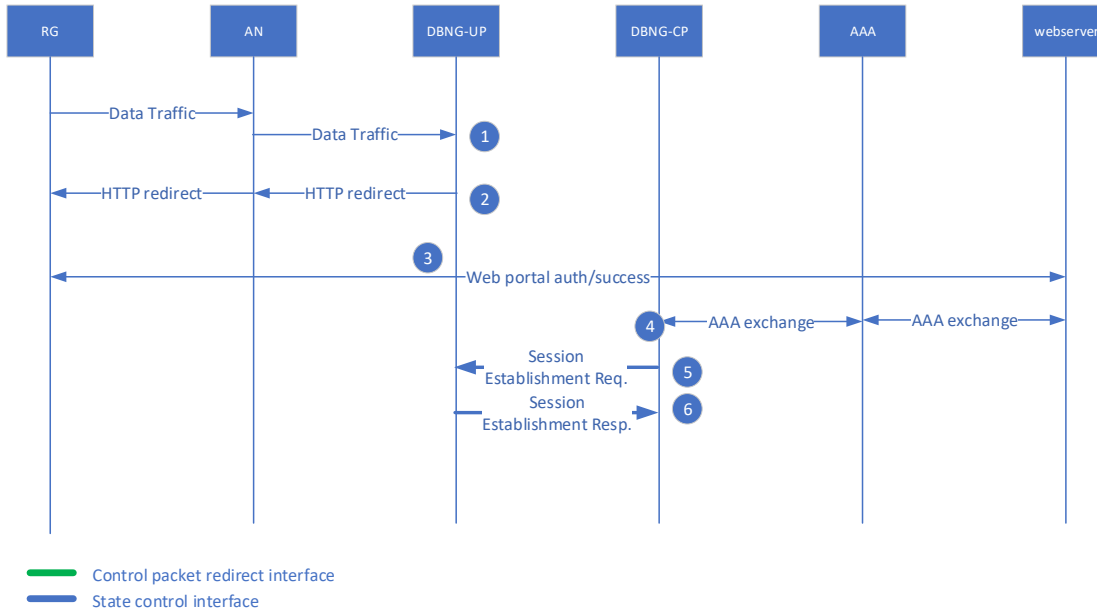


Figure 51: Public Wi-Fi Layer 3 Access call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common control packet redirection rule.
 Step 1-6: Follows the same procedure as section 4.5.28 step 8-13

4.5.30 TWAG Call Flows

The call flow for TWAG follows TR-291 [12] section 11.1 for S2a, where the TWAG and TWAP are integrated into the DBNG, the BBF access is considered as Trusted by the 3GPP network and single-connection mode with Enhanced Packet Core (EPC) access is provided to the UE. The DBNG terminates an S2a interface with the 3GPP PGW.

In this call flow, the TWAG is split up into DBNG-CP and DBNG-UP. Therefore, additional steps are added for DBNG-CP and DBNG-UP communications and packet redirection. Prior to step 1, a common rule would be installed on the DBNG-UP to redirect all control messages (such as DHCP, RS, and DHCPv6) to the DBNG-CP, see section 4.5.2. The procedure related to interaction with an external system stays the same as in TR-291 [12]. This procedure requires the GTP-c and GTP-u to utilize different IP address endpoints.

NOTE: 3GPP S2a signaling already supports that the Fully Qualified Tunnel Endpoint Identifier (F-TEID) for Control Plane can correspond to a different IP address than the S2a-U TWAN F-TEID in a Bearer Context (User Plane).

4.5.30.1 S2a initial attach based on layer 2 trigger: Ipv4 based on DHCPv4

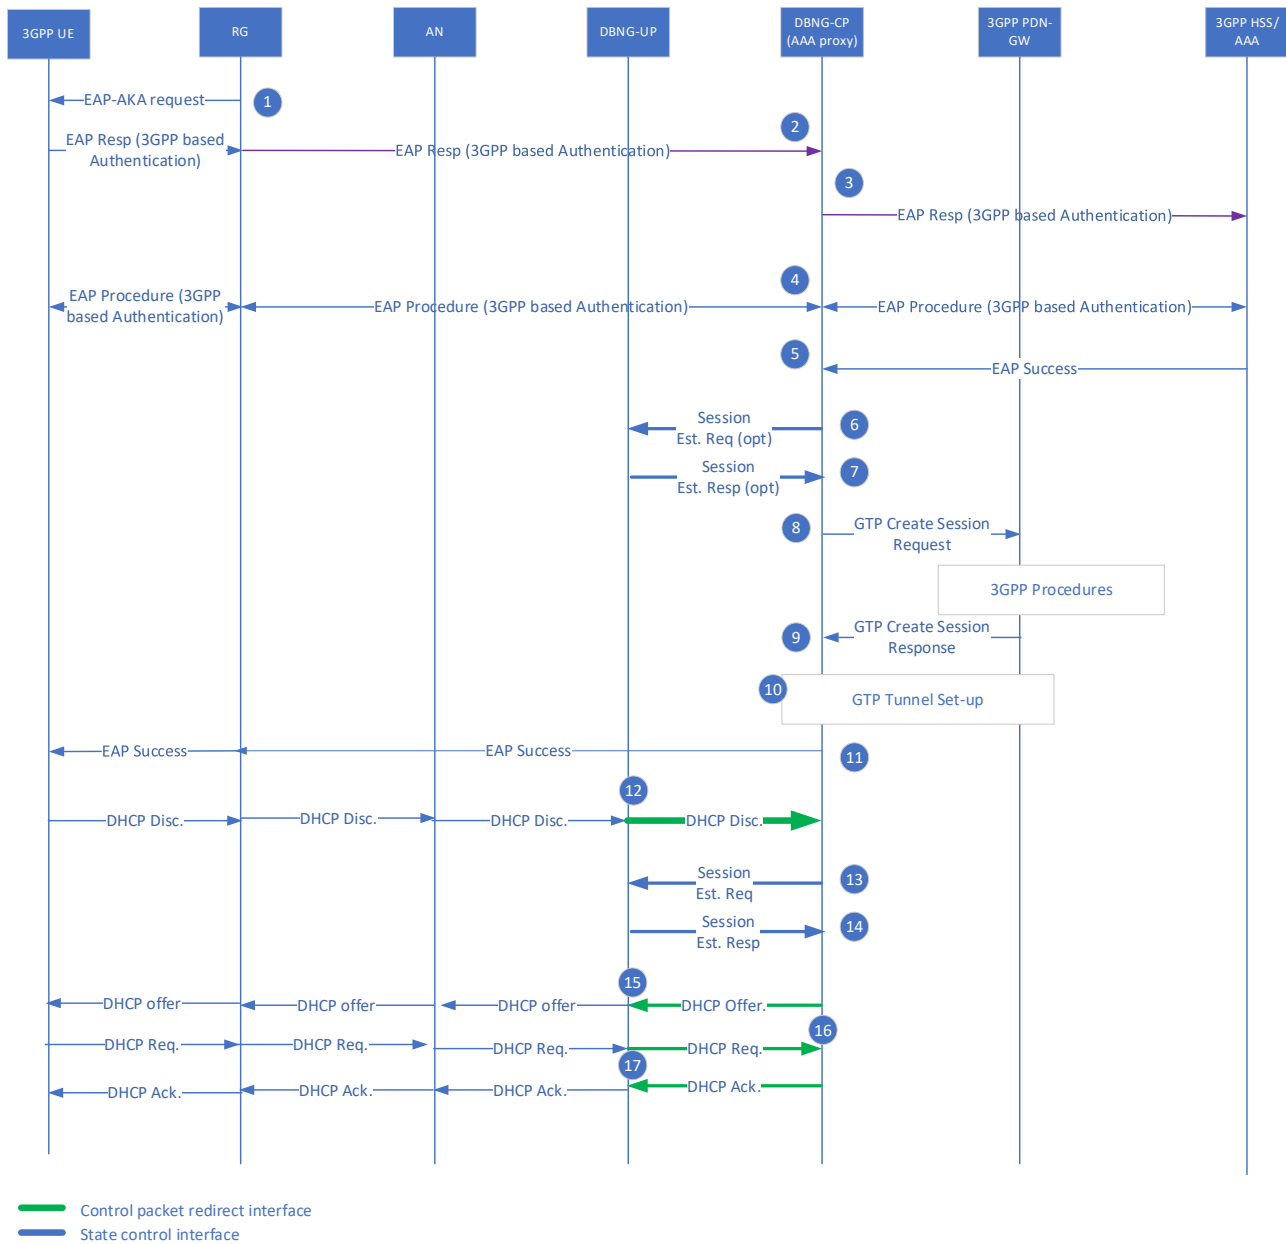


Figure 52: S2a initial attached based on layer 2 trigger: DHCPv4

1. The 3GPP UE attaches to the BBF access network. The RG initiates an EAP request to the 3GPP UE and thus initiates the EAP authentication process (see 3GPP TS 33.402 [30]). During the authentication phase, the RG acts as an 802.1X authenticator, and adds the MAC Address of the 3GPP UE to the RADIUS Request message.
2. The RG forwards the EAP response to the DBNG: the DBNG is invoked as an AAA proxy and DHCP address server for the 3GPP UE.

When the TWAG is deployed in a dedicated router, the DBNG is involved as an AAA proxy and for the 3GPP UE, distinguishing 3GPP UE signaling from fixed device signaling based on Network Access Identifier (NAI).

3. The DBNG-CP forwards the EAP response to the 3GPP AAA. 3GPP procedures between 3GPP AAA server and Home Subscriber Server (HSS) take place as described in TS 23.402 [27] clause 16.2 and 33.402 [30].
4. The DBNG-CP acting as a AAA proxy relays back and forth EAP signaling between the 3GPP AAA server and the RG.
5. Once the 3GPP UE successfully authenticates, the 3GPP AAA server creates an EAP-Success that it embeds into a AAA-Success sent to the DBNG-CP; The 3GPP UE is now authenticated, and the BBF access is considered as trusted by the 3GPP Network.

The rest of the procedure assumes that the DBNG-CP has been instructed by the AAA to provide EPC access (e.g., to establish an S2a GTP tunnel).

6. [Optional]* In the case where the TEID for User plane is assigned by the DBNG-UP, the DBNG-CP initiates a DBNG-UP Control session establishment request to retrieve the TEID for the GTP-u tunnel at DBNG-UP network (PGW) side. (3GPP TS 29.244 [28] section 5.5.3)
7. [Optional]* The DBNG-UP responds with a DBNG-UP Control session establishment response supplying the requested TEID. (3GPP TS 29.244 [28] section 5.5.3)
8. The DBNG-CP sends a GTP-c Create Session Request message to the PDN GW. The DBNG-CP selects the PDN GW and builds the Create Session Request as defined in 3GPP TS 23.402 [27] clause 16.2 8b; 3GPP procedures possibly including a Policy and Charging Rules Function (PCRF) take place. As a result, an IP address is allocated to the UE. It must be noted that the GTP-c and GTP-u for S2a can utilize two different IP endpoints.
9. The PDN GW returns a Create Session Response, including the IP address allocated for the 3GPP UE.
10. Both DBNG-CP and PGW finishes exchanging information to establish the GTP-u tunnel.
11. The DBNG-CP proxies the EAP Success message to the RG through the CPR interface. The RG sends the EAP Success to the 3GPP UE. The 3GPP UE is now authenticated, and the BBF access is considered as trusted by the 3GPP Network.
12. The 3GPP UE sends a DHCP Discovery message and is redirected to the DBNG-CP through the DBNG-UP.
13. The DBNG-CP provides the DBNG-UP with packet forwarding rules based on the DHCP discovery packet (which could include the encapsulation, MAC, and VLANs) and on the GTP-u F-TEID and QoS requirements received from the PDN GW in Create Session Response. These rules include the request to forward DHCP signaling from the UE to the DBNG-CP. They are sent in a session establishment request unless a session establishment request has been used in step 6, in which case a session modification request would be used.
NOTE: The DBNG-CP is responsible of the charging interface, if any.
14. The DBNG-UP sends a response back to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward the subscribers IP data packets
15. The DBNG-CP sends a DHCP offer including the Ipv4 Address allocated by the PDN GW to the 3GPP UE via the CPR interface.
16. The 3GPP UE sends a DHCP request message and is redirected to the DBNG-CP through the DBNG-UP via the CPR interface.
17. The DBNG-CP sends a DHCP ack through the CPR interface.

Note: The procedure for accounting message exchange used for charging purposes is not included in this flow.

Note: TEID allocation can be done by the DBNG-CP function (3GPP TS 29.244 [28] section 5.5.2) or optionally by the DBNG-UP function (3GPP TS 29.244 [28] section 5.5.3)

4.5.30.2 S2a initial attach based on layer 2 trigger: Ipv6 prefix based on SLAAC

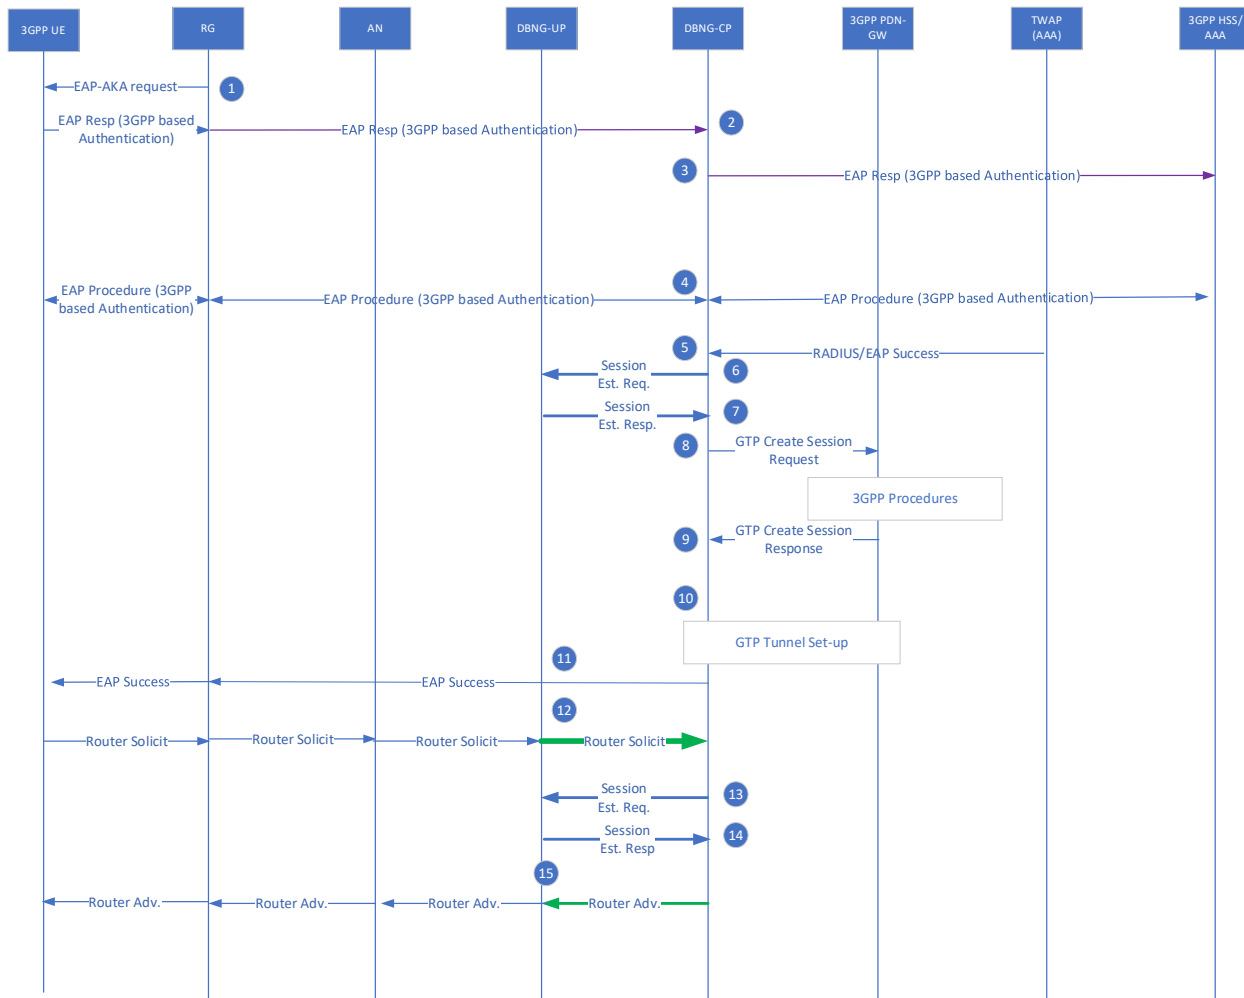


Figure 53: S2a initial attached based on layer 2 trigger: SLAAC

Procedure 1 to 11 would follow the same DHCPv4 procedure outlined in section 4.5.30.1

12. The 3GPP UE sends a router solicitation message and is redirected to the DBNG-CP through the DBNG-UP.

13. The DBNG-CP provides the DBNG-UP with packet forwarding rules based on the router solicitation packet (which could include the encapsulation, MAC, and VLANs) and on the GTP-u F-TEID and QoS requirements received from the PDN GW in Create Session Response. These rules include the request to forward router solicitations from the UE to the DBNG-CP. They are sent in a session establishment request unless a session establishment request has been used in step 6, in which case a session modification request would be used.

NOTE: The DBNG-CP is responsible of the charging interface, if any.

14. The DBNG-UP sends a response back to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward the subscribers IP data packets

15. The DBNG-CP sends a router advertisement including the Ipv6 prefix allocated by the PDN GW to the 3GPP UE via the CPR interface.

Note: The procedure for accounting message exchange used for charging purposes is not included in this flow.

4.5.30.3 S2a initial attach based on layer 3 trigger: Ipv4 based on DHCPv4

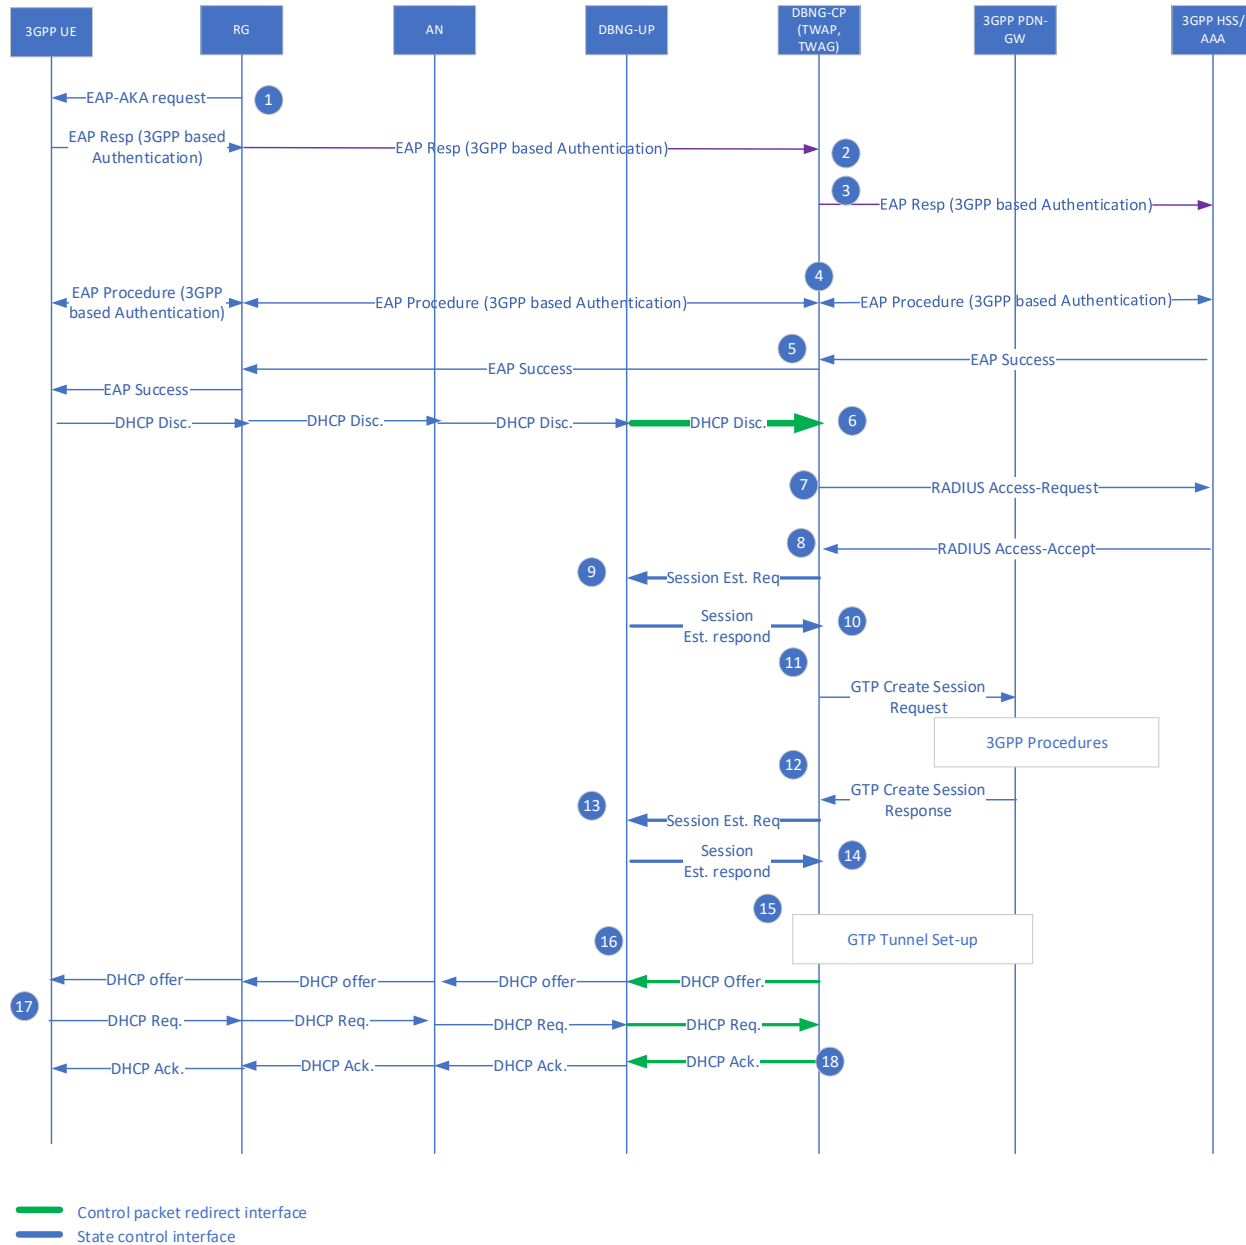


Figure 54: S2a initial attached based on layer 3 trigger: DHCPv4

1. The 3GPP UE attaches to the BBF access network. The RG initiates an EAP request to the 3GPP UE and thus initiates the EAP authentication process (see 3GPP TS 33.402 [30]). During the authentication phase, the RG acts as an 802.1X authenticator, and adds the MAC Address of the 3GPP UE to the RADIUS Request message. The DBNG is invoked as an AAA proxy and address server for the 3GPP UE.
2. When the DBNG is deployed as a dedicated router, the DBNG acts as an AAA proxy and for the 3GPP UE, distinguishing 3GPP UE signaling from fixed device signaling based on NAI. During the

authentication phase, the RG acts as an 802.1X authenticator, and adds the MAC Address of the 3GPP UE to the RADIUS message sent to the TWAG.

3. The DBNG-CP forwards the EAP response to the 3GPP AAA. 3GPP procedures between 3GPP AAA server and HSS take place as described in 3GPP TS 23.402 [27] clause 16.2 and 33.402 [30].
4. The DBNG-CP acting as a TWAP relays back and forth EAP signaling between the 3GPP AAA server and the RG.
5. Once successfully authenticated, the 3GPP AAA server creates an EAP-Success and sends to the DBNG-CP; The 3GPP UE is now authenticated, and the BBF access is considered as trusted by the 3GPP Network.
6. The 3GPP UE sends a DHCP Discover message including the MAC Address. The RG Relays the DHCP Discover message to the DBNG-CP through the DBNG-UP.
7. The DBNG-CP sends a RADIUS Access-Request to the AAA, including the MAC Address. The TWAG makes use of the MAC Address, which is stored during the authentication phase of the 3GPP UE, for correlating the information obtained from the 3GPP Domain during the authentication phase (Step 2) with the IP session.
8. The AAA responds with a RADIUS Access-Accept to the DBNG-CP, including an indication of the need to establish an S2a connection between the DBNG-UP and the 3GPP PDN GW. TWAP also provides information retrieved from the 3GPP Domain at Step 1, such as the APN, the selected PLMN Id and the 3GPP IMSI.

Note: Steps 7 and 8 may be avoided if the DBNG is a RADIUS proxy during the 3GPP UE authentication performed during Step 2.

9. [Optional]* In the case where the TEID for User plane is assigned by the DBNG-UP, the DBNG-CP must initiate a session establishment request to retrieve the TEID for the GTP-u tunnel at its network side.
10. [Optional]* The DBNG-UP responds with a session establishment response supplying the TEID for the PGW GTP-u tunnel to use.
11. If the DBNG-CP is instructed by the AAA to establish an S2a GTP tunnel, the DBNG-CP sends a Create Session Request message to the PDN GW. The DBNG-CP selects the PDN GW and builds the Create Session Request as defined in 3GPP TS 23.402 [27] clause 16.2 8b; 3GPP procedures possibly including a PCRF take place. As a result, an IP address is allocated to the UE. It must be noted that the GTP-c for S2a and GTP-u can utilize two different IP endpoints.
12. The PDN GW returns a Create Session Response, including the IP address allocated for the 3GPP UE.
13. The DBNG-CP provides the DBNG-UP with packet forwarding rules based on the DHCP discovery packet at step 6 (which could include the encapsulation, MAC, and VLANs) and on the GTP-u F-TEID and QoS requirements received from the PDN GW in Create Session Response. These rules include the request to forward DHCP signaling from the UE to the DBNG-CP. They are sent in a session establishment request unless a session establishment request has been used in step 6, in which case a session modification request would be used.
NOTE: The DBNG-CP is responsible of the charging interface, if any.
14. The DBNG-UP sends a response back to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward the subscribers IP data packets
15. Both DBNG and PGW finishes exchanging information to establish the GTP-u tunnel.
16. The DBNG-CP sends a DHCP offer including the Ipv4 Address allocated by the PDN GW to the 3GPP UE via the CPR interface.
17. The 3GPP UE sends a DHCP request message and is redirected to the DBNG-CP through the DBNG-UP via the CPR interface.
18. The DBNG-CP sends a DHCP ack through the CPR interface

Note: The procedure for accounting message exchange used for charging purposes is not included in this flow.

Note: TEID allocation can be done by the DBNG-CP function (3GPP TS 29.244 [28] section 5.5.2) or optionally by the DBNG-UP function (3GPP TS 29.244 [28] section 5.5.3)

4.5.31 Hybrid Access Gateway

The call flow covers Hybrid Access Gateway L3 Network-based Tunneling in both TR-348 [16] and TR-378 [17].

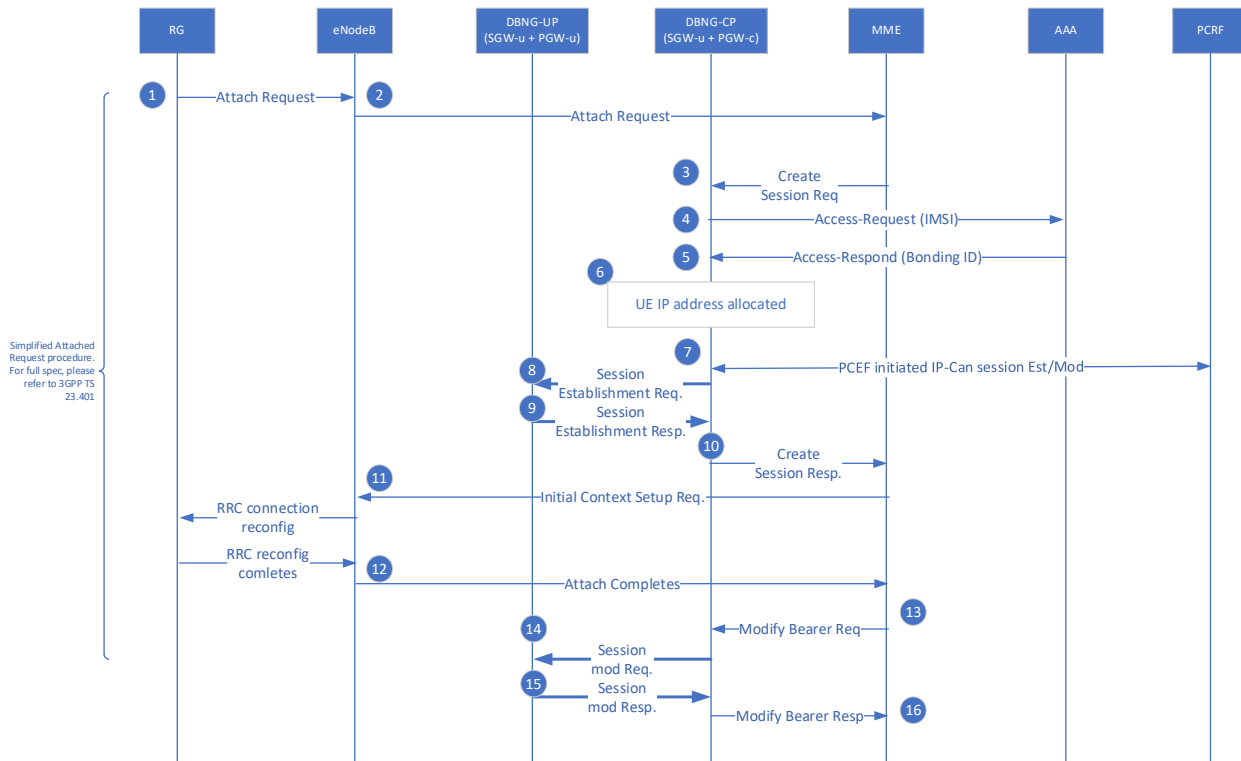


Figure 55: Hybrid Access Gateway L3 network-based Tunneling call flow

The above RG attach procedure is simplified, for full attachment procedure, refer to 3GPP TS 23.401 [26]. The HAG is a BBF defined element in TR-348 [16] where PGW, SGW, and MS-BNG are integrated into a single element that can serve both wireline and wireless access. For some deployment cases, it is possible for only the PGW and MS-BNG to be integrated into the HAG. The HAG in the call flow procedure is split between the DBNG-CP and DBNG-UP.

The Hybrid CPE (HCPE) can start with either the wireless or wireline connection. To bond the two connections, a policy or a local configuration is required. As an example, a RADIUS attribute named Bonding ID returned by the AAA.

- For wireline: during the DHCP or PPPoE authentication process the AAA server returns the RADIUS attribute: Bonding ID.
- For wireless: during the create session request procedure, the DBNG-CP sends an Access-Request to the AAA server which returns the RADIUS attribute: Bonding ID.

Utilizing the bonding ID, the DBNG-CP can correlate the wireline and wireless session as a single hybrid session. Regardless of the connection sequence, wireline or wireless, the DBNG-CP can identify the two connections as a single hybrid session.

Additional parameters might be required to load-balance traffic among wireline and wireless access. However, it is important to note that the HAG integration is transparent to other 3GPP components (MME, eNodeB, and PCRF).

Key information regarding load-balancing as documented in TR-378 [17]:

- Downstream:
 - o QoS or policy control on the HAG can be provided by local policies or via RADIUS
- Upstream:
 - o RG traffic load balancing is provided by a pre-defined policy.

This diagram depicts the case where the RG connects to Core over LTE first and then over Wireline

1. The Hybrid RG initiates an Attach Request to the eNodeB including a request for a PDN connection in an ESM container. This takes place as defined in step 1 of Figure 5.3.2.1 within 3GPP TS 23.401 [26].
2. The eNodeB forwards the Attach Request to the MME. This takes place as defined in step 2 of 3GPP TS 23.401 [26] Figure 5.3.2.1. The MME may carry out security features and retrieve subscription data as defined in step 3 to 11 of 3GPP TS 23.401 [26] Figure 5.3.2.1.
3. MME identifies the DBNG (HAG) to host the subscriber session as a PGW, where the SGW is also co-located: if MME has received one in the subscription data from HSS, the MME will select the DBNG in the subscription data, otherwise it selects one using the APN in the subscription data and as defined in 3GPP TS 23.401. The MME sends a session request to the DBNG-CP (acting as both the SGW and PGW from MME viewpoint). This takes place as defined in steps 12 to 13 of 3GPP TS 23.401 [26] Figure 5.3.2.1.
4. The DBNG-CP sends an Access-Request to the AAA server. The Request contains the subscriber IMSI.
5. The AAA server returns the bonding ID for the subscriber session. The AAA might contain other attributes which contains QoS parameter for load-balancing downstream.
6. DBNG-CP selects a DBNG-UP and allocates an IP address for the Hybrid RG
7. *[optional] QoS parameter might be provided by the PCRF to the DBNG-CP through the Gx interface.
8. DBNG-CP requests a session establishment from the DBNG-UP. The DBNG-UP could be used to provide the TEID for the GTP-u tunnel.
9. The DBNG-UP responds to the session establishment request.
10. The DBNG-CP sends a create session response back to the MME. This takes place as defined in steps 15 to 16 of 3GPP TS 23.401 [26] Figure 5.3.2.1;
11. The MME sends the NAS attach accept back to the HCPE through the eNodeB via DBNG-CP and DBNG-UP: the NAS attach accept is sent within a S1-AP Initial Context Setup Request as defined in steps 17 of 3GPP TS 23.401 [26] Figure 5.3.2.1.
12. The eNodeB forwards the Attach complete message to the MME through the DBNG-UP and DBNG-CP.
13. The MME sends a modify bearer request to the DBNG-CP. This informs the DBNG-CP the TEID that the eNodeB intends to use.
14. The DBNG-CP sends modify the session to update the DBNG-UP with known information for both uplink and downlink (TEID, RG IP, and the bonding ID)
15. The DBNG-UP respond to the session modification to the DBNG-CP
16. The DBNG-CP sends a modify bearer response to the MME.

For subsequent DHCP, SLAAC, or PPPoE wireline connections, please look at section 4.5 for the call flow. The only modification to the call flows of the wireline access is AAA Access-Accept includes the Bonding ID to allow DBNG-CP to bond the existing wireless session with the new wireline session. This is application to:

- DHCP call flow step 2
- DHCPv6 call flow step 2
- SLAAC call flow step 2
- PPPoE call flow step
- PPPoEv6 call flow step 10

Note: TEID allocation can be done by the DBNG-CP function (3GPP TS 29.244 [28] section 5.5.2) or optionally by the DBNG-UP function (TS 29.244 [28] section 5.5.3)

4.5.32 Lawful Intercept call flows

Support for Lawful Interception will be over SCi PFCP as described in Section 5.7.2 of 3GPP TS 29.244 [28].

Any of the encapsulations specified in Outer Header Creation (IE Type 84) shall be allowed in DBNG for encapsulating mirrored traffic.

To maintain privacy and confidentiality, the lawful interception over PFCP shall only be exchanged over secure transport or by encrypting the Lawful Interception IEs.

Note: Lawful Intercept (LI) is governed by local policy and regulations. This TR provides a basic mechanism for the DBNG to perform lawful intercept. The following call flow is an example which can realize Lawful Intercept using various PFCP mechanisms.

4.5.32.1 Lawful Intercept while subscriber is online

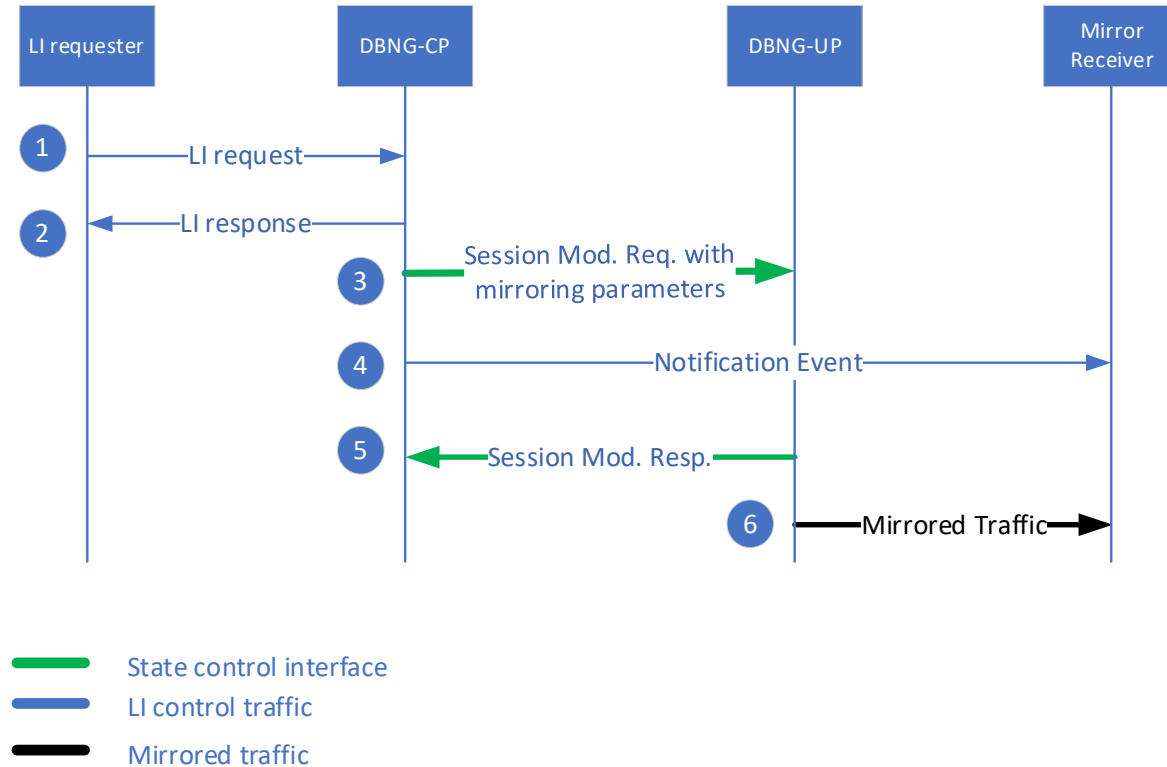


Figure 56: Example of Lawful intercept Request

1. The LI requester (unspecified in this document) contacts the DBNG-CP requesting to mirror traffic for a specific subscriber.
2. DBNG-CP sends back the answer to the LI-requester (note that this may optionally take place after step 4).
3. DBNG-CP performs a lookup for the specific subscriber and identifies the DBNG-UP handling that subscriber. DBNG-CP instructs the identified DBNG-UP and sends a session modification request message with mirroring parameters to the identified DBNG-UP.
4. DBNG-CP provides a LI event notification to the Mirror Receiver
5. DBNG-UP sends a session modification response message to the DBNG-CP.
6. DBNG-UP sends mirrored traffic to the mirror receiver.

If for any reason, the subscriber is moved to a different DBNG User Plane (due to a switchover), then the subscriber mirror configurations should also be moved accordingly.

4.5.32.2 Lawful Intercept while subscriber is offline

The following call flow is applicable when a target subscriber has not logged in yet.



Figure 57: Lawful intercept while subscriber is offline

1. The LI requester (unspecified in this document) contacts the DBNG-CP requesting to trigger mirroring for a specific subscriber.
2. DBNG-CP stores the request in a persistent database and sends back a response to the LI-requester acknowledging the request.

The interface between LI Requester and DBNG-CP in addition to adding triggers can be used to modify, delete, and query existing triggers. LI is performed when the targeted subscriber logs on to the DBNG.

4.5.32.3 Lawful Intercept triggered by AAA authentication

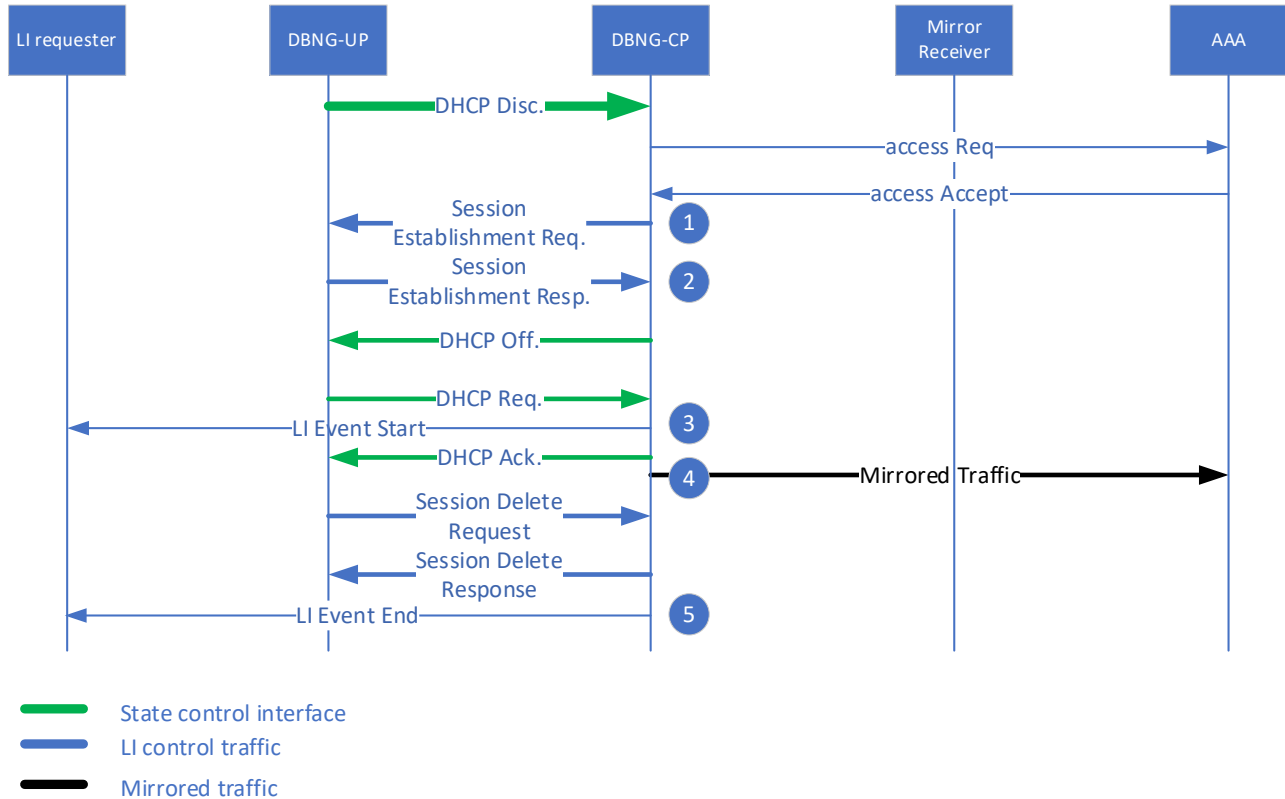


Figure 58: Lawful intercept triggered by AAA authentication

1. During login after AAA interaction, DBNG-CP performs a lookup for the specific subscriber and identifies the mirroring status. DBNG-CP sends a mirror request message to the identified DBNG-UP as part of the Session Establishment Request.
2. DBNG-UP sends a mirror response message to the DBNG-CP.
3. DBNG-CP sends a LI Event Start notification to the LI Requester on successful completion of login.
4. DBNG-UP sends mirrored traffic to the Mirror Receiver.
5. On logout, DBNG-CP sends a LI Event End notification to the LI Requester.

Note:

1. The LI Requester and Mirror Receiver functions could be done by a single Mediation Device function.
2. For each intercepted subscriber session, the DBNG-CP should interact with the Mediation device to inform the DBNG-UP source of the mirrored traffic in the case of subscriber switch over.
3. The example covers LI for immediate session creation. When delayed session is used, the mirror request message can be exchanged using Session Modification Request.

4.5.33 Subscriber session modification

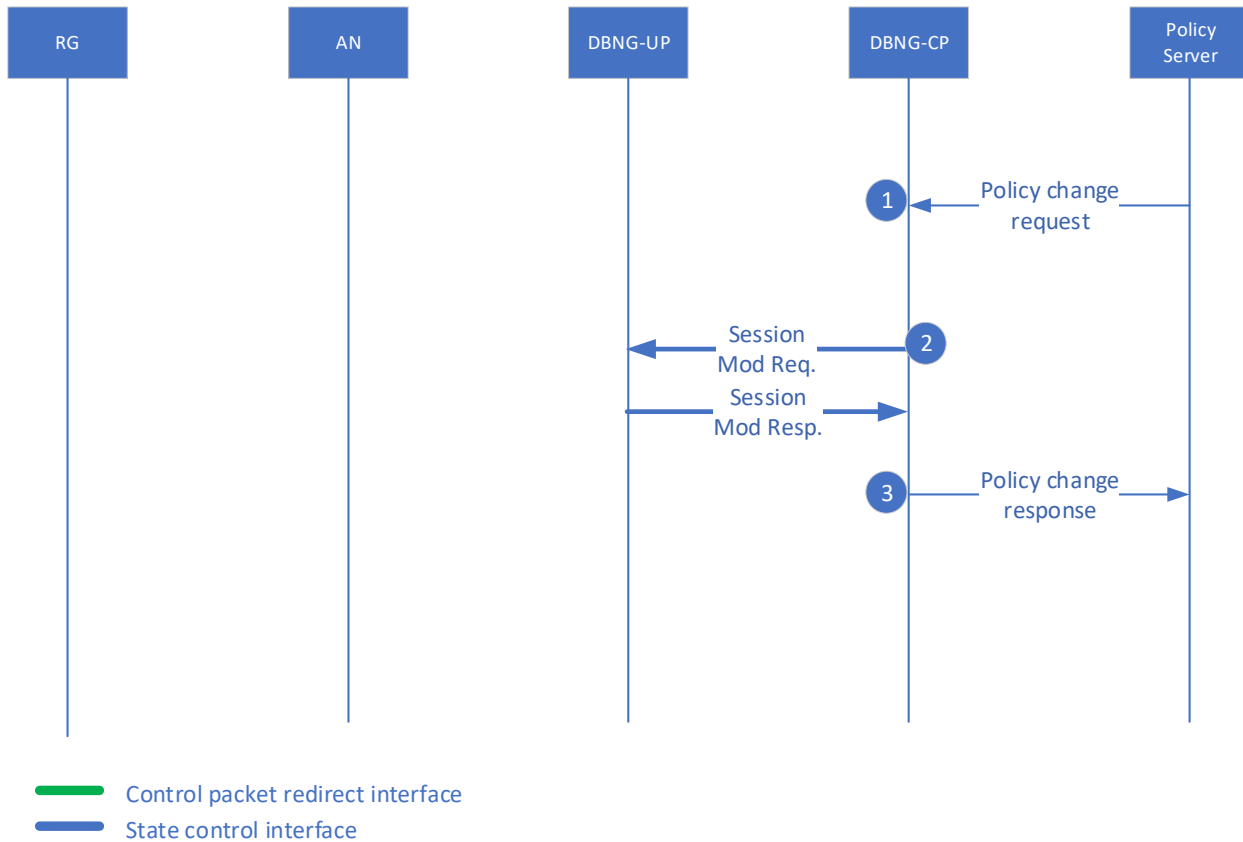


Figure 59: Subscriber session modification call flow

Steps:

1. [conditional] Policy change can be triggered by an external server or locally on the DBNG-CP. If the change is performed locally, skip to step 2. In this case, the DBNG-CP receives a policy request from an external policy server such as a AAA server through RADIUS Change of Authorization message. The change request can be related to a subscriber policy or other subscriber attribute.
 - a. [conditional] If the DBNG-CP cannot comply with the policy change, skip to step 3 to reply to the policy server.
2. Once the DBNG-CP receives the change request for the subscriber session, the DBNG-CP determines if the DBNG-UP requires any update. If no updates are required, skip to step 3. If the DBNG-UP requires an update:
 - a. The DBNG-CP initiates a subscriber session modification request to the DBNG-UP. The DBNG-UP responds to the session modification request if it was successful or not.
3. The DBNG-CP determines if the change has successfully taken place and responds accordingly to the policy server.

For a given subscriber session, the DBNG-CP may initiate multiple, successive modification requests to modify or add configuration for the session. As a modification request or the response to the modification request may be lost, the SCi protocol should retry a request to ensure the DBNG-CP receives a response from the DBNG-UP, confirming the modification request has been processed. This behavior, however, may result in session modification requests being applied to the DBNG-UP in a different order than they were initiated. When the information elements being modified do not overlap in the modification requests, this is considered a safe practice. If, however, one or more of the same information elements is being modified in

successive session modification requests, the DBNG-UP may not be provisioned with the desired value due to mis-ordering. For this case, the DBNG-CP should apply the session modification requests in strict order to ensure the information element value on the UP is deterministically and correctly configured.

When modifying one or more of the same information elements across successive session modification requests for a given subscriber session, the DBNG-CP should apply the update in strict order, meaning a subsequent session modification request is only initiated after the session modification response is received for the prior modification request.

4.5.34 DHCPv4 Release

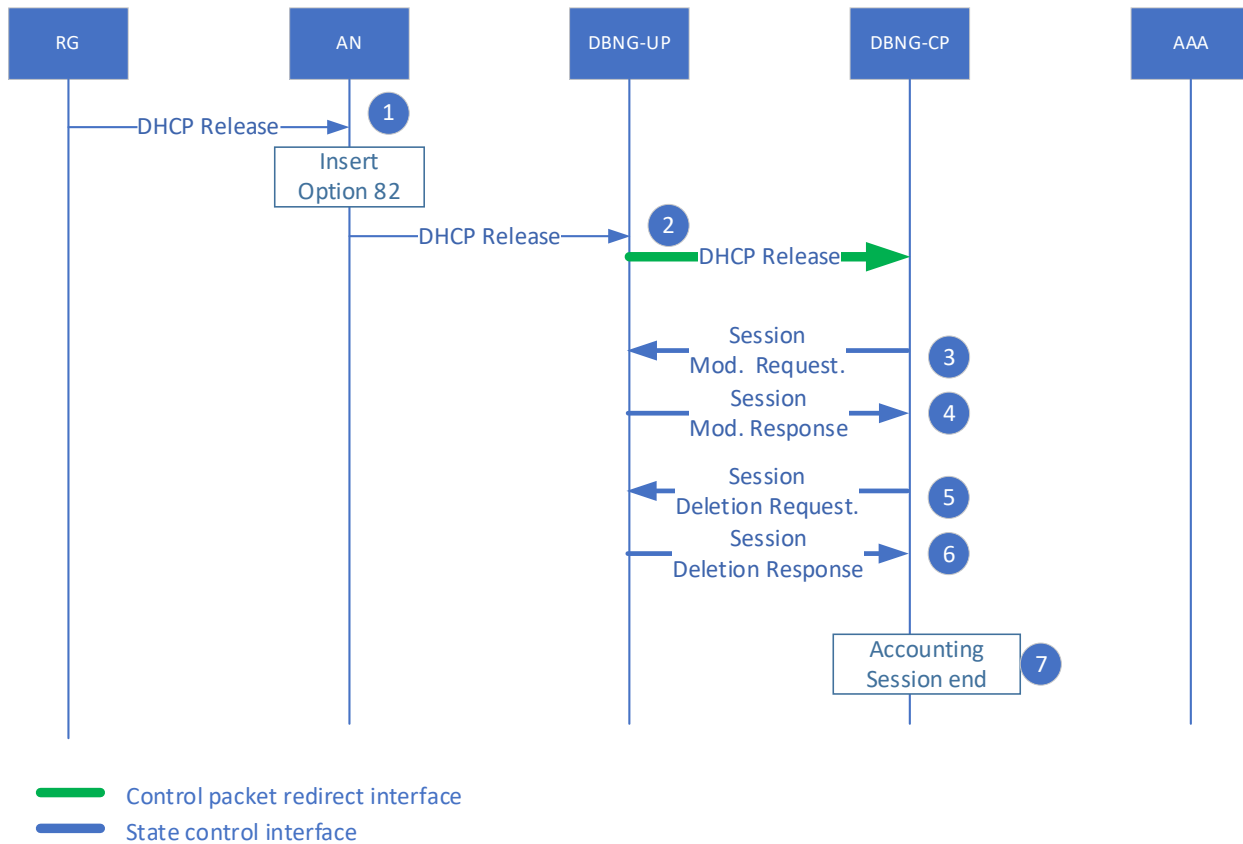


Figure 60: DHCPv4 Release Call flow

This call flow may be applicable to the following subscriber (single or dual stack) releasing a DHCPv4 address:

- IpoE subscriber
- Public Wi-Fi subscriber
- HAG subscriber

Steps:

1. RG initiates a DHCP release message to DBNG, the AN may optionally insert option 82 for circuit information,
2. DBNG-UP receives the DHCP release message and tunnels the message through the CPRi interface to DBNG-CP,
3. [conditional] If the subscriber has multiple stacks (otherwise skip to step5), the DBNG-CP initiates a session modification message to remove all traffic forwarding rules for this particular DHCPv4 session.
4. The DBNG-UP responds to DBNG-CP acknowledging that it has received the session modification message and has removed the traffic forwarding rules for this particular DHCPv4 session. The response message will also include accounting stats. Skip to step 5.
5. [conditional] If this is the subscribe only session, the DBNG-CP initiates a session delete request message to remove all traffic forwarding rules for the subscriber.
6. The DBNG-UP responds to DBNG-CP with a session delete response message after removing the traffic forwarding rules of the subscriber.
7. [optional] The DBNG-CP reports the collected subscriber statistics.

4.5.35 DHCPv4 Relay Release

The DHCPv4 server can either be connected to the DBNG-UP or directly to the DBNG-CP. Note that the call flows apply to both relay and relay-proxy modes

4.5.35.1 DHCPv4 Relay Release (via DBNG-UP)

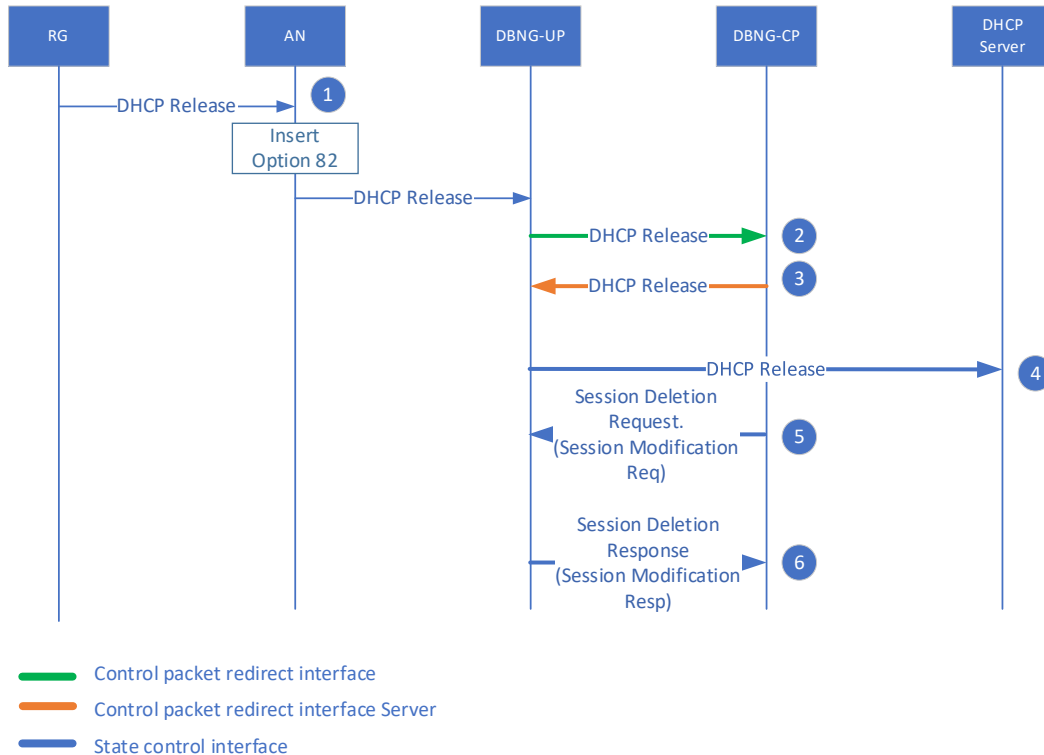


Figure 61: IpoE DHCPv4 Relay Release Call Flow (via DBNG-UP)

1. RG initiates a DHCP Release message, such that the AN inserts option 82 for circuit information.
2. The DHCP Release is sent from the RG to the DBNG-CP through the DBNG-UP utilizing the session control packet redirect tunnel.
3. The DBNG-CP sends the DHCP Release to the DBNG-UP utilizing the server control packet redirect tunnel.
4. The DBNG-UP forwards the DHCP Release to the external DHCP server.
5. For a single or the final DHCP session, the DBNG-CP initiates a Session Deletion Request message to remove all traffic forwarding rules for the subscriber. Otherwise, a Session Modification Request is initiated to remove all traffic forwarding rules for this DHCP session.
6. The DBNG-UP sends a Session Deletion Response (or Session Modification Response if not the final DHCP session)

4.5.35.2 DHCPv4 Relay Release (via DBNG-CP)

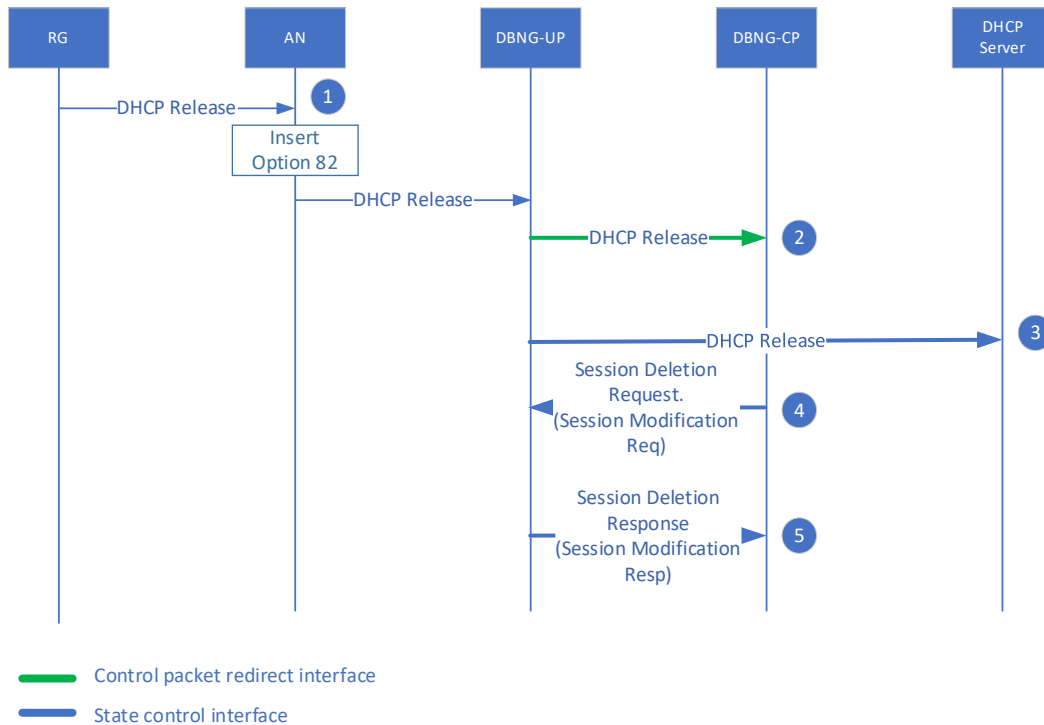


Figure 62: IpoE DHCPv4 Relay Release Call Flow (via DBNG-CP)

1. RG initiates a DHCP Release message, such that the AN inserts option 82 for circuit information.
2. The DHCP Release is sent from the RG to the DBNG-CP through the DBNG-UP utilizing the session control packet redirect tunnel.
3. The DBNG-CP sends the DHCP Release to the DHCP Server via a local control/management interface.
4. For a single or the final DHCP session, the DBNG-CP initiates a Session Deletion Request message to remove all traffic forwarding rules for the subscriber. Otherwise, a Session Modification Request is initiated to remove all traffic forwarding rules for this DHCP session.
5. The DBNG-UP sends a Session Deletion Response (or Session Modification Response if not the final DHCP session) to the DBNG-CP, acknowledging it has removed the traffic forwarding rules for the subscriber.

4.5.36 DHCPv6 Release

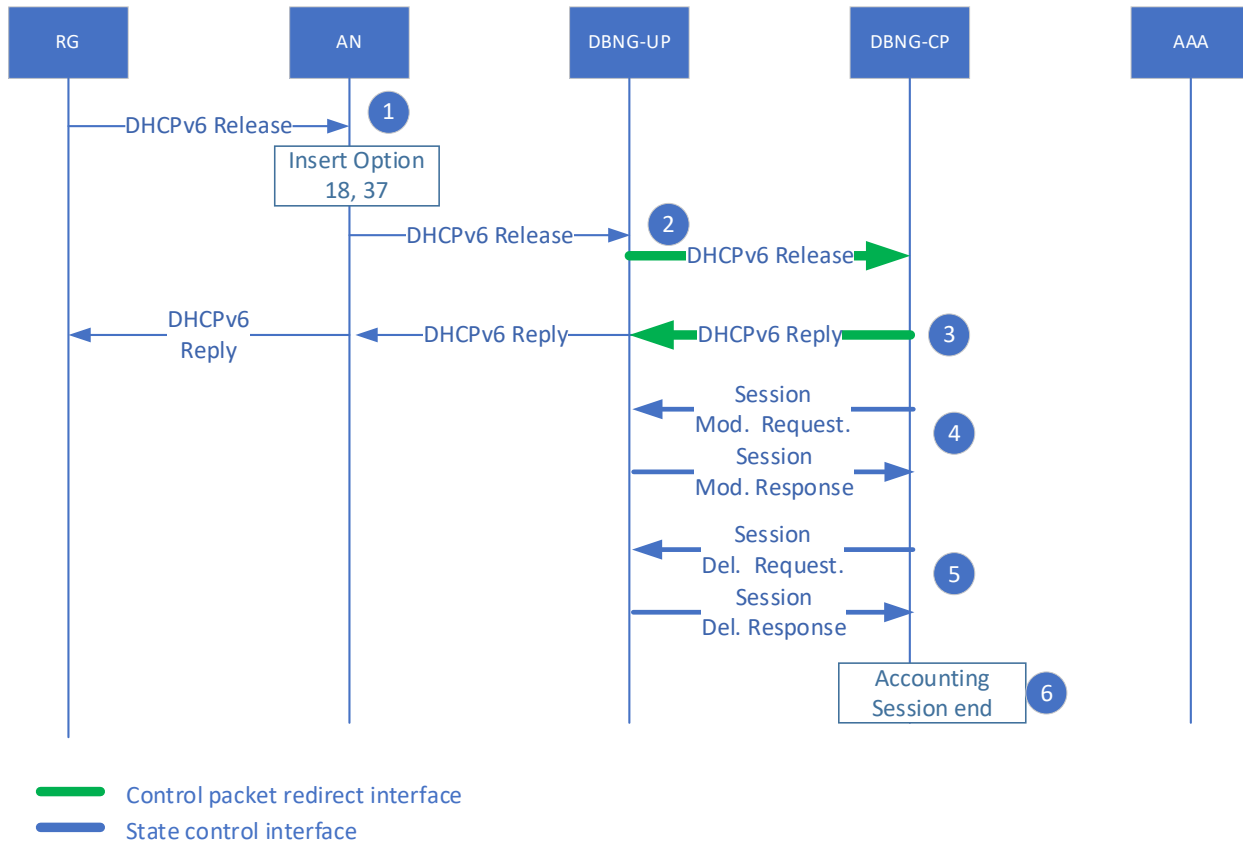


Figure 63: DHCPv6 Release Call Flow

This call flow may be applicable to the following subscriber (single or dual stack) releasing a DHCPv6 address and/or prefix:

- IpoE subscriber
- Public Wi-Fi subscriber
- HAG subscriber

Steps:

1. RG initiates DHCPv6 release message to DBNG, the AN may optionally insert option 82 for circuit information,
2. DBNG-UP receives the DHCPv6 release message within the DHCPv6 relay-forward message and tunnels the message through the CPRi interface to DBNG-CP,
3. DBNG-UP sends a reply back to the RG through CPRi interface.
4. [conditional] If the subscriber has multiple IpoE sessions (otherwise skip to step 5), the DBNG-CP initiates session modification message to remove all traffic forwarding rules for this particular DHCPv6 session.
The DBNG-UP responds to DBNG-CP acknowledging that it has received the session modification message and has removed the traffic forwarding rules for this particular DHCPv6 session. The response message will also include accounting stats. Skip to step 6.
5. [conditional] If this is the subscribe only session, the DBNG-CP initiates session delete request message to remove all traffic forwarding rules for the subscriber.
The DBNG-UP responds to DBNG-CP with a session delete response message after removing the traffic forwarding rules of the subscriber.
6. [optional] The DBNG-CP reports the collected subscriber statistics.

4.5.37 DHCPv6 Relay Release

The DHCPv6 relay server can either be connected to the DBNG-UP via the A10 interface or directly to the DBNG-CP via the B interface.

4.5.37.1 DHCPv6 Relay Release (via DBNG-UP)

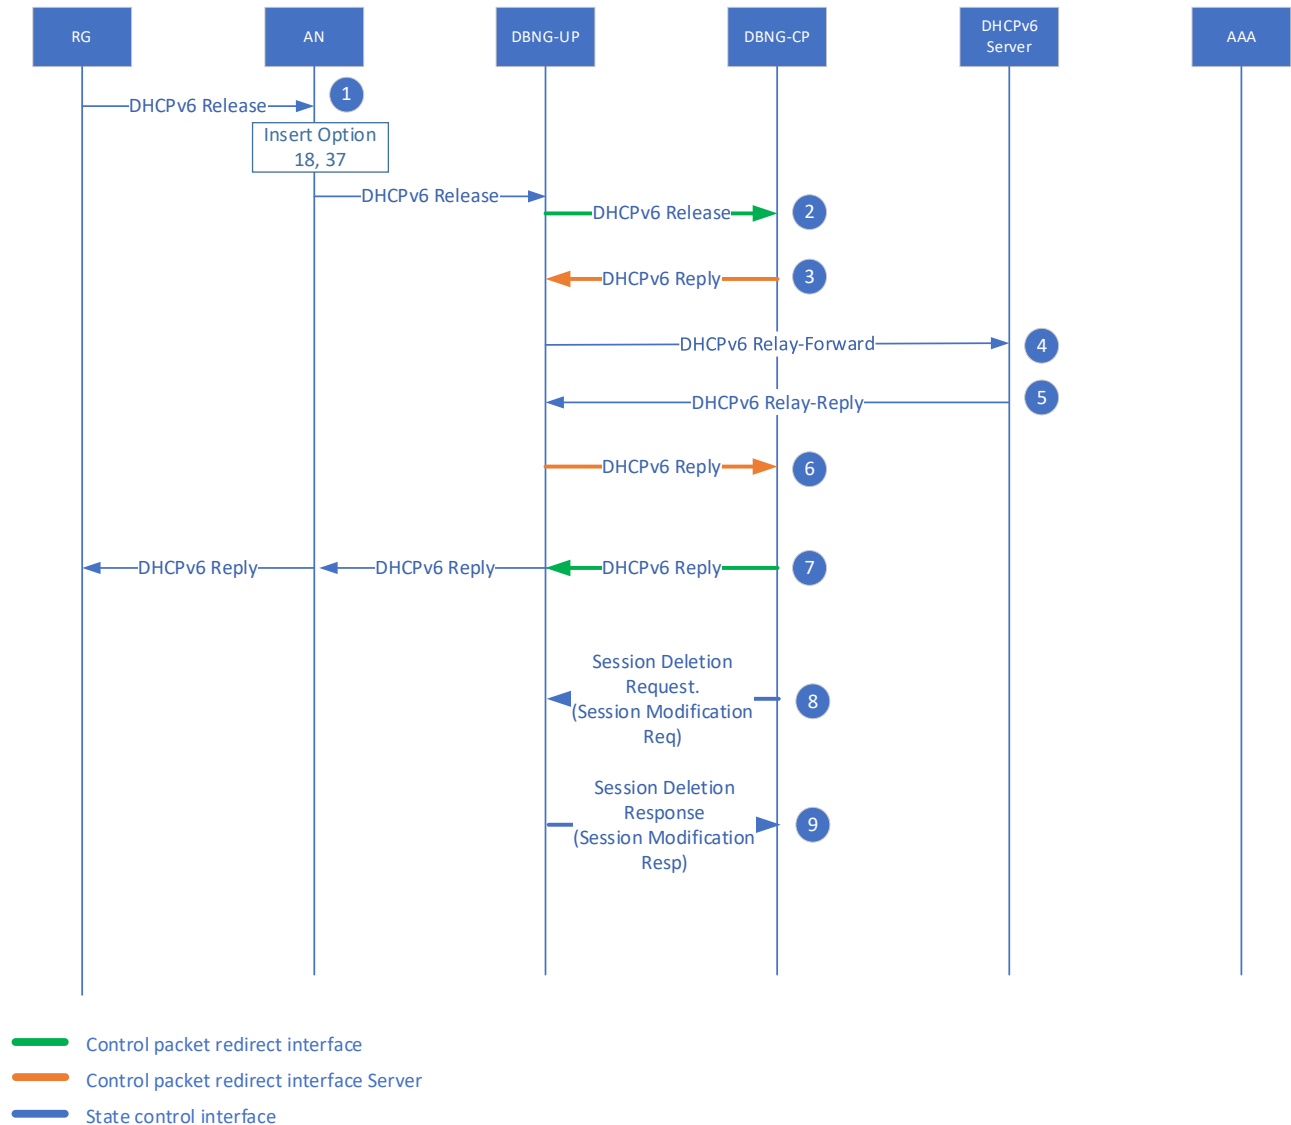


Figure 64: IpoE DHCPv6 Relay Release Call Flow (via DBNG-UP)

1. RG initiates a DHCPv6 Release message, such that the AN inserts option 18 and 37 for circuit information.
2. The DHCPv6 Release is sent from the RG to the DBNG-CP through the DBNG-UP utilizing the session control packet redirect tunnel.

3. The DBNG-CP encapsulates the DHCPv6 Release with a RELAY-FORW header and sends the DHCPv6 Release to the DBNG-UP utilizing the server control packet redirect tunnel.
4. The DBNG-UP forwards the DHCP Release to the external DHCPv6 server.
5. The external DHCPv6 server responds with a DHCPv6 RELAY-REPLY packet that is directed to the DBNG DHCPv6 relay.
6. The DBNG-UP detects the downstream DHCPv6 RELAY-REPLY packet from the external server and redirects the packet to the DBNG-CP through the server Control Packet Redirect tunnel.
7. The DBNG-CP decapsulates the RELAY-REPLY header from the DHCPv6 Reply message before sending it to the RG through the DBNG-UP, utilizing the session control packet redirect tunnel.
8. For a single or the final DHCPv6 session, the DBNG-CP initiates a Session Deletion Request message to remove all traffic forwarding rules for the subscriber. Otherwise, a Session Modification Request is initiated to remove all traffic forwarding rules for this DHCPv6 session.
9. The DBNG-UP sends a Session Deletion Response (or Session Modification Response if not the final DHCPv6 session) to the DBNG-CP, acknowledging it has removed the traffic forwarding rules for the subscriber.

4.5.37.2 DHCPv6 Relay Release (via DBNG-CP)

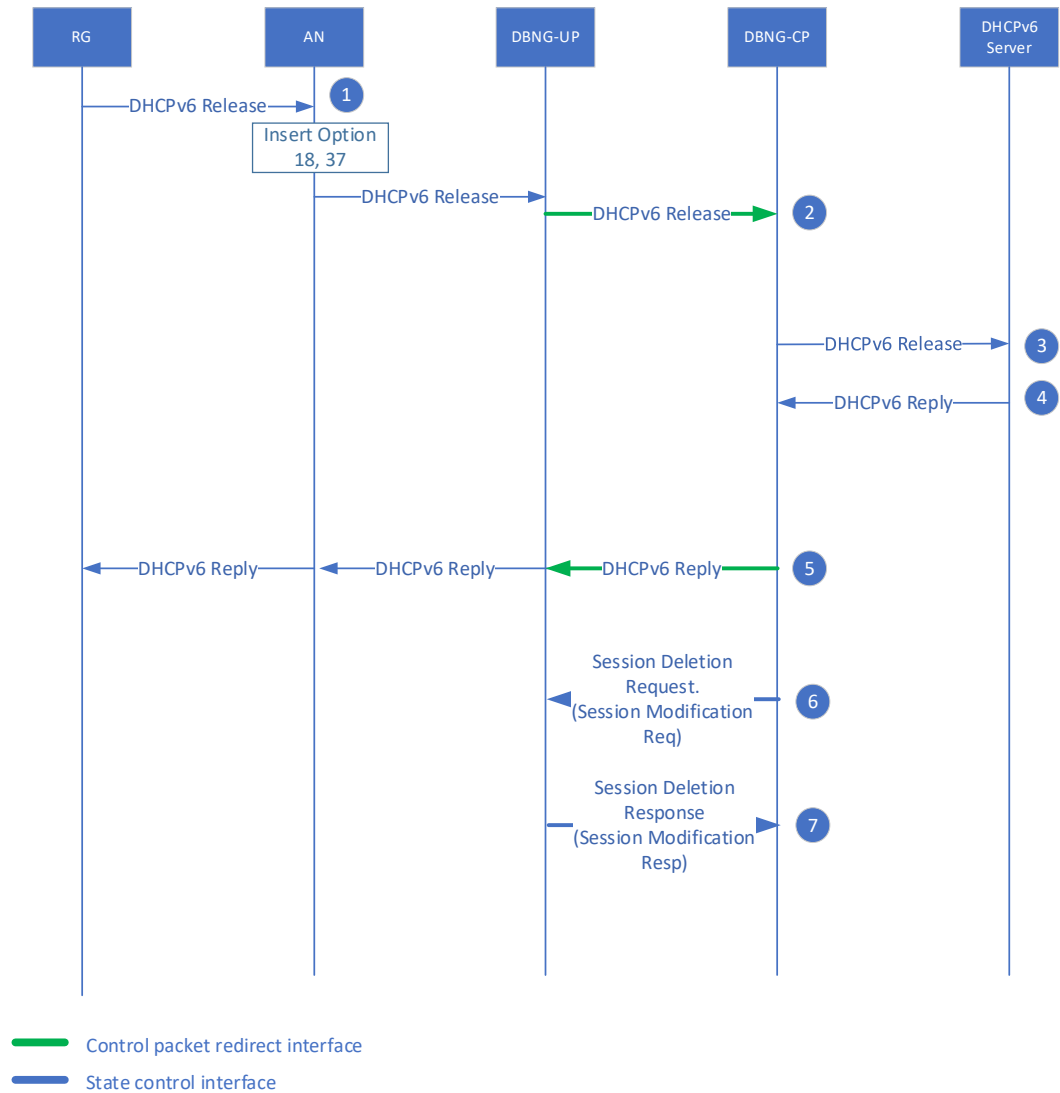


Figure 65: IpoE DHCPv6 Relay Release Call Flow (via DBNG-CP)

1. RG initiates a DHCPv6 Release message, such that the AN inserts option 18 and 37 for circuit information.
2. The DHCPv6 Release is sent from the RG to the DBNG-CP through the DBNG-UP utilizing the session control packet redirect tunnel.
3. The DBNG-CP encapsulates the DHCPv6 Release with a RELAY-FORW header and sends the DHCPv6 Release to the DHCPv6 Server via a local control/management interface.
4. The external DHCPv6 server responds with a DHCPv6 RELAY-REPLY packet that is directed to the DBNG DHCPv6 relay.
5. The DBNG-CP decapsulates the RELAY-REPLY header from the DHCPv6 Reply message before sending it to the RG through the DBNG-UP, utilizing the session control packet redirect tunnel.
6. For a single or the final DHCPv6 session, the DBNG-CP initiates a Session Deletion Request message to remove all traffic forwarding rules for the subscriber. Otherwise, a Session Modification Request is initiated to remove all traffic forwarding rules for this DHCPv6 session.

- The DBNG-UP sends a Session Deletion Response (or Session Modification Response if not the final DHCPv6 session) to the DBNG-CP, acknowledging it has removed the traffic forwarding rules for the subscriber.

4.5.38 Lease and Lifetime timeouts

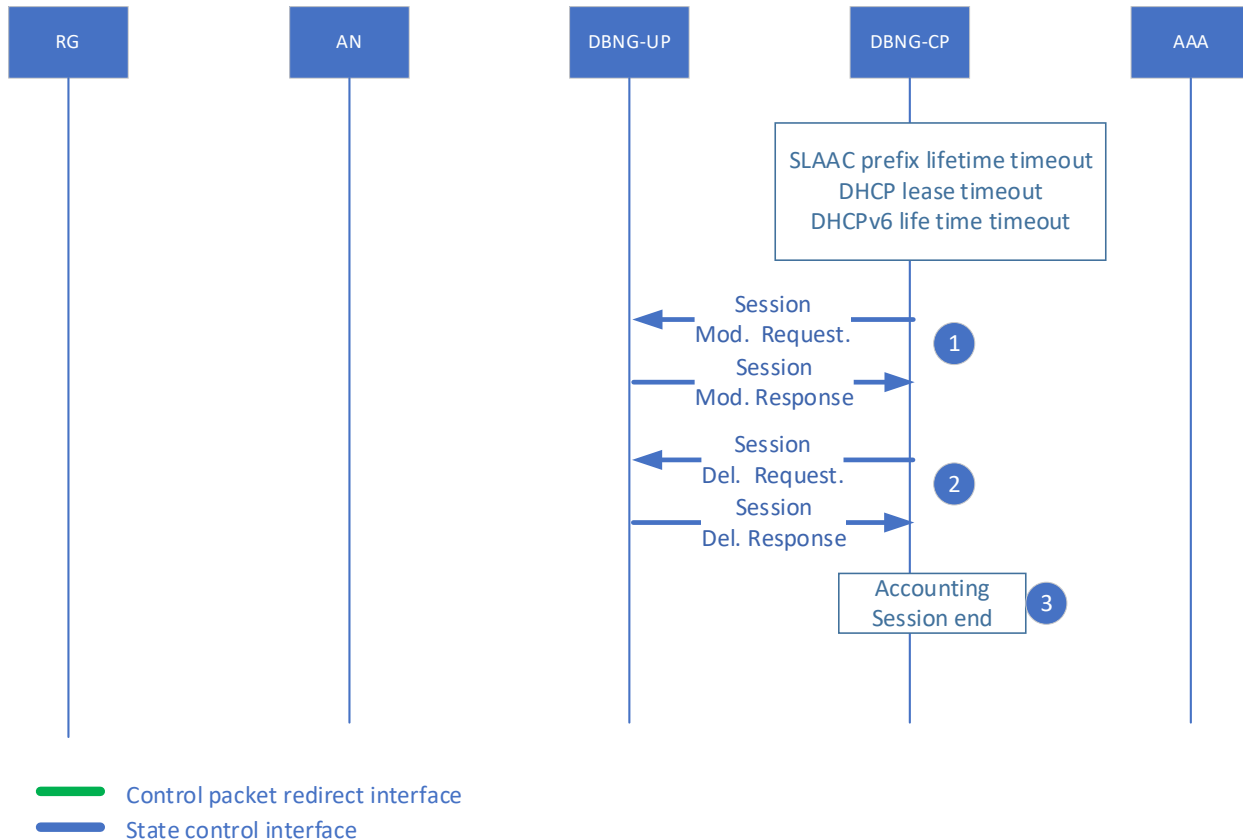


Figure 66: Lease and Lifetime timeout

Steps:

- [conditional] Only if the subscriber has multiple stacks (otherwise skip to step 2), as the lifetime or lease time of the address/prefix assigned to the RG times out, the DBNG-CP initiates session modification message to remove the traffic forwarding rules for the timed-out session.

The DBNG-UP responds to DBNG-CP acknowledging to that it has received the session modification message and has removed the traffic forwarding rules for the timed-out host. Skip to step 3.

- [conditional] If this is the subscriber’s only session, the DBNG-CP initiates a session delete request message to remove all traffic forwarding rules for the subscriber.

The DBNG-UP responds to DBNG-CP with a session delete response message after removing the traffic forwarding rules the subscriber.

- [optional] The DBNG-CP reports the collected subscriber statistics.

4.5.39 PPPoE termination due to LCP echo timeout

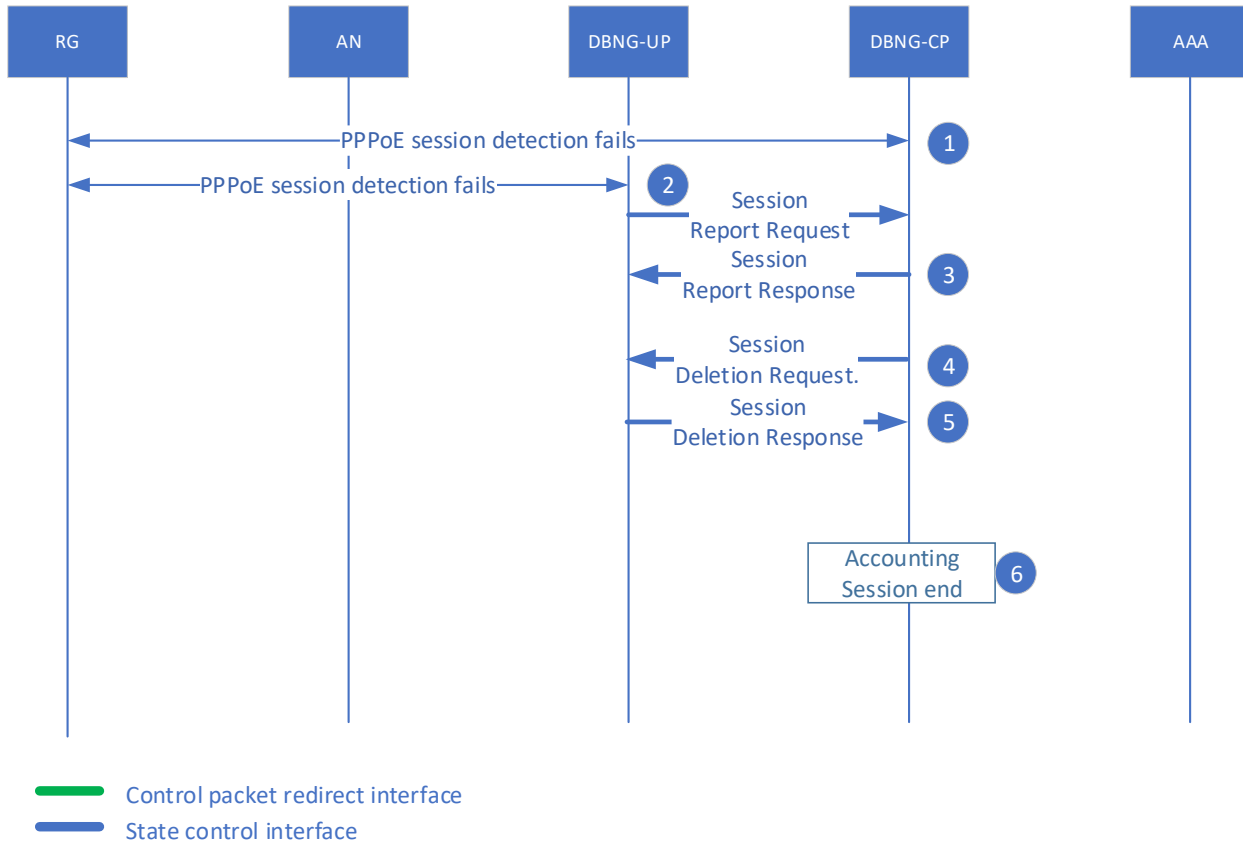


Figure 67: PPPoE termination due to LCP echo timeout

This call flow is applicable to both single stack and dual stack subscribers for the following use case:

- PPPoE subscriber
- HAG subscriber

Steps:

1. [conditional] If the LCP echo is offloaded to the DBNG-UP, then skip to step 2. If not, the DBNG-CP detects LCP echo failure and begins session deletion as explained in step 4.
2. [conditional] The LCP echo failure is detected by the DBNG-UP and the DBNG-UP initiates a session report request to the DBNG-CP to inform of a lost connection with the subscriber RG.
3. [conditional] The CP replies and acknowledges it received the session report
4. The DBNG-CP terminates the subscriber session by initiating a session deletion request message to remove all traffic forwarding rules for the subscriber.
5. The DBNG-UP responds to DBNG-CP with a session deletion response message after removing the traffic forwarding rules.
6. [optional] The DBNG-CP reports the collected subscriber statistics.

4.5.40 PPPoE Client initiated PPP LCP termination

This call flow applies directly to LCP termination but equally applies in the case where the RG sends a PADT to terminate the session.

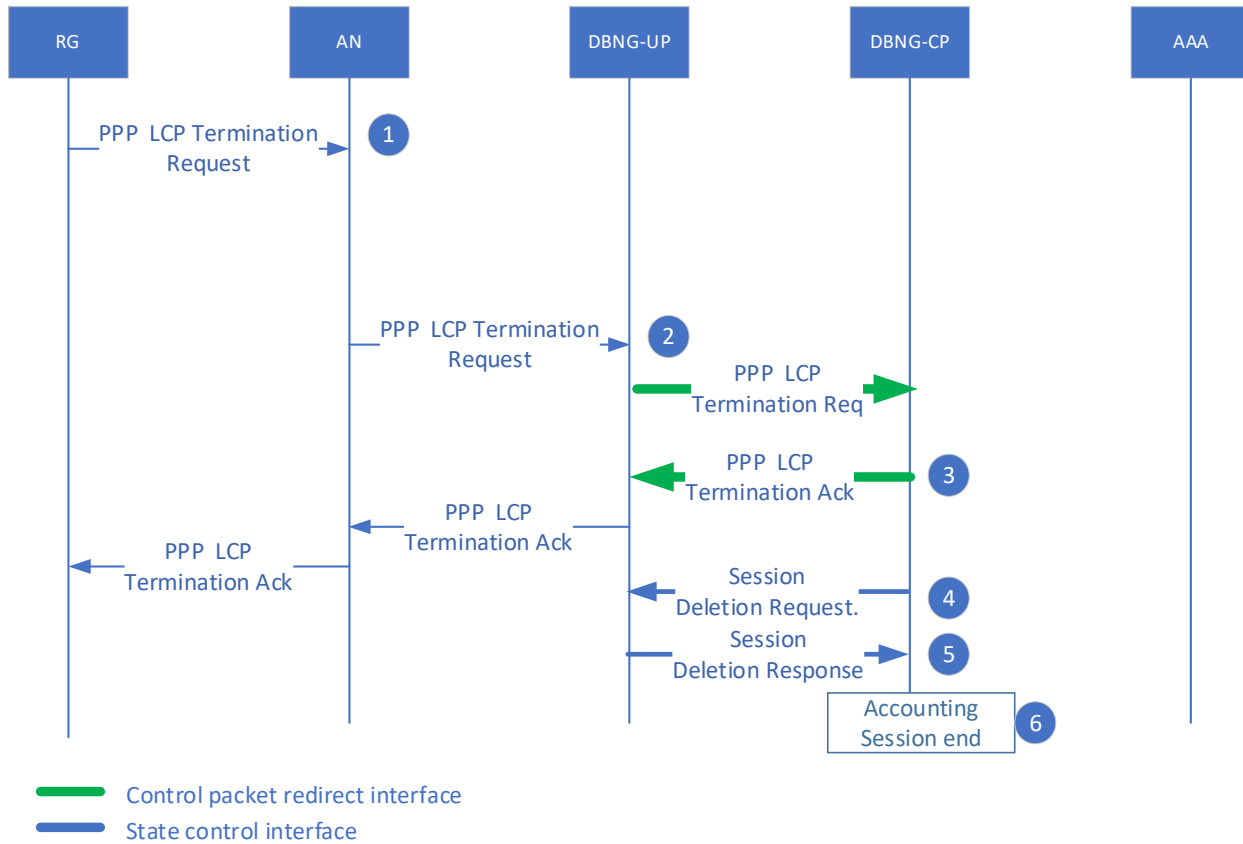


Figure 68: PPPoE client initiated PPP LCP termination

Steps:

1. The PPPoE client initiates a session termination.
2. DBNG-UP receives the LCP termination request message and tunnels the message through the CPRi interface to DBNG-CP
3. The DBNG-CP replies to the termination message request with an LCP termination acknowledgement message.
4. The DBNG-CP terminates the subscriber PPPoE session by initiating a session deletion request message to remove all traffic forwarding rules for the subscriber.
5. The DBNG-UP responds to the DBNG-CP with a session deletion response message after removing the traffic forwarding rules.
6. [optional] The DBNG-CP reports the collected subscriber statistics.

Note: PPPoE PADT may be sent by the RG on receipt of LCP Termination Ack and/or PPPoE PADT may be sent by DBNG-CP on transmitting LCP Termination Ack (Not shown in figure)

4.5.41 PPPoE Server initiated PPP LCP termination

This call flow applies directly to LCP termination but equally applies in the case where the DBNG-CP sends a PADT to terminate the session.

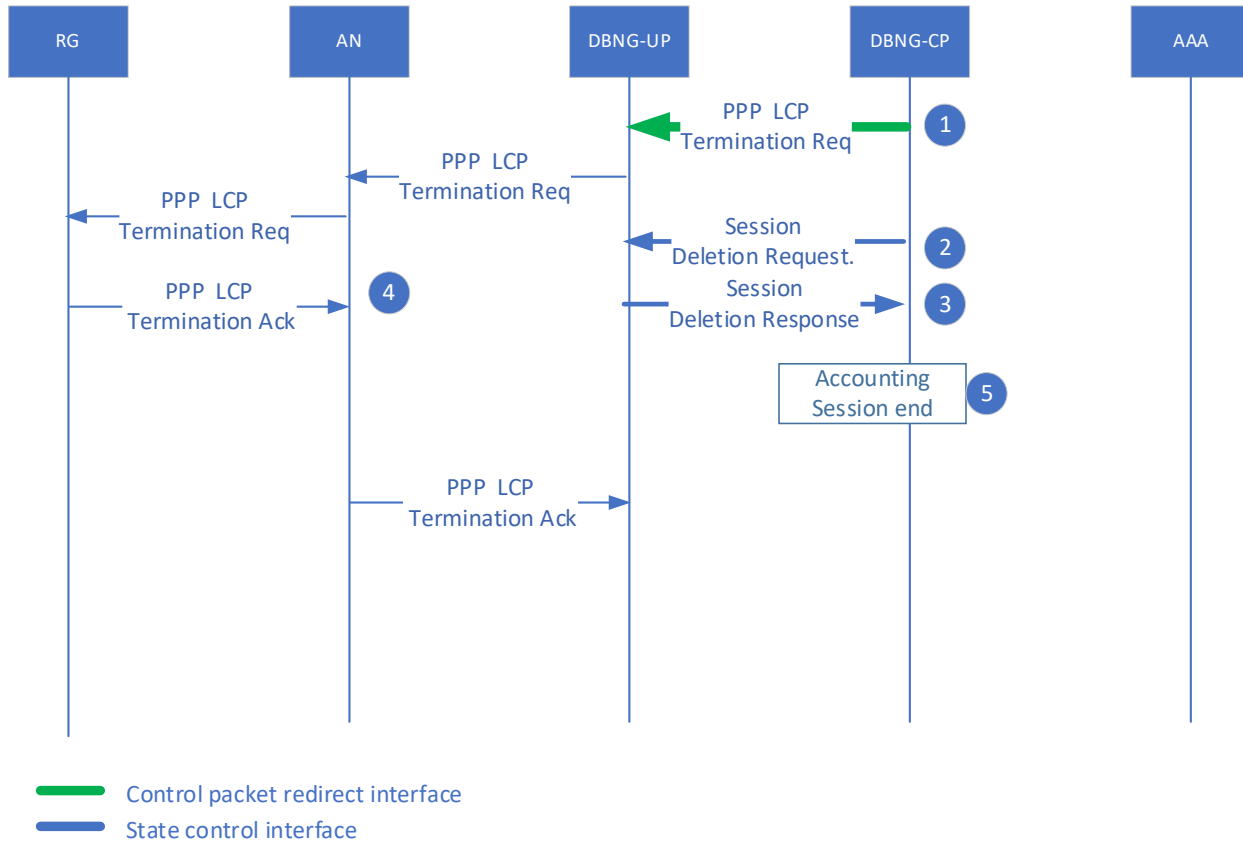


Figure 69: PPPoE Server initiated PPP LCP termination

Steps:

1. The DBNG-CP initiates a session termination by sending an LCP termination request. Message to the RG. This could be due to various network events such as LCP echo keepalive failure or exhaustion of credit. The message is tunneled to the DBNG-UP over the CPRi interface.
2. The DBNG-CP initiates session deletion request message to remove all traffic forwarding rules for the subscriber.
3. The DBNG-UP responds to the DBNG-CP with a session deletion response message after removing the traffic forwarding rules.
4. The RG replies with a PPP LCP Termination acknowledgement and proceeds to remove the PPPoE session.
5. [optional] The DBNG-CP reports the collected subscriber statistics.

Note: PPPoE PADT may be sent by the RG on receipt of LCP Termination Ack and/or PPPoE PADT may be sent by DBNG-CP on transmitting LCP Termination Ack (Not shown in Figure).

4.5.42 RG initiated PPP LCP termination for L2TP session

This call flow applies directly to LCP termination but equally applies in the case where the RG sends a PADT to terminate the session.

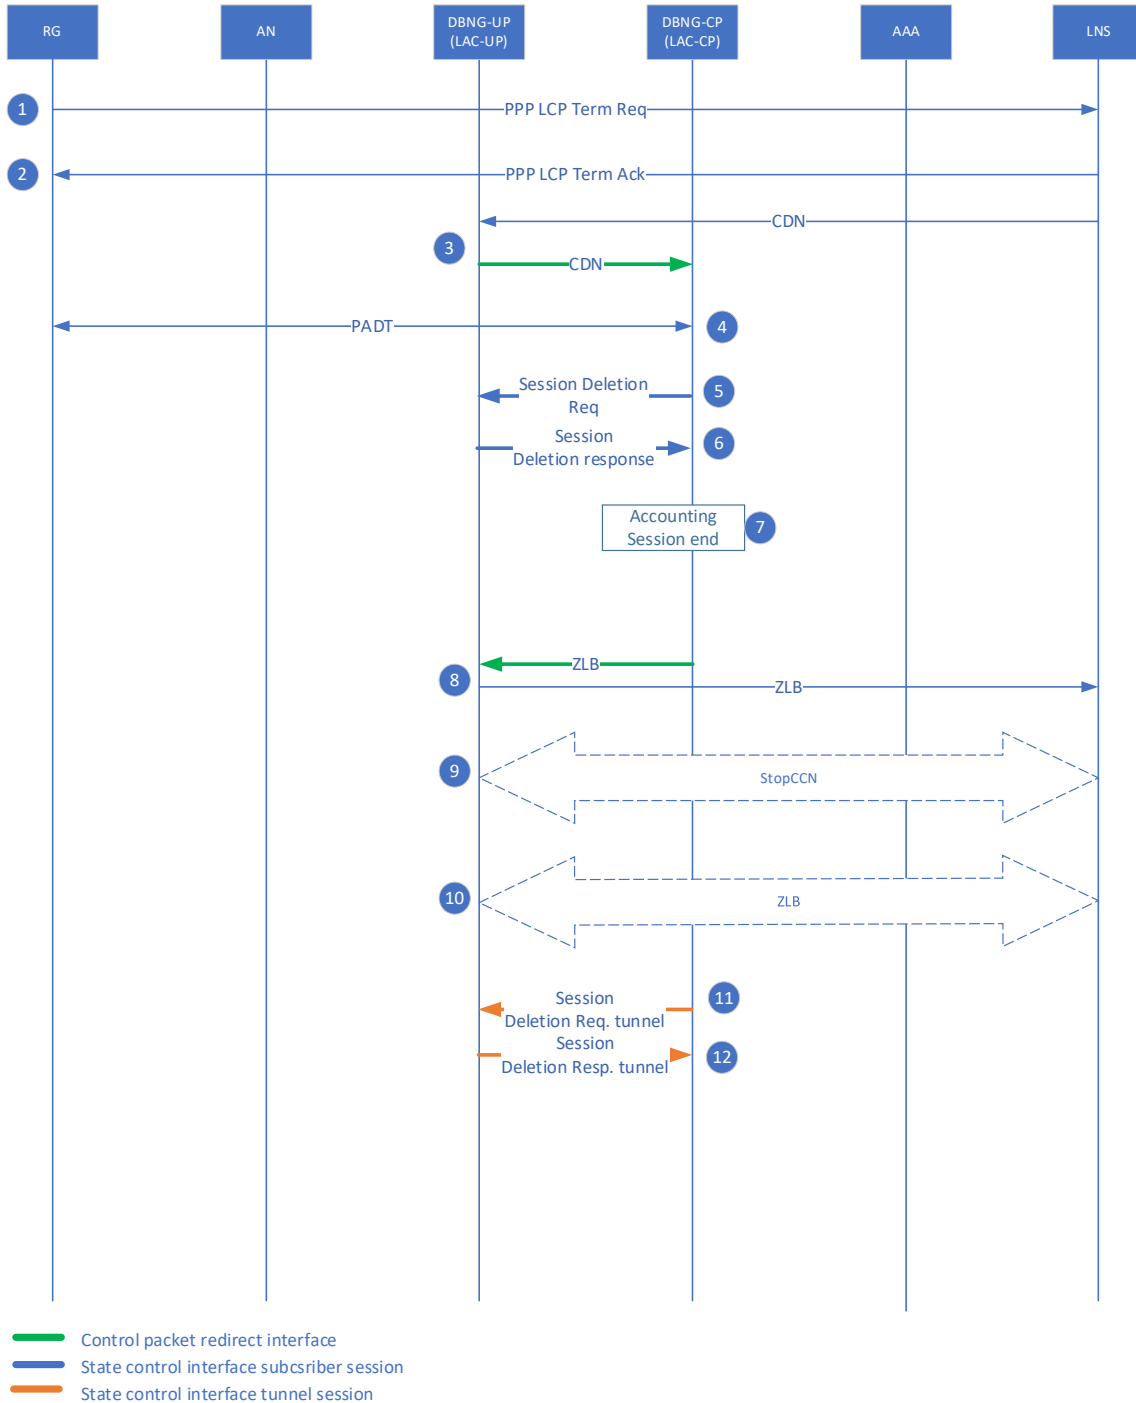


Figure 70: RG initiated PPP LCP termination for L2TP session

Steps:

1. The RG initiates a session termination to the LNS.
2. The LNS will acknowledge the termination request.
3. The LNS sends a CDN message which reaches the DBNG-CP through the CPRi.
4. PPPoE PADT may be sent by the RG on receipt of LCP Termination Ack and/or PPPoE PADT may be sent by LAC-CP on the reception of CDN
5. The DBNG-CP sends a session deletion message to remove all traffic forwarding rules for the subscriber.
6. The DBNG-UP responds to DBNG-CP acknowledging to that it has received the session deletion message and has removed the traffic forwarding rules.
7. [optional] The DBNG-CP reports the collected subscriber statistics.
8. The DBNG-CP acknowledges the CDN message with a ZLB message. The DBNG-UP sends the ZLB message to the LNS.
9. [conditional] If the LAC or LNS determines that there are no more L2TP sessions on the L2TP tunnel, the LAC or the LNS can terminate the tunnel by sending a L2TP stopCCN message. The stopCCN message can either be generated by the DBNG-CP and sent to the LNS via the CPR interface, or the LNS can send the message to the DBNG-CP through the CPR interface.
10. [conditional] The LNS or the LAC acknowledges the termination request by sending a ZLB message.
11. [conditional] The DBNG-CP sends a session deletion message to remove all traffic forwarding rules for the tunnel.
12. [conditional] The DBNG-UP responds to DBNG-CP acknowledging to that it has received the session deletion message and has removed the traffic forwarding rules for the L2TP tunnel.

Note: step 4 is the earliest point where Session Deletion can occur, but it is possible to trigger the session deletion in a later step.

4.5.43 LAC initiated termination for L2TP session

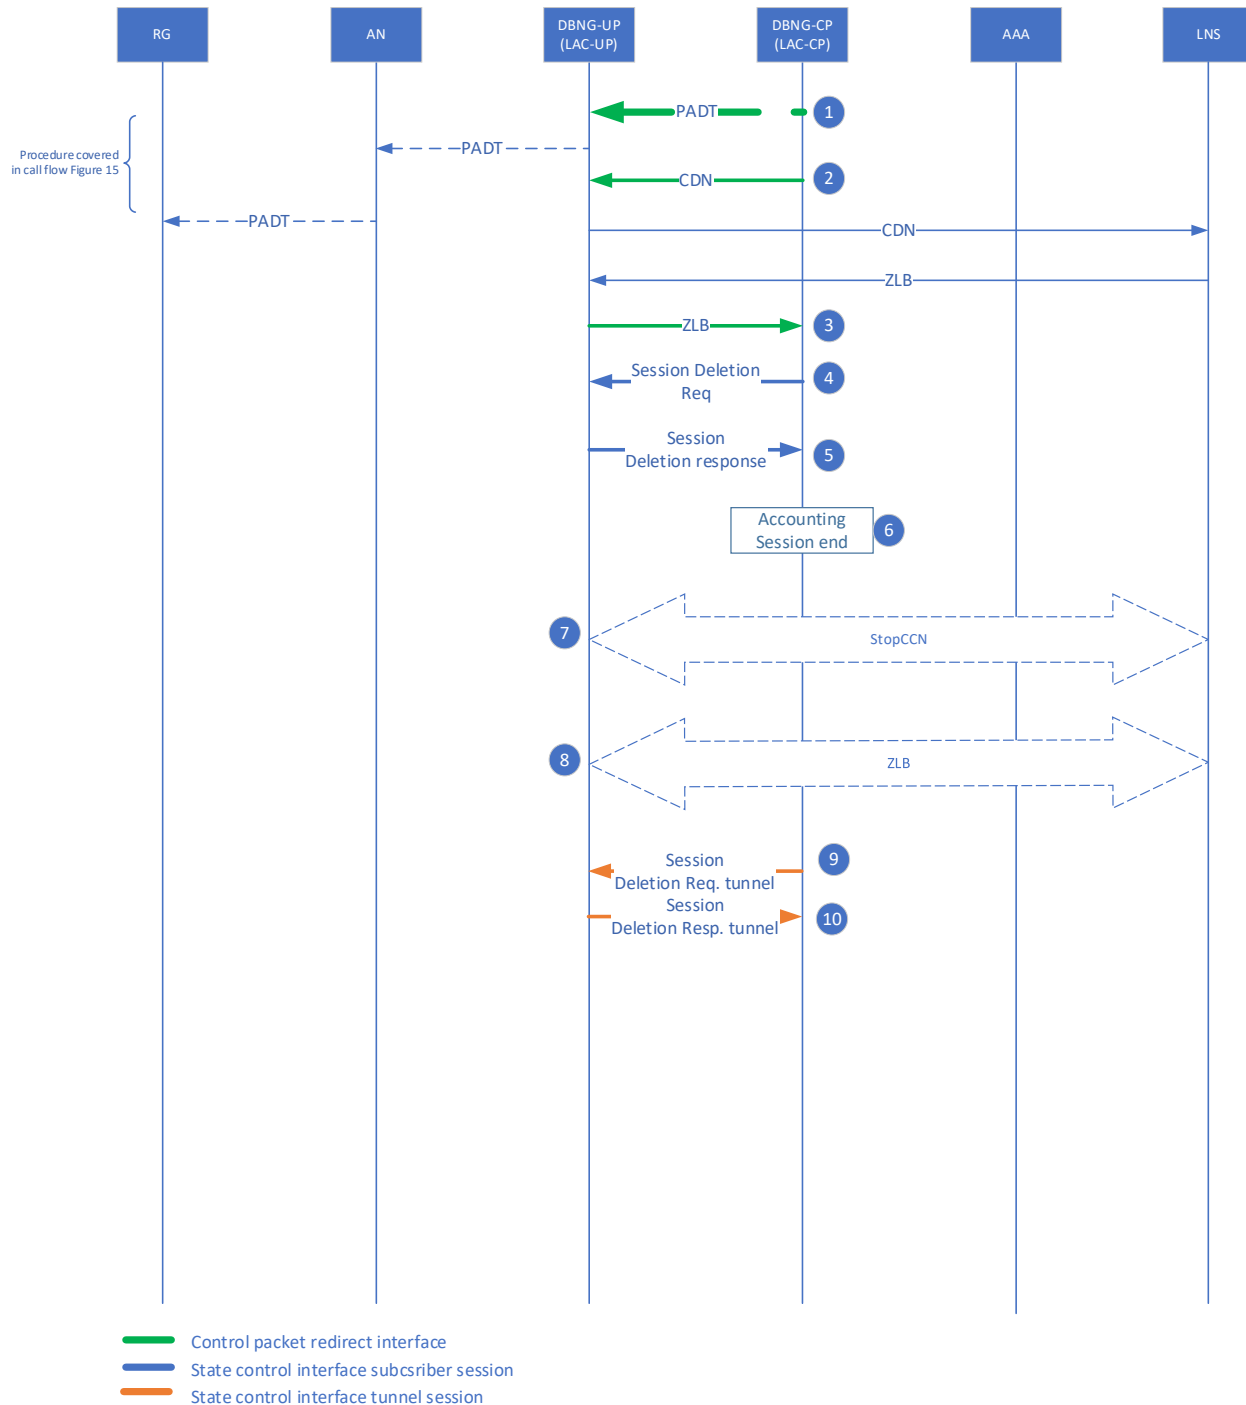


Figure 71: LAC initiated termination for L2TP session

Steps:

1. The LAC initiates a session termination with PADT. The DBNG-UP tunnels the message through the CPRi interface to RG.
2. The DBNG-CP initiates a L2TP session connection disconnect message (L2TP CDN message) towards the LNS.
3. The LNS will acknowledge the request with a ZLB. Once the DBNG-CP receives the ZLB message from the LNS, it acknowledges the PPP LCP termination message from the RG.
4. The DBNG-CP initiates a session deletion message to remove all traffic forwarding rules for the subscriber.
5. The DBNG-UP responds to DBNG-CP acknowledging to that it has received the session deletion message and has removed the traffic forwarding rules.
6. [optional] The DBNG-CP reports the collected subscriber statistics.
7. [conditional] If the LAC or LNS determines that there are no more L2TP sessions on the L2TP tunnel, the LAC or the LNS can terminate the tunnel by sending a L2TP stopCCN message. The stopCCN message can either be generated by the DBNG-CP and sent to the LNS via the CPR interface, or the LNS can send the message to the DBNG-CP through the CPR interface.
8. [conditional] The LNS or the LAC acknowledges the termination request by sending a ZLB message.
9. [conditional] The DBNG-CP sends a session deletion message to remove all traffic forwarding rules for the tunnel.
10. [conditional] The DBNG-UP responds to DBNG-CP acknowledging to that it has received the session deletion message and has removed the traffic forwarding rules for the L2TP tunnel.

Note: step 2 is the earliest point where Session Deletion can occur, but it is possible to trigger the session deletion in a later step.

4.5.44 LAC initiated termination for L2TP session on LNS

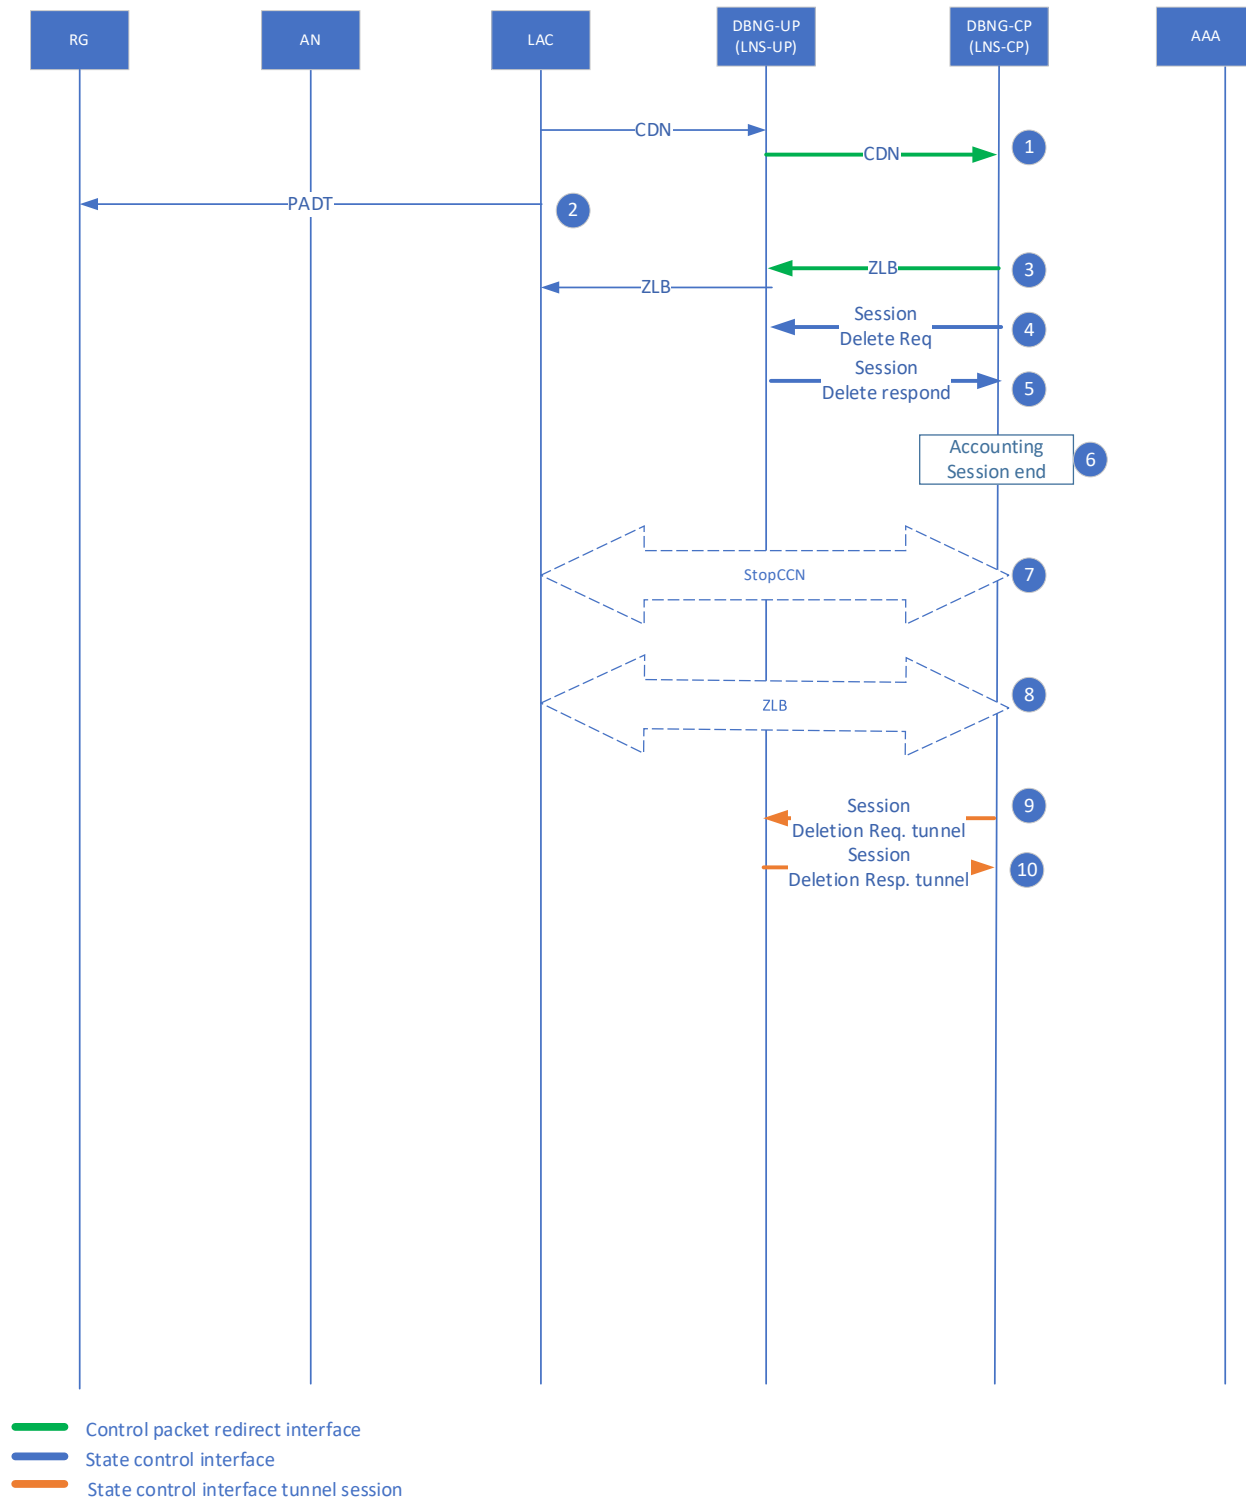


Figure 72: LAC initiated termination for L2TP session on LNS

Steps:

1. The LAC initiates a session termination. The DBNG-UP tunnels the message through the CPR interface to LNS.
2. A PADT may be sent by the LAC to the RG
3. The LNS acknowledges the session termination and sends a ZLB message back to the LAC.
4. The DBNG-CP initiates a session deletion message to remove all traffic forwarding rules for the subscriber.
5. The DBNG-UP responds to DBNG-CP acknowledging that it has received the session deletion message and has removed the traffic forwarding rules.
6. [optional] The DBNG-CP reports the subscriber statistics for billing purposes.
7. [conditional] If the LAC or LNS determines that there are no more L2TP sessions on the L2TP tunnel, the LAC or the LNS can terminate the tunnel by sending a L2TP stopCCN message. The stopCCN message can either be generated by the DBNG-CP and sent to the LNS via the CPR interface. Or the LNS would send message to the DBNG-CP through the CPR interface.
8. [conditional] The LNS or the LAC will then acknowledge the termination request by sending a ZLB message.
9. [conditional] The DBNG-CP initiates session deletion message to remove all traffic forwarding rules for the tunnel.
10. [conditional] The DBNG-UP responds to DBNG-CP acknowledging that it has received the session deletion message and has removed the traffic forwarding rules for the L2TP tunnel.

4.5.45 DBNG-CP initiated termination

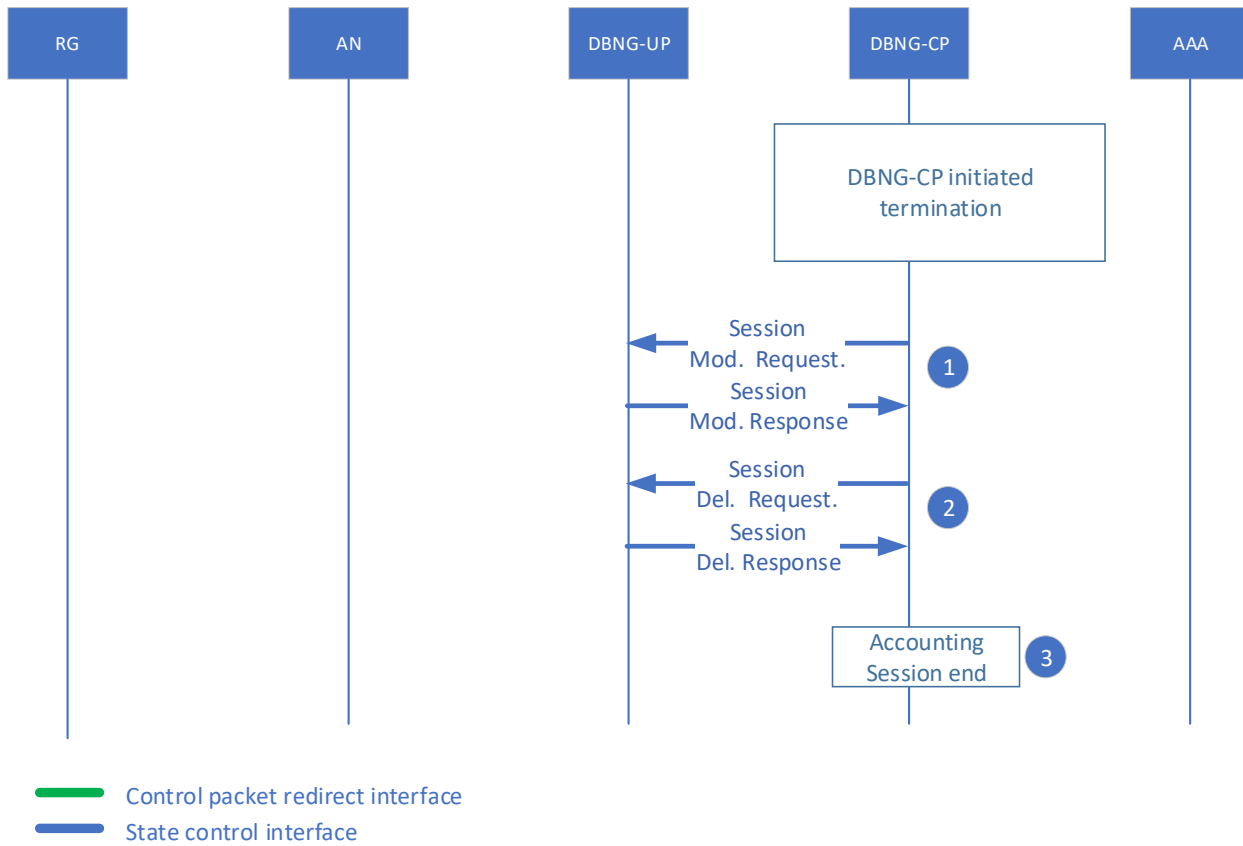


Figure 73: DBNG-CP initiated termination

Steps:

1. [conditional to subscriber with multiple stacks] The DBNG-CP, due to various error scenarios, terminates one of the subscriber session stacks, by initiating a session modification message. The session modification is used to remove the forwarding rules related to the terminated stack. The DBNG-UP responds to DBNG-CP acknowledging the session modification message and removes the traffic forwarding rules for the particular stack. The response message may include accounting stats if required.
2. [conditional to subscriber with single remaining stack] The subscriber only has one remaining stack, the DBNG-CP, due to various error scenarios, terminates the subscriber session by initiating a session delete message to remove all traffic forwarding rules for the subscriber. The DBNG-UP responds to DBNG-CP acknowledging the session delete message and removes the traffic forwarding rules of the subscriber.
3. [optional if accounting is required] The DBNG-CP reports the subscriber statistics for billing purposes.

Note: Terminating the Ipv4 or Ipv6 subscriber may or may not terminate an entire subscriber session, this is up to the vendor implementation.

4.5.46 DBNG subscriber statistics

There are two use cases for DBNG statistics

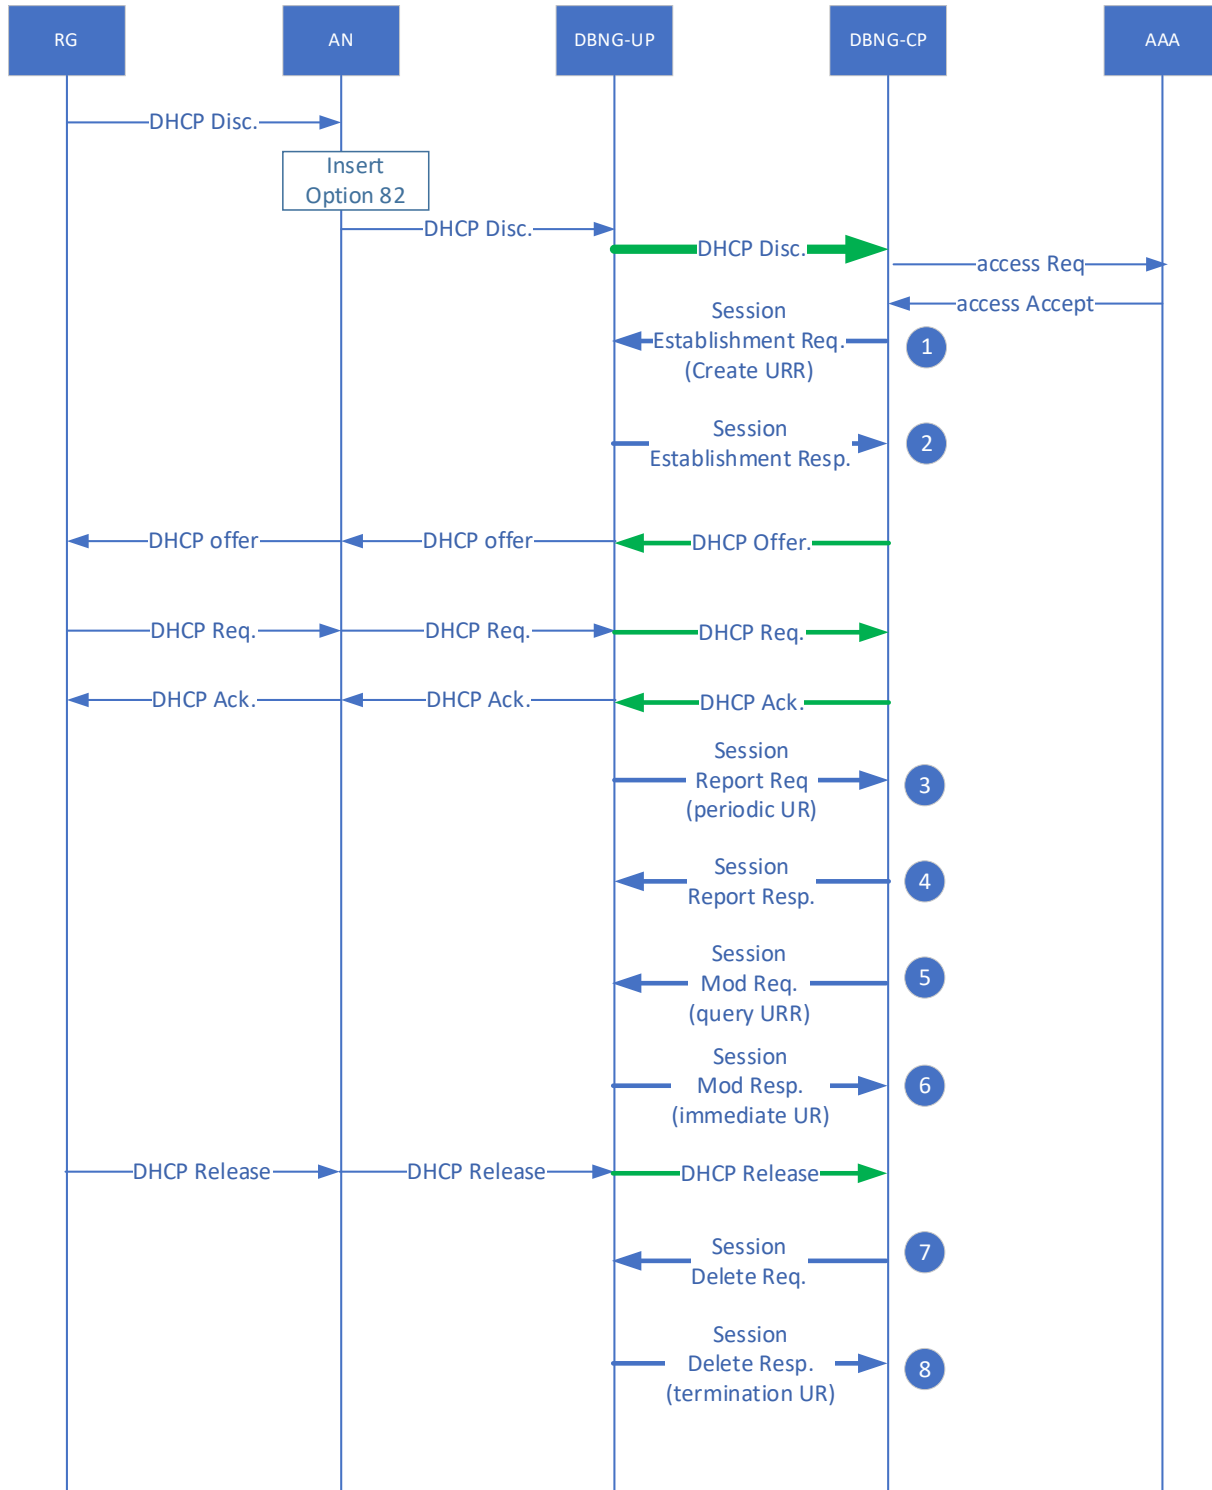
1. For accounting purposes. Basic Broadband Service accounting includes subscriber session start and subscriber session stop. The accounting starts after the subscriber session connects to the DBNG-UP before any packets are sent and does not require the DBNG-UP to report any statistics to the DBNG-CP. The accounting stops when the subscriber session disconnects from the DBNG-UP, which may be triggered manually by the subscriber or by the DBNG. In both cases, the DBNG-CP will end the PFCP session with the DBNG-UP. Upon receiving the session termination, the DBNG-UP will respond to the DBNG-CP acknowledging the termination with the final statistics included.

In some cases, interim reports are sent after the accounting starts and until the accounting stops. The interim is generally sent periodically to an accounting server. For interim, the DBNG-UP is required to report the instantaneous statistics for a particular subscriber session. The trigger for the report can be accomplished in two ways. 1) A polling method, where the DBNG-CP would trigger a query to the DBNG-UP and the DBNG-UP acknowledging the query with the statistics included. 2) A push method where the DBNG-UP is triggered by a predetermined timer to send the statistics to the DBNG-CP and the DBNG-CP sending a response to confirm the receipt.

2. Real time collection of statistics. For operational purposes, the DBNG-CP may request the DBNG-UP for subscriber session statistics, some examples are debugging traffic issues, displaying show command output in a terminal, responding to externally triggered ad-hoc accounting. In such cases, a polling method, where the DBNG-CP queries the DBNG-UP for subscriber session statistics is used.

The above use cases are supported by the standard 3GPP PFCP procedure and IEs. No feature flags are required. It should be noted that the polling method is required by both use cases while the push method can serve as an optional solution.

Below call flow describes subscriber statistics exchange over SCi using IpoE DHCPv4 as an example. The call flow captures the two essential mechanisms for statistics monitoring in TR-459 which can either be pushed or pulled from the DBNG-UP.



— Control packet redirect interface
— State control interface

Figure 74: DBNG subscriber statistics for IpoE

IpoE DHCPv4 subscriber connects and on successful authentication is assigned an Ipv4 address on the DBNG-CP.

1. DBNG-CP sends Usage Reporting Rules and reporting triggers in Session Establishment/ Modification Request message and the DBNG-UP to push the statistics in a defined interval.
2. DBNG-UP sends the Session Establishment/ Modification Response acknowledging URR creation.

During the life of the session, multiple usage reports are generated based on reporting trigger. The DBNG-CP will inform the DBNG-UP to push the statistics in a defined interval.

3. DBNG-UP sends interim Usage Report in Session Report Request message
4. DBNG-CP replies with the Session Report Response.

Note: Steps 3,4 represent the push method and can occur many times during the life of the session.

During the life of the session, the DBNG-CP may request usage report information for a subscriber anytime. The DBNG-CP will pull from the DBNG-UP the instantaneous statistics. This can be done periodically to support interim accounting updates.

5. DBNG-CP sends URR query in Session Modification Request message
6. DBNG-UP replies with Usage Report in the Session Modification Response

Note: Steps 5,6 represent the polling method and can occur many times during the life of the session.

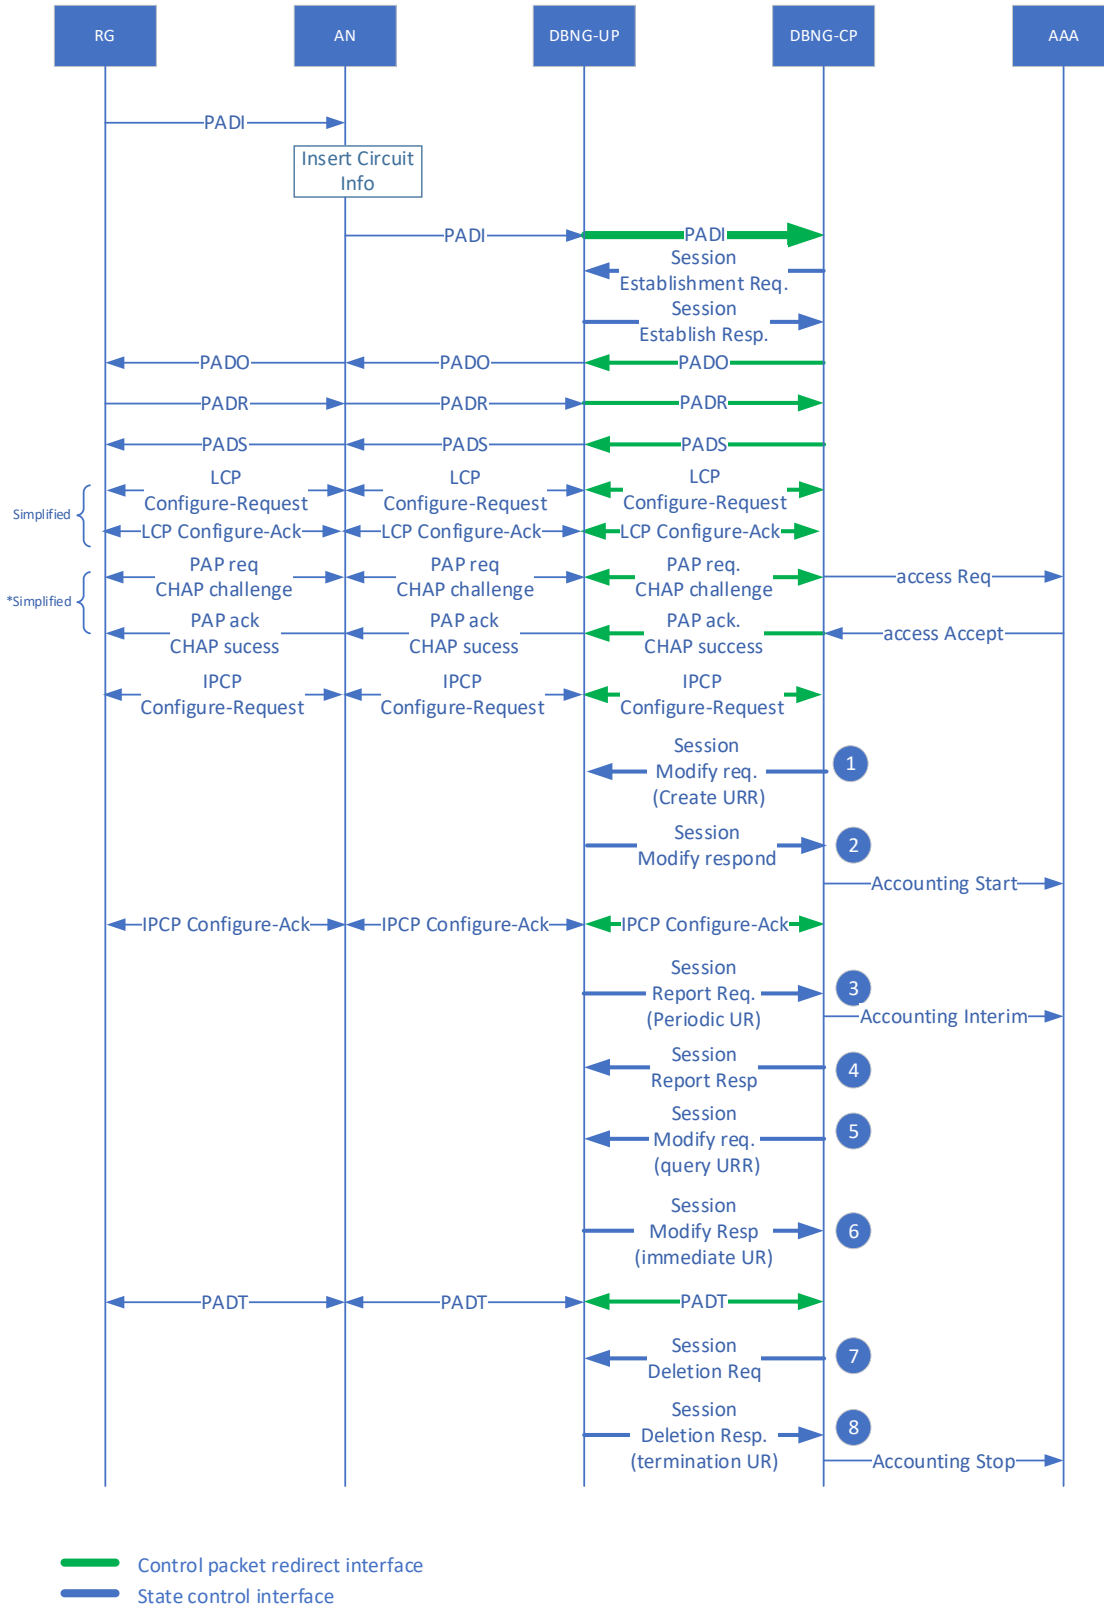
During subscriber logout, the DBNG-UP collects final statistics information as part of session tear down and sends it inline in Session Deletion Response.

7. DBNG-CP sends the Session Deletion Request message
8. DBNG-UP replies with final Usage Report in the Session Deletion Response. Alternatively, the DBNG-UP can choose to send final Usage Reports in separate Session Report Request messages if the DBNG-CP indicates support using the ARDR flag in the CP Function Features flags. This is an optional PFCP feature described in 3GPP TS 29.244 [28], clause 5.2.2.3.1.

DBNG-CP collects the subscriber session final statistics and deletes the session.

The above are native PFCP IE statistics monitoring features as defined in 3GPP TS 29.244 [28] and provide ways subscriber session statistics may be collected. It is up to the DBNG-CP to elaborate the statistics received by the DBNG-UP in order to formulate the correct reporting to the accounting (for example, as defined in RFC 2866 [37]) and operational systems. It is also up to the DBNG-CP to send the reporting at the right time to the external systems such as AAA server.

Below is an additional call flow example with PPPoEv4 (termination report sent inline with Session Delete Response).



*Simplified Call Flow: PAP is unidirectional while CHAP is bidirectional

Figure 75: DBNG Subscriber Statistics for PPPoE

Except for the Usage Reporting Rules (URR) and reporting triggers being sent in Session Modification Request by the DBNG-CP and acknowledged by the DBNG-UP in Session Modification Response, the remaining steps of statistics handling are the same as IpoE DHCPv4 described above.

4.5.47 Resiliency Call Flows

The call flow in section 4.5.47.1 refers to the establishment of a resilient session without the delegation by DBNG-CP to DBNG-UP of tracking the operational state of a logical port. The call flow in section 4.5.47.2 is a continuation of section 4.5.47.1, showing a switchover triggered by the DBNG-CP.

The call flow in section 4.5.47.3 refers to the establishment of resilient session with the use of track-logical-port solution described in section 4.4.7. This solution, which allows the switchover to be independently triggered by the DBNG-UP, requires that the DBNG-CP set the SGRP state to “Track-logical-port” on both the active and the backup DBNG-UP instances, specifying the logical ports whose state the DBNG-Ups have to track. The DBNG-UP will track the operational state of the access network connectivity on those logical ports, to decide whether and when to execute the switchover.

The call flow in section 4.5.47.4 is a continuation of section 4.5.47.3, showing a switchover independently initiated by the DBNG-UP instances.

4.5.47.1 Establishment of a resilient session

The DBNG-CP can establish a resilient session if it controls at least two distinct DBNG-UP instances and if the access network provides a dual-homing active-backup connectivity onto the two DBNG-UP instances. In the following figure, we identify the two DBNG-UP instances under consideration as “DBNG-UP1” and “DBNG-UP2”.

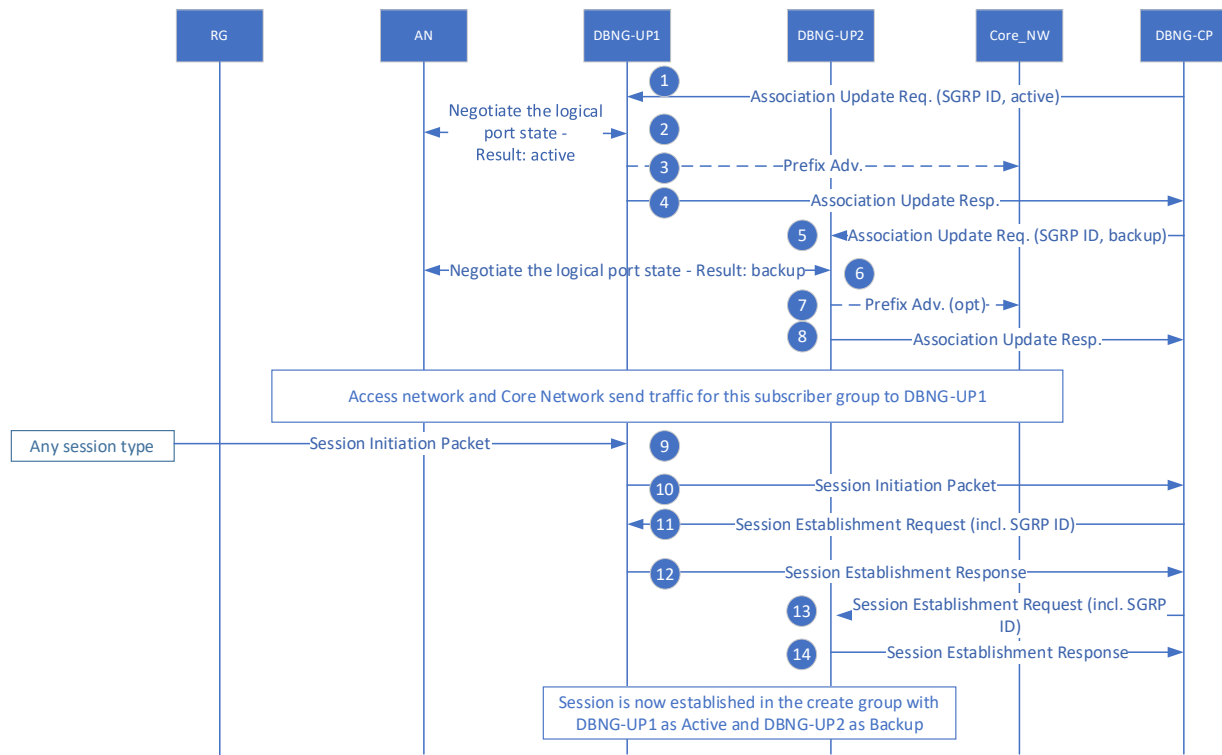


Figure 76: Establishment of a resilient session call flow

As per section 4.4, resiliency requires that there is a mechanism in the access network connectivity for traffic to switchover from one UP to another. Whatever is the specific mechanism, the following call flow

assumes that there are ways for the DBNG to indicate at any time to the Access Network which link should be Active.

1. The DBNG-CP sends the Association Update Request to DBNG-UP1 for creating the SGRP and set its state as active. Subsequent to the SGRP creation, the DBNG-CP associates one or more prefixes to the SGRP and programs the DBNG-UP1 on how to advertise them to the core network.
2. The DBNG-UP1 indicates to the Access Network that the forwarding state of the link to DBNG-UP1 has to be active. The exact mapping of the preferred forwarding state to the underlying resiliency mechanism is subject to local configuration on the DBNG-UP. If the Access Network is not able to set the link forwarding state to active, an error needs to be propagated back to the DBNG-CP in the Association Update Response.
3. The DBNG-UP1 advertises the SGRP prefix(es) towards the core network according to the programming issued by the DBNG-CP.
4. The DBNG-UP1 sends back to DBNG-CP the Association Update Response.
5. The DBNG-CP sends the Association Update Request to DBNG-UP2 for creating the SGRP and set its state as backup. Subsequent to the SGRP creation, the DBNG-CP associates one or more prefixes to the SGRP and may program the DBNG-UP2 on how to advertise them to the core network. In absence of programming, the DBNG-UP2 will not advertise the prefixes to the core network (see point 7 below).
6. The DBNG-UP2 indicates to the Access Network that the forwarding state of the link to DBNG-UP2 has to be "backup". If the Access Network is unable to set the link forwarding state to backup, an error needs to be propagated back to the DBNG-CP in the Association Update Response.
7. Optionally DBNG-UP2 advertises the SGRP prefix(es) towards the core network. If so, this is accomplished according to the prefixes programming issued by the DBNG-CP, that makes the routes to DBNG-UP2 as the less preferred by the core network.
8. The DBNG-UP2 sends back to DBNG-CP the Association Update Response. At this point all access network and core network traffic for this subscriber group (SGRP) flows through DBNG-UP1.
9. An RG sends a Session Initiation Packet (e.g., PADI, DHCP discover, ...) to DBNG-UP1.
10. The DBNG-UP1 forwards the Session Initiation Packet to DBNG-CP.
11. After authentication, the DBNG-CP sends Session Establishment Request to DBNG-UP1. The Session Establishment Request includes the Subscriber Group Identifier (SGRP ID).
12. The DBNG-UP1 sends Session Establishment Response to the DBNG-CP. Session is now established with the SGRP in state Active on DBNG-UP1.
13. The DBNG-CP sends Session Establishment Request to DBNG-UP2. The Session Establishment Request includes the Subscriber Group Identifier (SGRP ID).

Note: if the DBNG-UP2 was programmed by the DBNG-CP to be oversubscribed, it may skip installing the full session state and, at its discretion, it may or may not install its full forwarding state.
14. The DBNG-UP2 sends Session Establishment Response to the DBNG-CP. Session is now established with the SGRP in state "Backup" on DBNG-UP2.

The Subscriber Group ID has to be sent in every Session Establishment Request described in all call flows of this Technical Report.

4.5.47.2 DBNG-CP Managed Resilience Switchover

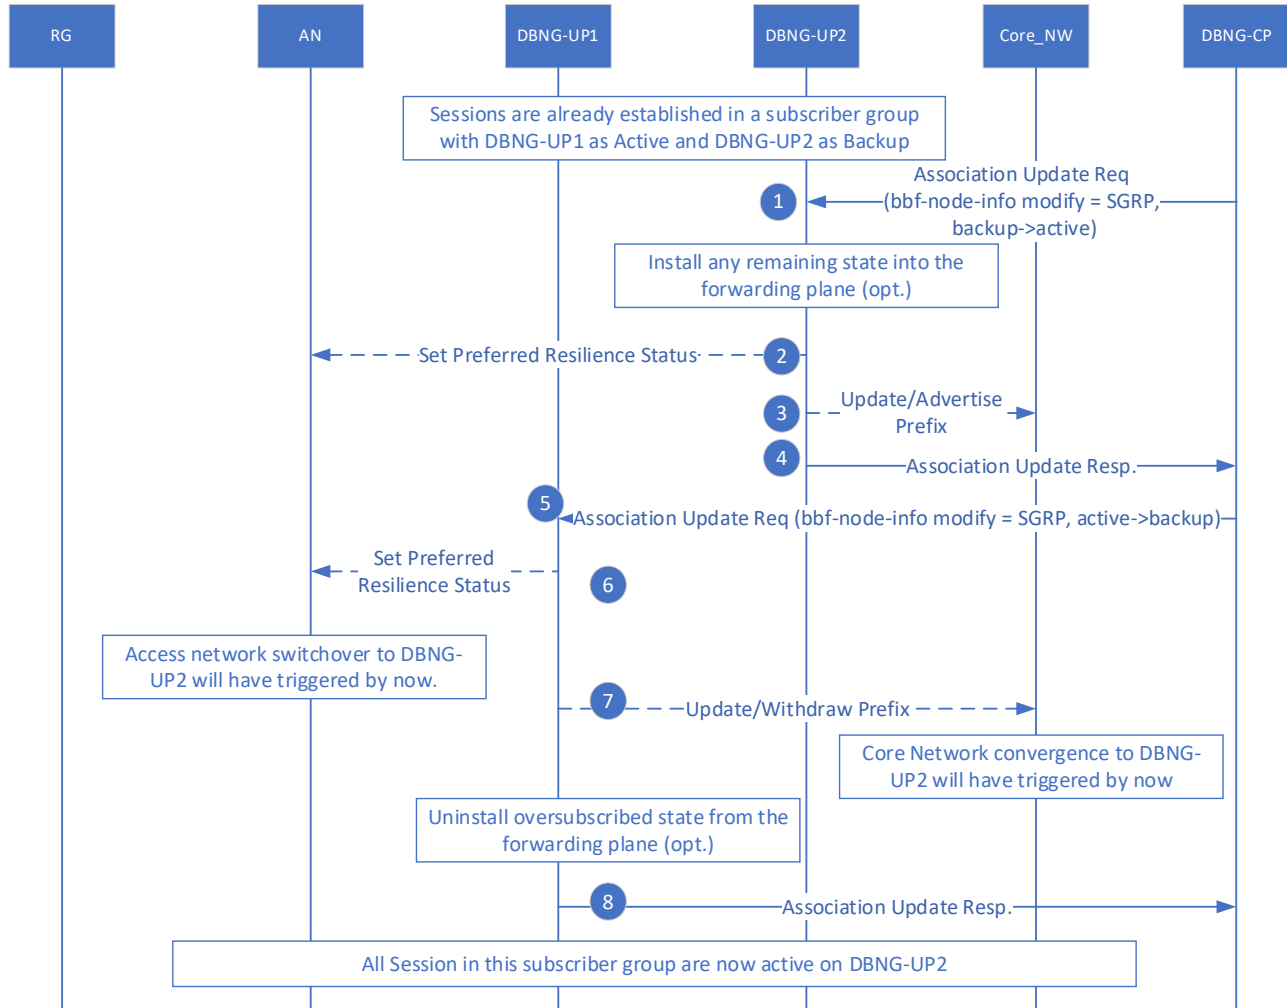


Figure 77: DBNG-CP Managed Resilience Switchover call flow

Note: It is left to implementation when reversion occurs

Prior to step 1, the call flow in section 4.5.47.1 assures that sessions are already established in a subscriber group in state Active on DBNG-UP1 and in state Backup on DBNG-UP2.

1. The DBNG-CP sends Association Update Request to DBNG-UP2 to activate the SGRP currently in state Backup. If not already done, the DBNG-UP2 upon receiving this message installs any remaining state into the forwarding plane.
2. The DBNG-UP2 indicates to the Access Network that the forwarding state of the link to DBNG-UP2 has to be active. The exact mapping of the preferred forwarding state to the underlying resiliency mechanism is subject to local configuration on the DBNG-UP. If the Access Network is not able to set the link forwarding state to active, an error needs to be propagated back to the DBNG-CP in the Association Update Response.
3. The DBNG-UP2 advertises or updates the previous advertisement of the SGRP prefix(es) towards the core network according to the previously signaled prefix programming by DBNG-CP.

4. The DBNG-UP2 sends back to DBNG-CP the Association Update Response to confirm that the SGRP state has been changed to active.
5. The DBNG-CP sends the Association Update Request to DBNG-UP1 to signal backup state for the SGRP.

Note: it is not required for step 1 to 4 to complete before initiating step 5 and onward when the fastest switchover is preferred over potential traffic loss.

6. The DBNG-UP1 upon receiving the SGRP Backup state programming, indicates to the Access Network that the forwarding state of the link to the DBNG-UP1 has to be "backup". If the Access Network is unable to set the link forwarding state to backup, an error needs to be propagated back to the DBNG-CP in the Association Update Response.
7. The DBNG-UP1 either withdraws or updates the previous advertisement of the SGRP prefix(es) towards the core network accordingly to the prefixes programming issued by the DBNG-CP. This makes the routes to DBNG-UP1 as the less preferred by the core network.
8. The DBNG-UP1 sends back to DBNG-CP the Association Update Response.
Note: if the DBNG-UP1 was programmed by the DBNG-CP to be oversubscribed, it may skip installing the full session state and, at its discretion, it may choose to remove its full forwarding state.

At this point all sessions in the subscriber group are now active on DBNG-UP2.

4.5.47.3 Establishment of a resilient session with track-logical-port

The DBNG-CP can establish a resilient session if it controls at least two distinct DBNG-UP instances and if the access network provides a dual-homed logical-port active-backup connectivity onto the two DBNG-UP instances. The key difference compared to Section 4.5.47.1 (non-track-logical-port resilient session) is that the active state of the SGRP is determined by the DBNG-UP instances themselves based on the active state of the logical port. Whenever the logical port's state changes, the DBNG-UP instances adjust the mastership of the SGRP as they track it at the same time and thus SGRP mastership changes can occur without the intervention of the DBNG-CP. In the following figure, we identify the two DBNG-UP instances under consideration as "DBNG-UP1" and "DBNG-UP2".

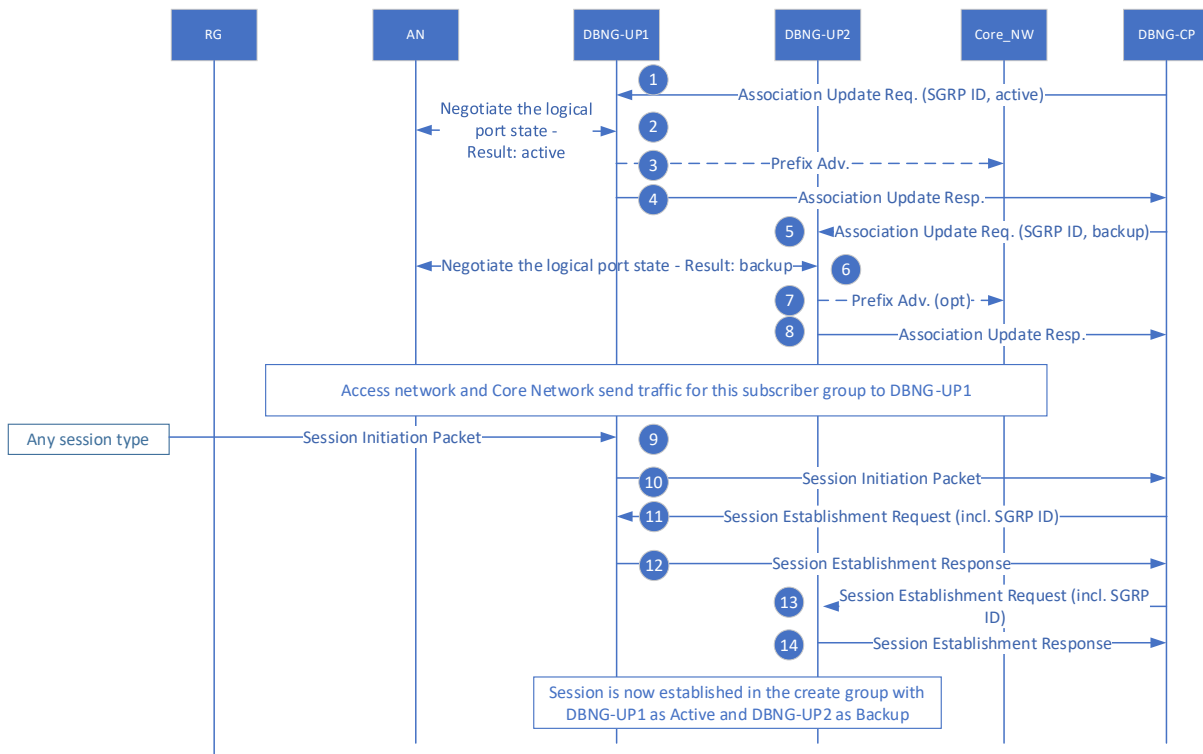


Figure 78: Establishment of a resilient session call flow with Track-Logical-Port

1. The DBNG-CP sends the Association Update Request to DBNG-UP1 for creating the SGRP and set its state to "Track Logical Port"; in the same message, the DBNG-CP indicates the logical port that has to be tracked for that SGRP by DBNG-UP1. Subsequently, the DBNG-CP associates one or more prefixes to the SGRP and programs the DBNG-UP1 on how to advertise them to the core network. The DBNG-UP1 now has the SGRP programmed with its associated information (prefixes, logical port to track, instructions on advertisements).
Note: the logical port is specific on DBNG-UP1 and can have a different name from the corresponding logical port on DBNG-UP2.
2. The DBNG-UP1, upon receiving the SGRP create message, finds the given logical port to be in active state on DBNG-UP1 and determines that SGRP should be active.
3. The DBNG-UP1 advertises the SGRP prefix(es) towards the core network as the preferred prefix(es) according to the programming issued by the DBNG-CP.

4. The DBNG-UP1 sends back to DBNG-CP the Association Update Response, in which it includes the BBF SGRP Notification Report: this report notifies to the DBNG-CP that the SGRP operational condition is active.
5. The DBNG-CP sends the Association Update Request to DBNG-UP2 for creating the SGRP and set its state to "Track Logical Port"; in the same message, the DBNG-CP indicates the logical port that has to be tracked for that SGRP by DBNG-UP2. Subsequently, the DBNG-CP programs on the DBNG-UP2 the prefix(es) and how they have to be advertised to the core network. The DBNG-UP2 now has the SGRP programmed with its associated information (prefixes, logical port to track, instructions on advertisements).
6. The DBNG-UP2, upon receiving the SGRP create message, finds the given logical port to be in backup state on DBNG-UP2 and determines that SGRP should be backup.
7. Optionally, DBNG-UP2 advertises the SGRP prefix(es) towards the core network. If so, this is accomplished according to the prefixes programming issued by the DBNG-CP, that makes the routes to DBNG-UP2 as the less preferred by the core network.
8. The DBNG-UP2 sends back to DBNG-CP the Association Update Response, in which it includes the BBF SGRP Notification Report: this report notifies to the DBNG-CP that the SGRP operational condition is backup. At this point all access network and core network traffic for the SGRP flows through DBNG-UP1.

Note: From this point forward, both DBNG-UP instances continue tracking the logical port's state locally, and autonomously adjust the local SGRP state depending on the state of the logical port as seen by respective DBNG-UP. If there is any change, DBNG-UP instances generate notifications containing the new SGRP operational condition to the DBNG-CP.

9. An RG sends a Session Initiation Packet (e.g., PADI, DHCP discover, ...) to DBNG-UP1.
10. The DBNG-UP1 forwards the Session Initiation Packet to DBNG-CP.
11. After authentication, the DBNG-CP sends Session Establishment Request to DBNG-UP1. The Session Establishment Request includes the Subscriber Group Identifier (SGRP ID).
12. The DBNG-UP1 sends Session Establishment Response to the DBNG-CP. Session is now established with the SGRP in state Active on DBNG-UP1.
13. The DBNG-CP sends Session Establishment Request to DBNG-UP2. The Session Establishment Request includes the Subscriber Group Identifier (SGRP ID).
Note: if the DBNG-UP2 was programmed by the DBNG-CP to be oversubscribed, it may skip installing the full session state and, at its discretion, it may or may not install its full forwarding state.
14. The DBNG-UP2 sends Session Establishment Response to the DBNG-CP. Session is now established with the SGRP in state Backup on DBNG-UP2.

Notes:

- The Subscriber Group ID has to be sent in every Session Establishment Request described in all call flows of this Technical Report.
- With regards steps no. 11 and no. 13, the DBNG-UP must respond with an error if the logical port in the Session Establishment Request does not match the logical port that was previously communicated to be part of this Subscriber Group.
- Steps 1-4 can run in parallel with steps 5-8. In the described call flow, it is assumed that DBNG-UP1 becomes active, while DBNG-UP2 becomes backup; however, the result of the negotiation on the dual-homed connectivity may be the opposite. There may be control mechanisms to make a link to be preferred over the other, but the way to configure these mechanisms work is out-of-scope for this document.

4.5.47.4 Independent DBNG-UP Switchover

This call flow describes a case where the DBNG-UP takes action without requiring intervention from the DBNG-CP.

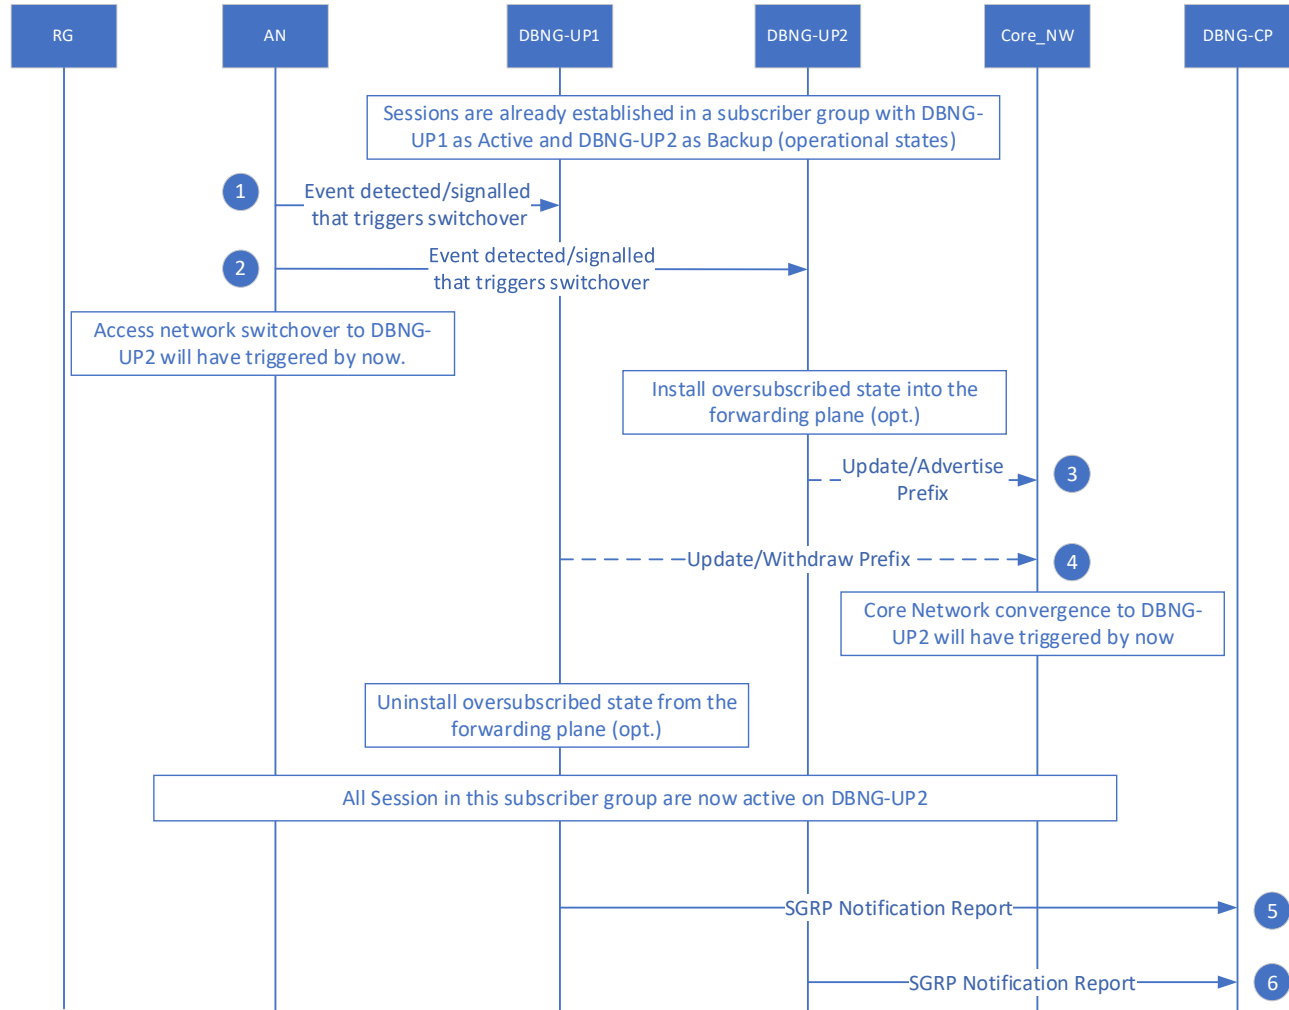


Figure 79: Independent DBNG-UP Switchover call flow

Note: depending upon the exact trigger and timing of response, the items may take place in a different order.

Prior to step 1, the call flow in section 4.5.47.3 assures that sessions are already established in a subscriber group and are working in operational state Active on DBNG-UP1 and in operational state Backup on DBNG-UP2.

In the call flow described in this section, it is assumed that there occurs an event – detected by the AN or signaled to the AN – which impacts the connectivity between the RG and DBNG-UP1, triggering a switchover from DBNG-UP1 to DBNG-UP2 without an intervention from the DBNG-CP.

However, an independent DBNG-UP switchover can be also triggered by an event that one or both the DBNG-UP are able to directly detect: for example, the active DBNG-UP crashes or there is a fault on its (active) logical port. In these cases, steps from 1 to 2 are skipped.

1. The Access Network notifies to the DBNG-UP1 that a logical port switchover is needed exploiting the underlying protocol that regulates the dual-homed connectivity. This automatically changes the operational state of the logical-port and of the SGRP on the DBNG-UP1 from active to standby.
2. The Access Network notifies to the DBNG-UP2 that a logical port switchover is needed exploiting the underlying protocol that regulates the dual-homed connectivity. This automatically changes the operational state of the logical-port and of the SGRP on the DBNG-UP2 from standby to active. At this point the access network switchover to DBNG-UP2 has already been initiated.
3. The DBNG-UP2 become aware of handling the SGRP in active mode. If not already done, DBNG-UP2 installs any remaining state into the forwarding plane. Moreover, the DBNG-UP2 advertises or updates the previous advertisement of the SGRP prefix(es) towards the core network, so that the routes become the preferred ones.
4. The DBNG-UP1 either withdraws or updates the previous advertisement of the SGRP prefix(es) towards the core network, so that the routes become the less preferred ones. At this point the core network convergence to DBNG-UP2 has already been initiated.
Note: if the cause of the switchover is a crash of the active DBNG-UP, this step may not occur. The core network will detect the fault via other mechanisms, which are out-of-scope for this document.
5. The DBNG-UP1 sends back to DBNG-CP the SGRP Notification in Node Report: this report notifies to the DBNG-CP that the SGRP operational condition is now backup on DBNG-UP1.

Notes:

- if the DBNG-UP1 was programmed by the DBNG-CP to be oversubscribed, it may skip installing the full session state and, at its discretion, it may choose to remove its full forwarding state.
 - if the cause of the switchover is a crash of the active DBNG-UP, step 5 may not occur. The DBNG-CP may detect the fault via the heartbeat mechanism; in the meantime, it should receive a complementary PFCP Node Report from the DBNG-UP2.
6. The DBNG-UP2 sends back to DBNG-CP the SGRP Notification in Node Report: this report notifies to the DBNG-CP that the SGRP operational condition is now active on DBNG-UP2.

After step 4, all sessions in this subscriber group are active on DBNG-UP2.

Note: Steps 1 and 2 may occur in parallel. Steps 3 and 4 may occur in parallel. Steps 5 and 6 may occur in parallel.

5 Technical Requirements

5.1 State Control Interface requirements

The section lists the functional requirements for a SCi protocol that DBNG-CP and DBNG-UP use to communicate and maintain the required functions and services.

- [R-1] The State Control Interface (SCi) protocol MUST support the communication of traffic detection and matching forwarding rules to allow the handling of subscriber traffic over the V-interface.
- [R-2] The SCi protocol MUST support the communication of forwarding rules to allow the handling of subscriber traffic over the V-interface.
- [R-3] The SCi protocol MUST support the communication of traffic detection and matching forwarding rules to allow the handling of subscriber traffic over the A10 interface. For example, Ethernet, MPLS, L2TP, and GTP.
- [R-4] The SCi protocol MUST support the communication of forwarding rules to allow the handling of subscriber traffic over the A10 interface. For example, Ethernet, MPLS, L2TP, and GTP.
- [R-5] The SCi protocol MUST enable the DBNG to support all the functions and services supported by a single node-based MS-BNG as documented in Table 1.
- [R-6] The SCi protocol MUST support the multiple access technologies used by functions and services required by Table 1.
- [R-7] The SCi protocol MUST provide reliable communication between the DBNG-CP and the DBNG-UP entities of the DBNG.
- [R-8] The SCi protocol MUST be extensible, e.g., allow introduction of the new information elements in a backward compatible manner.
- [R-9] The SCi protocol MUST be able to monitor reachability and liveness of the entities in the association.
- [R-10] The SCi protocol MUST support the redundancy function requirements [R-112] thru [R-117].

Security and privacy are critical for the SCi protocol. The following are detailed requirements to ensure secure communication and minimize the risk of attacking DBNG-CP and/or DBNG-UP.

- [R-11] It MUST be possible to encrypt data exchanged by the protocol.
- [R-12] It MUST be possible to authenticate the source of data exchanged by the protocol.
- [R-13] The SCi protocol MUST be able to convey rules to DBNG-UP to determine types of control packets to be re-directed to DBNG-CP.

To fulfill MS-BNG functionality, the DBNG-CP and DBNG-UP form associations. It is possible for each association to support a different subset of functions. The protocol should be able to exchange information on supported features.

[R-14] The SCi protocol MUST support capabilities determination between the DBNG-CP and DBNG-UP.

[R-15] The SCi protocol MUST be able to support a DBNG-CP to:

- add per-subscriber access routes on a DBNG-UP
- update per-subscriber access routes on a DBNG-UP
- delete per-subscriber access routes on a DBNG-UP

Note: access routes in [R-16] refer to framed routes, host IP address(es), host Ipv6 address(es), and host Ipv6 prefix(es).

[R-16] For all of the access types documented on fourth column of Table 2, the SCi protocol MUST support a DBNG-CP to:

- add per-subscriber session forwarding states on DBNG-UP(s).
- update per-subscriber session forwarding states on DBNG-UP(s).
- delete per-subscriber session forwarding states on DBNG-UP(s).

- [R-17] The SCi protocol MUST support the programming of DBNG-UP prefix(es) through the use of SGRP
- [R-18] The SCi protocol MUST be able to support session status synchronization between a DBNG-CP and a DBNG-UP.
- [R-19] The SCi protocol MUST support the programming of subscriber groups (SGRP) by DBNG-CP on DBNG-Ups, in terms of creating, updating, and deleting.
- [R-20] The SCi protocol MUST be able to support a DBNG-CP to selectively apply lawful interception on DBNG-UP(s).

As specified in [R-1], a MS-BNG MUST support a variety of access types: fixed wireline, fixed wireless, and hybrid. Therefore, the DBNG-UP must be just as flexible in supporting the access types. To accomplish this, the DBNG-CP is responsible to program DBNG-UP forwarding information. The programming of forwarding states simply consists of:

- 1) Find the subscriber and match on packet types (**matching criteria**)
- 2) Perform one or more actions. (perform **action**)

The match and action rules would provide platform independent abstraction to derive data-path state and processing by the DBNG-UP. Table 5 shows examples of DBNG-CP using a SCi protocol instructing the DBNG-UP to install traffic detection and forwarding rules on the DBNG-UP.

Table 5: Examples of traffic detection and traffic forwarding rules

| Access Types | Upstream (access to network) | | Downstream (network to access) | |
|--------------------|--|---|----------------------------------|--|
| | Match | Action | Match | Action |
| PPPoE single stack | Port ETH header PPPoE Session ID Ipv4 | Remove ETH + PPPoE Apply QoS and Filter(s) Forward to network | Ipv4 | Encap. ETH + PPPoE Apply QoS and Filter(s) Forward to access |
| PPPoE dual Stack | Port ETH header PPPoE Session ID Ipv4 Ipv6 NA/PD Ipv6 SLAAC | Remove ETH + PPPoE Apply QoS and Filter(s) Forward to network | Ipv4 Ipv6 NA/PD Ipv6 SLAAC | Encap. ETH + PPPoE Apply QoS and Filter(s) Forward to access |
| IpoE single stack | Port ETH header Ipv4 | Remove ETH Apply QoS and Filter(s) Forward to network | Ipv4 | Encap. ETH Apply QoS and Filter(s) Forward to access |
| IpoE dual Stack | Port ETH header Ipv4 Ipv6 NA/PD Ipv6 SLAAC | Remove ETH encap. Apply QoS and Filter(s) Forward to network | Ipv4 Ipv6 NA/PD Ipv6 SLAAC | Encap. ETH Apply QoS and Filter(s) Forward to access |
| Public Wi-Fi | Port GRE tunnel | Remove GRE/ETH Apply QoS and Filter(s) | IP | Encap. ETH/GRE Apply QoS and Filter(s) |

| | | | | |
|------|--------------------------|---|-----|--|
| | ETH header IP | Forward to network | | Forward to access |
| TWAG | Port ETH header IP | Remove ETH Encap. GTP Apply QoS and Filter(s) Forward to network | GTP | Remove GTP Encap. ETH Apply QoS and Filter(s) Forward to access |
| HAG | Port ETH header IP | Remove ETH Apply QoS and Filter(s) Forward to network | IP | Encap. ETH Apply QoS and Filter(s) Forward to access |
| | GTP | Remove GTP Apply QoS and Filter(s) Forward to network | | Encap. GTP Apply QoS and Filter(s) Forward to s1u |

Supporting data forwarding for different types of access, as shown in Table 5, can be represented by a list of matching criteria and action rules. The traffic detection and forwarding rules can be combined to accomplish the following:

- Flexible to cover all different access types for current MS-BNG deployments
- Flexible to cover the A10 transport protocol as specified on Table 3 used for current MS-BNG deployments.
- Extensible to handle future types of access and transport protocols
 - o Covering future types of access and transport SHOULD NOT require versioning

[R-21] The SCi protocol MUST be able to signal a set of packet matching rules and set of actions for each individual subscriber session from the DBNG-CP to the DBNG-UP.

Note: the method of how rules and actions are signaled is up to the protocol, for example, expressing the relationship between the set of rules and actions is protocol specific.

A native MS-BNG function includes subscriber accounting, to account for both time and volume usage per subscriber. It is expected that the DBNG architecture will provide the equivalent capability. Based on the accounting function, service provider can offer services such as monthly flat rate or credit based broadband access. Credit control refers to subscriber utilizing monetary means to exchange credits in the form of time, data volume, or a combination of both for broadband service. The DBNG-UP would be instructed to monitor the subscriber usage. As the credits are exhausted, the DBNG-UP would inform the DBNG-CP, followed by a subsequent action including, degrading the service, redirecting to a web portal, or stopping the service entirely.

Accounting is the ability to report periodic or upon demand the subscriber usage based on time and data. The accounting record can also contain the amount of time and/or volume consumed by the subscriber.

[R-22] The SCi protocol MUST be able to perform credit control for usage monitoring and reporting.

[R-23] In case UP based volume quota and thresholds capability and SGRP resiliency are combined, the DBNG-CP MUST coordinate the consumed quota between the active and backup DBNG-Ups based on received usage reports.

As specified in TR-178 [10] section 7.1.4 and 7.1.6, MS-BNG has a variety of QoS mechanisms such as scheduling, traffic shaping, and traffic policing. The DBNG is expected to provide the same set of QoS mechanisms. The SCi protocol must provide means for the DBNG-CP to communicate QoS parameters to the DBNG-UP.

[R-24] The SCi protocol MUST be able to set and modify QoS policy per subscriber session.

A MS-BNG that is also acting as a PE (as per TR-178 [10]) may have multiple virtual networks to which a subscriber may be connected.

- [R-25] The SCi protocol MUST support the identification of virtual network to which the subscriber should be connected.
- [R-26] The SCi protocol MUST support the communication of an action rule with a recursive lookup.
- [R-27] The SCi protocol MUST support the ability to request and confirm the offload of PPP keep-alive generation and processing including the generation and response to LCP Echo-Request and the processing of LCP Echo-Reply messages from the DBNG-CP to the DBNG-UP on a per session basis.
- [R-28] The SCi protocol MUST support the ability to disable offload of PPP keep-alive generation and processing including the generation and response to LCP Echo-Request and the processing of LCP Echo-Reply messages on a per session basis.
- [R-29] The SCi protocol MUST support the ability to advertise the presence of the PPP keep-alive generation and processing capability including the generation and response to LCP Echo-Request and the processing of LCP Echo-Reply messages offloaded to the DBNG-UP.

Keepalive offload is the recommended mode of operation for the DBNG architecture. Offload allows the DBNG to achieve higher scalability and mitigate connectivity faults between the DBNG-CP and DBNG-UP.

- [R-30] The SCi protocol MUST support the communication of the desired resilience state for groups of subscriber sessions.
- [R-31] At the request of DBNG-CP received via SCi protocol, DBNG-UP must be able to execute traffic mirroring for a specific session towards the LEA DF.
- [R-32] At the request of DBNG-CP received via SCi protocol, DBNG-UP must be able to stop traffic mirroring for a specific session towards the LEA DF.
- [R-33] The SCi protocol MUST support signaling of SGRP operational and error conditions.
- [R-34] The SCi protocol MUST support signaling ACL chaining for a subscriber session from DBNG-CP to DBNG-UP.
- [R-35] The SCi protocol MUST support signalling a logical port status from DBNG-UP to DBNG-CP, which can be used to derive a desired resilience state on the DBNG-CP
- [R-36] The SCi protocol MUST support signalling a Network Instance status from DBNG-UP to DBNG-CP, which can be used to derive a desired resilience state on the DBNG-CP.
- [R-37] The SCi protocol SHOULD support DBNG-CP to indicate the mode of session statistics: only Ipv4 statistics, only Ipv6 statistics or both Ipv4&Ipv6 statistics to DBNG-UP.
- [R-38] The SCi protocol MUST support reporting only the Ipv4 session statistics from DBNG-UP to DBNG-CP.
- [R-39] The SCi protocol MUST support reporting only the Ipv6 session statistics from DBNG-UP to DBNG-CP.
- [R-40] The SCi protocol MUST support reporting both Ipv4&Ipv6 session statistics simultaneously from DBNG-UP to DBNG-CP.
- [R-41] The SCi MUST support signalling a logical port usage from DBNG-UP to DBNG-CP, which for example can be used by the DBNG-CP for subscriber placement.

5.2 Requirements to support Control packet redirect interface

The CPR Interface is an interface between DBNG-CP and DBNG-UP. It is used to tunnel control packets between the V or the A10 interface to the DBNG-CP via the DBNG-UP.

- [R-42] Subscriber control messages to and from RGs MUST be tunneled between DBNG-CP and DBNG-UP through the CPR interface.

- [R-43] In the case where the DBNG is performing LNS function, Control messages to and from LAC MUST be tunneled between DBNG-CP and DBNG-UP through the CPR interface.
- [R-44] In the case where the DBNG is performing LAC function, Control messages to and from LNS MUST be tunneled between DBNG-CP and DBNG-UP through the CPR interface.
- [R-45] In the case of DHCP relay, where the DHCP server is connected to the DBNG-UP, control messages to and from the DHCP server MUST be tunneled between DBNG-CP and DBNG-UP through the CPR interface.
- [R-46] The DBNG-CP MUST be able to communicate rules to the DBNG-UP to match specific control packet types and forward these to the DBNG-CP over a default CPR interface when no subscriber session context exists on the DBNG-UP
- [R-47] The DBNG-UP MUST pass access interface information (e.g., port or virtual-port on which the control traffic is received) together with the tunneled subscriber control packet to the DBNG-CP where applicable. E.g., control packets: DHCP, PPP, RA, data-trigger
- [R-48] The DBNG-CP MUST pass access interface information (e.g., port or virtual-port on which the control traffic is received) together with the tunneled control packet to the DBNG-UP where applicable.
- [R-49] The DBNG-CP MUST be able to signal the DBNG-UP to update the forwarding action matching specific control messages from the V interface per subscriber
- [R-50] The DBNG-CP MUST be able to send control packets to the A10 interface through the DBNG-UP
- [R-51] The DBNG-CP MUST be able to signal the DBNG-UP to update the forwarding action matching specific control messages from the A10 interface per subscriber

The selection and redirection of messages to the DBNG-CP must be flexible in accommodating the many types of services provided by the DBNG. E.g., some subscribers connect to the internet using PPPoE, IPTV multicast services through DHCP, and only enterprise customers are using data-trigger service.

- [R-52] The SCi protocol MUST allow the DBNG-CP to signal DBNG-UP to redirect only specific control packets per match criteria. E.g., match criteria could be as granular as per subscriber or as general as per DBNG-UP node

Subscriber control packets can be rate limited to protect the DBNG-CP. The DBNG may apply rate limiting per subscriber basis in order to limit the traffic of malicious users and protect the rest of the subscribers: e.g., it may apply rate limiting to data-trigger redirected packets and it may prioritize control messages for authenticated subscribers versus subscribers yet to be authenticated. The rate limiting is programmed so that undesired control packets are filtered out at the DBNG-UP before reaching the CPR interface. For the Requirements below, the match criteria can be as specific as a subscriber or a logical port or as general as DBNG-UP.

- [R-53] The CUPS protocol MUST allow the DBNG-CP to signal the DBNG-UP the priority of specific control messages per match criteria.
- [R-54] The DBNG-UP MUST be able to perform rate limiting on the control packets per match criteria
Note: control packets include data-trigger packets.

When operating as a DHCP relay or relay-proxy or DHCPv6 relay with an external DHCP/DHCPv6 server connected to the DBNG-UP (and not reachable directly from the DBNG-CP), the selection and redirection of messages between DBNG-CP and the external DHCP/DHCPv6 server must be flexible in supporting on a per network instance basis for IpoE subscribers assigned to the network instance.

- [R-55] When operating as a DHCP/DHCPv6 relay with the external DHCP/DHCPv6 server reachable via the DBNG-UP, the DBNG-CP MUST be able to signal control packet redirection rules to instruct the DBNG-UP to forward subscriber-originated DHCP/DHCPv6 control packets from the DBNG-CP to the external DHCP/DHCPv6 server via the A10 interface on a per network instance basis.
- [R-56] When operating as a DHCP/DHCPv6 relay with the external DHCP/DHCPv6 server reachable via the DBNG-UP, the DBNG-CP MUST be able to signal control packet redirection rules to instruct the DBNG-UP to tunnel subscriber-bound DHCP/DHCPv6 control messages received from the external DHCP/DHCPv6 server via the A10 interface to the DBNG-CP on a per network instance basis.

5.3 Management Interface requirements

The Management interface on DBNG-CP and DBNG-UP provides two main functions: configuration and operational state retrieval of DBNG-UP.

One of the main functions of the CUPS Management Interface is configuration of functions and services. Data models present the most powerful and flexible approach to configure devices, services, and to monitor their operational state. Also, it is advantageous to support the ability to use machine tools to automate generation, manipulation, and parsing of the configuration data received state information. Extensible Markup Language (XML) was designed to store and transport data. XML was designed to be self-descriptive and is human-readable.

[R-57] The Management Interface **MUST** support transactional configuration from DBNG-CP to DBNG-Ups based on YANG data model

Operational data can either be streamed (published) at configured cadence or on-change/event. With on-change/event, data is streamed only when a change in the data occurs. Operational data can be transported using different protocols and in different encoding formats.

[R-58] The Management Interface **MUST** support operational state retrieval based on YANG data model

[R-59] The Management Interface **MUST** support operational state retrieval in XML via NETCONF

[R-60] The Management Interface **SHOULD** support operational state retrieval in JSON via RESTCONF

[R-61] The Management Interface **MUST** support publishing of state information at configurable cadence or on-event based on YANG data model

[R-62] The Management Interface **MUST** support publishing of state information in XML via NETCONF

[R-63] The Management Interface **SHOULD** support publishing of state information in JSON via RESTCONF

5.4 Disaggregated MS-BNG control plane requirements

DBNG-CP is responsible for the wireline access subscriber management as well as its address management, and user plane management.

[R-64] The DBNG-CP **MUST** support Ipv4/Ipv6 address pool management;

[R-65] The DBNG-CP **MUST** support Ipv4/Ipv6 address prefix and prefix length allocation

[R-66] The DBNG-CP **MUST** support subscriber state management of DBNG-UP;

[R-67] The DBNG-CP **MUST** support the establishment and release of association with DBNG-UP;

In section 4.5.4, in step 3, after the subscriber has successfully authenticated, the DBNG-CP would assign an IP address to the subscriber. It is possible for the DBNG-CP to either have a pre-allocated static IP for the subscriber or dynamically assign the next IP address available for the subscriber. This is applicable to all call flows in section 4.5.

[R-68] The DBNG-CP **MUST** be able to support static address prefix allocation.

[R-69] The DBNG-CP **MUST** be able to support dynamic address prefix allocation on behalf of DBNG-UP upon subscriber access control procedure.

[R-70] In the case that the DBNG-CP has an integrated User Plane Selection Function, the DBNG-CP **MUST** support communication with a Traffic Steering Function as per requirements 0 to [R-74].

[R-71] In the case that the DBNG-CP has an integrated User Plane Selection Function, the DBNG-CP **MUST** select the DBNG-UP that is to serve a newly connecting subscriber.

[R-72] In the case that the DBNG-CP has an integrated User Plane Selection Function, the DBNG-CP **MUST** identify required changes in serving DBNG-UP for an established subscriber.

[R-73] In the case that the DBNG-CP has an integrated User Plane Selection Function, , the DBNG-CP MUST signal the target DBNG-UP to the Traffic Steering Function.

Note: The exact mechanism by which the DBNG-UP Selection Function signals the Traffic Steering Function is not defined here, but it is expected that this may be based on communicating the requirement via an SDN controller.

[R-74] Where the DBNG-CP has an integrated User Plane Selection Function, the User Plane Selection Function, the DBNG-UP Selection Function MUST be able to identify the DBNG-UP which will serve a specific subscriber based on policy. The policy may take into account the following criteria:

- Subscriber policy information from AAA.
- Current load of the DBNG-UP elements.
- DBNG-Ups that are not available (such as undergoing maintenance).
- DBNG-UP to which the subscriber cannot be connected (for network topology reasons).
- Network performance (such as latency) between the subscriber and the DBNG-UP.
- Placement of other subscribers with common attributes such as IP Subnet.

[R-75] The DBNG-CP MUST support PPP keep-alive processing, including the generation and response to LCP Echo-Request and the processing of LCP Echo-Reply messages

[R-76] The DBNG-CP MUST support NETCONF for integration with EMS

[R-77] The DBNG-CP SHOULD support RESTCONF for integration with EMS

[R-78] The DBNG-CP MAY support SNMP for integration with EMS

[R-79] DBNG-CP MUST be able to handle the signaling with the LEA MD to receive instructions for lawful interception actions on specific targets.

[R-80] The DBNG-CP MUST keep reliable tables of Lawful Interception targets.

[R-81] The DBNG-CP MUST generate LI start, LI end and LI failure event notifications towards the MD for all Lawful Interception targets.

[R-82] The DBNG-CP MUST generate an event notification towards the MD when it is ready for servicing LI requests.

[R-83] The DBNG-CP MUST be able to create Ipv4 or Ipv6 prefix information for a group of subscriber sessions on the DBNG-UP via SCi.

[R-84] The DBNG-CP MUST be able to modify Ipv4 or Ipv6 prefix information for a group of subscriber sessions on the DBNG-UP via SCi.

[R-85] The DBNG-CP MUST be able to delete Ipv4 or Ipv6 prefix information for a group of subscriber sessions on the DBNG-UP via SCi.

[R-86] The Ipv4 or Ipv6 prefix information message MUST allow a vendor specific option to be included with the subscriber prefix information to convey additional information beyond that defined by standard attributes.

[R-87] The DBNG-CP, when assigning a prefix to DBNG-UP via SCi, SHOULD indicate the Network Instance to which the prefix belongs. Note: if the Network Instance is not indicated, the prefix is intended to belong to the default Network Instance

[R-88] The DBNG-CP, when assigning an Ipv4 prefix to DBNG-UP via SCi, MAY specify a /32 address belonging to the prefix as default gateway. If so, this /32 address MUST be used by the DBNG-UP as source address to reply to ARP requests. Note: specifying an Ipv4 default gateway is not necessary for PPPoE-based access.

[R-89] The DBNG-CP MAY assign an Ipv6 link-local-address to the DBNG-UP by signaling a prefix with prefix-length 128. If so, this /128 prefix MUST be used by the DBNG-UP to respond to NS requests. Note: specifying an Ipv6 default link-local-address is not necessary for PPPoE-based access

[R-90] The DBNG-CP MUST be able to establish a default CPR tunnel on the DBNG-UP to redirect all subscriber initial control packets.

- [R-91] The DBNG-CP **MUST** support the filtering requirements [R-137], [R-138], [R-139] through local configuration on the default CPR interface. Note: If [R-111] is supported, the DBNG-CP is not expected to perform the filtering.
- [R-92] The DBNG-CP **MAY** install both a default and per-logical port tunnel on the same DBNG-UP. In this case, the per-logical-port PDRs **MUST** be signaled with a higher precedence than the default tunnel PDRs.
- [R-93] In the case the DBNG-CP receives indication for NSH header insertion support from the DBNG-UP, the DBNG-CP **MAY** choose to request for NSH header insertion through BBF outer header creation.
- [R-94] The DBNG-CP **MUST** have the ability to assign a different SGRP for each subscriber session within a logical port.
- [R-95] The DBNG-CP **MUST** be able to query the DBNG-UP for statistics of a particular subscriber session.
- [R-96] The DBNG-CP **MUST** be able to utilize the statistics queries for accounting purposes.
- [R-97] The DBNG-CP **MUST** be able to utilize the statistics queries for operational purposes.
- [R-98] The DBNG-CP **MUST** be able to acknowledge subscriber session periodic statistics reports from the DBNG-UP.
- [R-99] The DBNG-CP **MUST** be able to utilize the statistics reports for accounting and operational purposes. The DBNG-CP **MAY** use the periodic statistics report as a trigger to send a periodic accounting report.
- [R-100] Upon completion of Session Deletion procedure, if Accounting is enabled, the DBNG-CP **MUST** generate AAA Accounting Stop message related to the deleted session.
- [R-101] Upon completion of Session Deletion procedure, if Accounting is enabled, the Accounting Stop counters **SHOULD** be aligned with the latest statistics updates available to the DBNG-CP.

5.5 Disaggregated MS-BNG user plane requirements

DBNG-UP is responsible for routing, forwarding subscriber data, and terminating subscriber L2 data traffic. Apart from subscriber data termination and forwarding, user plane runs as a gateway between the user and the control plane. It could reside in dedicated devices which are designed specifically for forwarding performance or in virtual machines.

- [R-102] The DBNG-UP **MUST** support association and dissociation with DBNG-CP
- [R-103] The DBNG-UP **SHOULD** support network management interfaces to the operator's EMS.
- [R-104] The DBNG-UP **MUST** support V-interface encapsulations.
- [R-105] The DBNG-UP **MUST** support A10 interface encapsulations.
- [R-106] The DBNG-UP **MUST** be able support the offload of PPP keep-alive generation and processing including the generation and response to LCP Echo-Request and the processing of LCP Echo-Reply messages.
- [R-107] In the case that the DBNG-UP supports PPP offload as per [R-106], this **MUST** be supported on a per session basis.
- [R-108] A DBNG-UP **MUST NOT** generate offloaded keepalive messages for a session that is in backup state.
- [R-109] The DBNG-UP **MUST** be able to reply to ARP and ICMP destined for the Ipv4 gateway address
- [R-110] The DBNG-UP **MUST** be able to reply to ND and ICMPv6 destined for the Ipv6 gateway address
- [R-111] The DBNG-UP **SHOULD** support offload of the filtering requirements [R-137], 0, [R-139] configured via the SCI.

- [R-112] The DBNG-UP MAY support per logical port CPR tunnel. In this case, the DBNG-UP MUST advertise this support using a BBF UP function feature flag, and the DBNG-CP MAY choose to install a per logical port tunnel to that DBNG-UP.
- [R-113] The DBNG-UP MAY support NSH header insertion on any CPR tunnel.
- [R-114] In the case that the DBNG-UP supports NSH header insertion, the DBNG-UP MUST signal this support as part of BBF UP functional feature flag.
- [R-115] The DBNG-UP SHOULD include its local time when sending any subscriber statistics report.
- [R-116] In the case DBNG-UP supports ACL, DBNG-UP MUST apply SDF filters and subsequently apply ACL or ACL Chain.
- [R-117] Where the DBNG supports [R-129], the DBNG-UP MUST generate a logical port status report as soon as the status changes. This report update may be delayed only for message throttling purposes.
- [R-118] Where the DBNG supports [R-129], the DBNG-UP MUST generate a Network Instance status report as soon as the status changes. This report update may be delayed only for message throttling purposes.
- [R-119] The DBNG-UP MUST support signaling of SGRP operational and error conditions to DBNG-CP.
- [R-120] The DBNG-UP MUST be able to respond to a DBNG-CP statistics query for a particular subscriber session.
- [R-121] The DBNG-UP MUST be able to periodically report statistics for a particular subscriber session to the DBNG-CP, when provisioned to do so by the DBNG-CP.
- [R-122] The DBNG-UP MUST send final accounting statistics when receiving a PFCP session Deletion message. By default, these are sent in the PFCP session Deletion response message, but the DBNG-UP MAY send these in separate Session Report Request messages when the DBNG-CP indicates the capability for it.

5.6 Disaggregated MS-BNG functional requirements

- [R-123] The DBNG MUST support cold standby for redundant DBNG-CPs.
- [R-124] The DBNG MUST support cold standby for redundant DBNG-Ups.
- [R-125] The DBNG SHOULD support warm standby for redundant DBNG-CPs.
- [R-126] The DBNG SHOULD support warm standby for redundant DBNG-Ups.
- [R-127] The DBNG MAY support hot standby for redundant DBNG-CPs.
- [R-128] The DBNG MAY support hot standby for redundant DBNG-Ups.
- [R-129] The DBNG MUST be able to support the termination of resilient connections to the same Access Network split across two or more DBNG-Ups in an Active/Backup manner.
- [R-130] Where the DBNG supports [R-129], it MUST be possible for the resilience to be at a logical port level

It is also possible for resilience to be applied without being tied to a logical port.

- [R-131] Where the DBNG supports [R-129], it MUST be possible for the resilience to apply to a group of sessions that is not tied to a particular logical port.
- [R-132] Where the DBNG supports [R-129], it MUST be possible for the relevant subscriber session forwarding state to simultaneously exist on all relevant DBNG-Ups.

For the requirement [R-133], the resilient connection to the Access Network may include a signaling capability to dynamically select the active path or link. For example, the IEEE 802.3 Link Aggregation Control Protocol (LACP) can be used to negotiate active links within a link group and to detect component link failures.

- [R-133] Where the DBNG supports [R-129], the DBNG MUST take any relevant signaling from the Access Network into account when selecting which DBNG-UP is active for the relevant subscriber sessions.

The following requirement is relevant to CP-Triggered resilience switchover:

- [R-134] Where the DBNG supports [R-129] and there is a signaling capability between the DBNG and the access network, the DBNG MUST be able to influence the selection of the active DBNG-UP.
Note: The access network may need to be configured to respect the indication of active DBNG-UP in [R-134] above.
- [R-135] Where the DBNG supports [R-129], the DBNG MUST advertise prefixes for the subscribers in such a way that it attracts traffic to the currently active DBNG-UP for relevant subscriber sessions.
- [R-136] If a target session, for resilience purposes, is moved from a DBNG-UP to another DBNG-UP under the control of the same DBNG-CP, the DBNG as a whole must guarantee the operational continuity of the lawful interception.
- [R-137] The DBNG MUST support the ability to accept specific initial control packets (DHCPv4 Discover, DHCPv6 Solicit, PPPoE PADI, Ipv4 data packet, Ipv6 data packet) on a logical port.
- [R-138] The DBNG MUST support filtering of VID values in the range 1-4093 on a logical port based on IEEE 802.1Q.
- [R-139] The DBNG MUST support filtering of VID values in the range 1-4093 for stacked VLANs on a logical port based on IEEE 802.1AD(QinQ).
- [R-140] The DBNG SHOULD support multiple Logical Ports in a SGRP to support CP-driven switchover.
- [R-141] DBNG SHOULD support ACL for filtering of traffic on subscriber sessions.
- [R-142] DBNG SHOULD support ACL chaining for a subscriber session.

5.7 Requirements for 3GPP PFCP node messages

- [R-143] The PFCP heartbeat request procedure MUST comply with 3GPP TS 29.244 [28] sections 6.2.2.1 and 6.2.2.2.
- [R-144] The PFCP heartbeat response procedure MUST comply with 3GPP TS 29.244 [28] sections 6.2.2.1 and 6.2.2.3.
- [R-145] The DBNG MUST support PFCP association request and response initiated by the CP Function as defined in 3GPP TS 29.244 [28] sections 6.2.6.1 and 6.2.6.2.
- [R-146] The DBNG MUST support PFCP association request and response initiated by the UP Function as defined in 3GPP TS 29.244 [28] sections 6.2.6.1 and 6.2.6.3.
- [R-147] The DBNG-CP or DBNG-UP sending PFCP association setup request and response message MUST include the IEs specified as mandatory in Table 7.
- [R-148] The DBNG MUST support PFCP Association Update procedure initiated by the DBNG-CP as defined in 3GPP TS 29.244 [28] sections 6.2.7.1 and 6.2.7.2.
- [R-149] The DBNG MUST support PFCP Association Update procedure initiated by the DBNG-UP as defined in 3GPP TS 29.244 [28] sections 6.2.7.1 and 6.2.7.3.
- [R-150] The DBNG-CP and DBNG-UP sending PFCP Association Update Request and Response message MUST include the IEs specified as mandatory in Table 7.
- [R-151] The DBNG MUST support PFCP Association Release Procedure as defined in section 3GPP TS 29.244 [28] section 6.2.8.
Note: the PFCP Association Release Procedure can be initiated only by the DBNG-CP Function.
- [R-152] The DBNG-CP and DBNG-UP sending PFCP Association Release Request and Response messages respectively MUST include the IEs specified as mandatory in Table 7.
- [R-153] The DBNG-CP MUST support PFCP message PFCP Version Not supported as defined in 3GPP TS 29.244 [28] 7.4.4.7
- [R-154] The DBNG MUST support PFCP Node Report Procedure as defined in section 3GPP TS 29.244 [28] 6.2.9 PFCP Node Report Procedure
- [R-155] The DBNG-UP and DBNG-CP sending PFCP Node Report Request and Response respectively MUST include IEs as specified in Table 7

5.8 Requirements for 3GPP PFCP session messages

[R-156] The DBNG MUST support PFCP session establishment procedure as defined in 3GPP TS 29.244 [28] section 6.3.2.

Note: To interpret the text in 3GPP TS 29.244 [28] sections 6.3.2.2, 6.3.3.2 and 6.3.4.2 appropriately for DBNG, consider the terms “PDN connection” and “IP-CAN session” as if they have the meaning of DBNG subscriber session. This applies to [R-156], [R-157] and [R-159].

[R-157] The DBNG MUST support PFCP session Report procedure as defined in 3GPP TS 29.244 [28] section 6.3.5 PFCP Session Report Procedure

[R-158] The DBNG MUST support PFCP session modification procedure as defined in 3GPP TS 29.244 [28] section 6.3.3.

[R-159] The DBNG MUST support PFCP session deletion procedure as defined in 3GPP TS 29.244 [28] section 6.3.4.

[R-160] The DBNG-CP sending PFCP session establishment request MUST include the IEs specified as mandatory in Table 9.

[R-161] The DBNG-UP sending PFCP session establishment response MUST include the IEs specified as mandatory in Table 10.

[R-162] The DBNG-CP sending PFCP session modification request MUST include the IEs specified as mandatory in Table 11.

[R-163] The DBNG-UP sending PFCP session modification response MUST include the IEs specified as mandatory in Table 12.

[R-164] The DBNG-CP sending PFCP session deletion request MUST include the IEs specified as mandatory in Table 13.

[R-165] The DBNG-UP sending PFCP session deletion response MUST include the IEs specified as mandatory in Table 14.

[R-166] The DBNG Control plane MUST program via SCi the same PDI, PDR, FAR, URR and QER on the backup DBNG-UP that are installed on the active DBNG-UP for all resilient sessions.

6 PFCP CUPS protocol

PFCP is the selected CUPS protocol for the DBNG SCi and is used to program subscriber forwarding state and control packet redirection rules. PFCP is a 3GPP protocol standardized since release 14. Details of the protocol can be found in TS 29.244 [28] “Interface between the Control Plane and User Plane Node”. PFCP addresses the technical and functional requirements listed in this document.

PFCP contains two main message types: node messages and session messages. Node messages are mainly used to form and to maintain association between DBNG-CP and DBNG-UP; for DBNG, node messages have been enhanced to program subscriber groups and IP prefixes. Session messages are mainly used to program the subscriber forwarding state. Both node and session messages utilize information elements (IEs) for communication between DBNG-CP and DBNG-UP. Most IEs are extensible, details on IE extensibility are covered in TS 29.244 [28] Table 8.1.2-1. The following section describes the IE extensions required to support various MS-BNG use cases.

Note: In this section, each PFCP session uniquely maps to a subscriber forwarding state and “subscriber forwarding state” is hereinafter referred to as simply “session”.

6.1 PFCP messages

The following is a brief introduction for common PFCP messages used by the DBNG. For the complete list of PFCP messages and their details, please refer to 3GPP TS 29.244 [28].

All 3GPP PFCP IEs based in TR-459 [23] are based on 3GPP release 17 PFCP IEs from TS 29.244 [28].

6.1.1 PFCP node messages

PFCP node message includes:

- Association Setup: Used to signal node level information such as capabilities.
- Association Update: Used to signal a change of node level information, e.g., due to reconfiguration or an upgrade.
- Association Release: Used to remove a node from the DBNG function.
- Heartbeat: Used to detect unexpected failures.
- Node Report: Used to report information from the DBNG-UP to the DBNG-CP.

3GPP 29.244 [28] Table 7.3.1 provides the standard PFCP node messages. The following table identifies mandatory PFCP node messages that MUST be implemented to support fixed wireline use case covered in TR-459. In the table, “Y” indicates this PFCP message MUST be implemented, “N” indicates that vendor may choose to implement for experimental purpose and is not required by TR-459, and “NA” indicates these messages are not applicable and MUST be left alone to comply with 3GPP.

Table 6: Mandatory 3GPP PFCP Node messages for TR-459

| Message Type value (Decimal) | Message | Mandatory |
|---------------------------------|---------|-----------|
| | | SCi |

| | | |
|------------|-------------------------------------|----|
| 0 | Reserved | NA |
| | PFCP Node related messages | |
| 1 | PFCP Heartbeat Request | Y |
| 2 | PFCP Heartbeat Response | Y |
| 3 | PFCP PFD Management Request | N |
| 4 | PFCP PFD Management Response | N |
| 5 | PFCP Association Setup Request | Y |
| 6 | PFCP Association Setup Response | Y |
| 7 | PFCP Association Update Request | Y |
| 8 | PFCP Association Update Response | Y |
| 9 | PFCP Association Release Request | Y |
| 10 | PFCP Association Release Response | Y |
| 11 | PFCP Version Not Supported Response | Y |
| 12 | PFCP Node Report Request | Y |
| 13 | PFCP Node Report Response | Y |
| 14 | PFCP Session Set Deletion Request | N |
| 15 | PFCP Session Set Deletion Response | N |
| 16 to 49 | For future use | NA |
| 100 to 255 | For future use | NA |

High level overview of PFCP Node messages

The following section provide a high-level overview of mandatory PFCP Node messages which are relevant to TR-459.

PFCP Heartbeat Messages

The use of the PFCP heartbeat procedure is to detect the health status of the DBNG-CP from the DBNG-UP perspective and vice versa. The procedure is also used to detect the health status of the SCi.

PFCP Association Messages

PFCP Association message allows the DBNG-CP and DBNG-UP to form an association.

The PFCP Association Setup message is used to signal node level information such as capabilities.

The PFCP Association update message allows the DBNG-UP to inform the DBNG-CP of updates in its capability and vice versa. For example, a hardware modification on the DBNG-UP will trigger a PFCP Association update message from the DBNG-UP to the DBNG-CP informing its new capability.

The PFCP Association Release message is used to remove a DBNG-UP node from the DBNG function.

The PFCP Association update message has been enhanced to allow the DBNG-CP to program onto the DBNG-Ups subscriber groups and prefixes, with their associated attributes.

PFCP Version Not Supported Message

PFCP version not supported response message is included only in PFCP Association setup procedure if the DBNG-CP notices that the DBNG-UP is using an unsupported PFCP version.

Table 7: 3GPP PFCP Node Messages IEs Applicable to TR-459

| PFCP message type | PFCP IEs | Source Reference |
|--------------------------------|-----------------------------------|-----------------------|
| PFCP Association Setup Request | Node ID | 3GPP TS 29.244 8.2.38 |
| | Recovery Time Stamp | 3GPP TS 29.244 8.2.65 |
| | UP Function Features ¹ | 3GPP TS 29.244 8.2.25 |

| | | | | |
|--|---------------------------------------|--------------------------|---|---|
| | CP Function Features | | 3GPP TS 29.244. 8.2.58 | |
| | BBF UP Function Features ² | | TR-459 6.9.1 | |
| | Maximum ACL Chain Length | | TR-459 6.9.24 | |
| PFCP Association Setup Response | Node ID | | 3GPP TS 29.244 8.2.38 | |
| | Cause | | 3GPP TS 29.244 8.2.1 | |
| | UP Function Features ¹ | | 3GPP TS 29.244 8.2.25 | |
| | CP Function Features | | 3GPP TS 29.244. 8.2.58 | |
| | BBF UP Function Features ² | | TR-459 6.9.1 | |
| | Maximum ACL Chain Length | | TR-459 6.9.24 | |
| | Recovery Time Stamp | | 3GPP TS 29.244. 8.2.65 | |
| PFCP Association Update Request | Node ID | | 3GPP TS 29.244 8.2.38 | |
| | UP Function Features ¹ | | 3GPP TS 29.244 8.2.25 | |
| | CP Function Features | | 3GPP TS 29.244. 8.2.58 | |
| | BBF UP Function Features ³ | | TR-459 6.9.1 | |
| | BBF-Node-Info create (grouped) | BBF SGRP (grouped) | BBF SGRP ID BBF SGRP State BBF virtual MAC BBF Logical Port BBF SGRP Flags | TR-459 6.9.15 TR-459 6.9.16 3GPP TS 29.244 8.2.93 TR-459 6.9.2 TR-459 6.9.17 |
| | | BBF UP Subscriber Prefix | BBF SGRP ID BBF Ipv4 prefix BBF Ipv6 prefix Network Instance BBF Active Prefix Tag BBF Backup Prefix Tag | TR-459 6.9.15 TR-459 6.9.19 TR-459 6.9.20 3GPP TS 29.244 8.2.4 TR-459 6.9.21 TR-459 6.9.21 |
| | BBF-Node-Info modify (grouped) | BBF SGRP (grouped) | BBF SGRP ID BBF SGRP State BBF SGRP Flags | TR-459 6.9.15 TR-459 6.9.16 TR-459 6.9.17 |
| | | BBF UP Subscriber Prefix | BBF SGRP ID BBF Ipv4 prefix BBF Ipv6 prefix Network Instance BBF Active Prefix Tag BBF Backup Prefix Tag | TR-459 6.9.15 TR-459 6.9.19 TR-459 6.9.20 3GPP TS 29.244 8.2.4 TR-459 6.9.21 TR-459 6.9.21 |
| | BBF-Node-Info delete (grouped) | BBF SGRP (grouped) | BBF SGRP ID | TR-459 6.9.15 |
| | | BBF UP Subscriber Prefix | BBF SGRP ID BBF Ipv4 prefix BBF Ipv6 prefix Network Instance | TR-459 6.9.15 TR-459 6.9.19 TR-459 6.9.20 3GPP TS 29.244 8.2.4 |
| | Maximum ACL Chain Length | | TR-459 6.9.24 | |
| PFCP | Node ID | | 3GPP TS 29.244 8.2.38 | |

| | | | | |
|--|--|--------------------------------|--------------------------|-----------------------|
| Association Update Response | Cause | | 3GPP TS 29.244 8.2.1 | |
| | UP Function Features ¹ | | 3GPP TS 29.244 8.2.25 | |
| | CP Function Features | | 3GPP TS 29.244. 8.2.58 | |
| | BBF UP Function Features ² | | TR-459 6.9.1 | |
| | BBF SGRP Notification (grouped) | BBF SGRP ID | | TR-459 6.9.15 |
| | | BBF SGRP Operational Condition | | TR-459 6.9.18 |
| | | BBF SGRP Error (grouped) | BBF SGRP Error Code | TR-459 6.9.22 |
| | | | BBF SGRP Error Message | TR-459 6.9.23 |
| | | BBF Prefix Error (grouped) | BBF Prefix Error Code | TR-459 6.9.22 |
| | | | BBF Prefix Error Message | TR-459 6.9.23 |
| | | | BBF Ipv4 prefix | TR-459 6.9.19 |
| BBF Ipv6 prefix | | | TR-459 6.9.20 | |
| Network Instance | | 3GPP TS 29.244 8.2.4 | | |
| PFCP Association Release Request | Node ID | | 3GPP TS 29.244 8.2.38 | |
| PFCP Association Release Response | Node ID | | 3GPP TS 29.244 8.2.38 | |
| | Cause | | 3GPP TS 29.244 8.2.1 | |
| PFCP Version Not Supported Response | N/A ⁴ | | 3GPP TS 29.244 7.4.4.7 | |
| Heartbeat Request | Recovery Time Stamp | | 3GPP TS 29.244 8.2.65 | |
| Heartbeat Response | Recovery Time Stamp | | 3GPP TS 29.244 8.2.65 | |
| PFCP Node Report Request | Node ID | | 3GPP TS 29.244 8.2.38 | |
| | Node Report Type | | 3GPP TS 29.244 8.2.69 | |
| | Vendor-Specific Node Report Type | | 3GPP TS 29.244 8.2.217 | |
| | BBF Logical Port Report (Grouped) | Logical Port | | TR-459 6.9.2 |
| | | Forwarding Capability | | TR-459 6.9.25 |
| | | MAC Address | | 3GPP TS 29.244 8.2.93 |
| | | Usage | | 3GPP TS 29.244 8.2.34 |
| | BBF SGRP Notification Report (Grouped) | BBF SGRP ID | | TR-459 6.9.15 |
| | | BBF SGRP Operational Condition | | TR-459 6.9.18 |
| | | BBF SGRP Error | BBF SGRP Error Code | TR-459 6.9.22 |
| | | | BBF SGRP Error Message | TR-459 6.9.23 |

| | | | |
|----------------------------------|---------------------------------------|---------------------|-----------------------|
| | | Message | |
| | BBF Network Instance Report (Grouped) | Network Instance | 3GPP TS 29.244 8.2.4 |
| | | Connectivity Status | TR-459 6.9.26 |
| PFCP Node Report Response | Node ID | | 3GPP TS 29.244 8.2.38 |
| | Cause | | 3GPP TS 29.244 8.2.1 |
| | Offending IE | | 3GPP TS 29.244 8.2.22 |

[R-167] ¹UP Function Features MUST be sent when PFCP Association message is sent from the DBNG-UP.

[R-168] ²BBF UP Function Features MUST be sent when PFCP Association message is sent from the DBNG-UP.

[R-169] ³BBF UP Function Features MUST be sent if and only if DBNG-UP has at least one DBNG-UP function to update.

⁴Conditional as referred to in 3GPP 29.244 [28] Section 7.2.3.2.

6.1.2 PFCP session messages

3GPP 29.244 [28] Table 7.3.1 provides the standard PFCP session messages. The following table identifies mandatory PFCP session messages that MUST be implemented to support fixed wireline use case covered in TR-459. In the table, “Y” indicates this PFCP message MUST be implemented.

Table 8: Mandatory 3GPP PFCP Session messages for TR-459

| Message Type value (Decimal) | Message | Mandatory |
|------------------------------|--------------------------------------|-----------|
| | | SCi |
| | PFCP Session related messages | |
| 50 | PFCP Session Establishment Request | Y |
| 51 | PFCP Session Establishment Response | Y |
| 52 | PFCP Session Modification Request | Y |
| 53 | PFCP Session Modification Response | Y |
| 54 | PFCP Session Deletion Request | Y |
| 55 | PFCP Session Deletion Response | Y |
| 56 | PFCP Session Report Request | Y |
| 57 | PFCP Session Report Response | Y |

PFCP Session Messages

PFCP session messages are divided into the following 4 message types:

- **Session Establishment:** programs subscriber forwarding state, typically used when a subscriber initiates a connection to the DBNG. The message allows the DBNG-CP to program subscriber forwarding rules to the DBNG-UP. Each subscriber in wireline access as specified in section 6.2 requires four forwarding rules: two rules for bi-directional control message forwarding and two rules for bi-directional data message forwarding. Each subscriber has its own PFCP session.
- **Session Modification:** updates subscriber forwarding state. The message allows the DBNG-CP to modify subscriber forwarding rules existing on the DBNG-UP. This is typically needed when Immediate Session Creation option is adopted to complete the set of rules for a PFCP session, adding new PDR and FAR rules for data traffic. Another typical use of this message is when RADIUS CoA is received and DBNG-CP needs to update subscriber attributes upon receiving a new policy from AAA. Modification is typically triggered but not limited to changing subscriber QoS policies, filtering policies, forwarding actions. Modification also supports DBNG-CP querying the DBNG-UP for statistics.

- Session Deletion: removes a subscriber forwarding state, typically used when a subscriber terminates the broadband session or logs off the network. The message allows the DBNG-CP to delete the subscriber session on the DBNG-UP.
- Session Report: reports information about the session, allows the DBNG-UP to report subscriber session information to the DBNG-CP.

Within Session Establishment, Session Modification, and Session Deletion, rules are used to program the subscriber forwarding state:

- Packet Detection Rule (PDR) is a rule that contains a selection of the objects below
 - Packet Detection Identifier (PDI) specifies the matching criteria for packets
 - Forward action Rule (FAR) specifies the action (e.g., forward/drop/mirror) to be taken based on the matching PDI.
- QoS Enforcement Rule (QER) specifies the QoS treatment based on the matching PDI.
- Usage Reporting Rule (URR) specifies the usage reporting and charging rule based on the matching PDI.

The following tables list both 3GPP and BBF IEs that are used in context of DBNG. Most IEs are supported as-is by a compliant implementation. As an exception, certain IEs are only supported if a specific BBF UP Function Feature Flag is negotiated using a PFCP Association message. The list of these IEs and their mapping to the specific Feature Flag is defined in Table 34.

Table 9: 3GPP PFCP Session Establishment Messages IEs Applicable to TR-459

| PFCP message type | PFCP Grouped IEs and IEs | | Source Reference | |
|------------------------------------|--------------------------------|-----------------------|------------------------|-----------------------|
| PFCP Session Establishment Request | Node ID | | 3GPP TS 29.244 8.2.38 | |
| | CP F-SEID | | 3GPP TS 29.244 8.2.1 | |
| | PPP LCP Connectivity (Grouped) | Traffic Endpoint ID | 3GPP TS 29.244 8.2.92 | |
| | | Verification Timers | TR-459 6.9.7 | |
| | | PPP LCP Magic Number | TR-459 6.9.8 | |
| | BBF ACL (Grouped) | BBF ACL Name | 3GPP TS 29.244 8.2.72 | |
| | | BBF ACL Direction | TR-459 6.9.13 | |
| | | BBF ACL Family | TR-459 6.9.14 | |
| | Create PDR (Grouped) | Precedence | | 3GPP TS 29.244 8.2.11 |
| | | PDR ID | | 3GPP TS 29.244 8.2.36 |
| | | PDI (Grouped) | Source Interface | 3GPP TS 29.244 8.2.2 |
| | | | Traffic Endpoint ID | 3GPP TS 29.244 8.2.92 |
| | | | SDF Filter | 3GPP TS 29.244 8.2.5 |
| | | | Ethernet Packet Filter | 3GPP TS 29.244 8.2.98 |
| | | | L2TP type | TR-459 6.9.12 |
| Outer Header Removal | | 3GPP TS 29.244 8.2.64 | | |
| BBF Outer Header Removal | | TR-459 6.9.4 | | |
| FAR ID | | 3GPP TS 29.244 8.2.74 | | |

| | | | | |
|-----------------------------------|---------------------------------|----------------------------|------------------------|-----------------------|
| Create FAR (Grouped) | FAR ID | | 3GPP TS 29.244 8.2.74 | |
| | Apply Action | | 3GPP TS 29.244 8.2.26 | |
| | Forwarding Parameters (Grouped) | Destination interface | | 3GPP TS 29.244 8.2.24 |
| | | Outer header creation | | 3GPP TS 29.244 8.2.56 |
| | | Linked Traffic Endpoint ID | | 3GPP TS 29.244 8.2.92 |
| | | BBF Outer Header Creation | | TR-459 6.9.3 |
| | | Network instance | | 3GPP TS 29.244 8.2.4 |
| | | MTU | | TR-459 6.9.9 |
| Create Traffic Endpoint (Grouped) | Traffic Endpoint ID | | 3GPP TS 29.244 8.2.92 | |
| | Local F-TEID | | 3GPP TS 29.244 8.2.3 | |
| | Network instance | | 3GPP TS 29.244 8.2.4 | |
| | UE IP address | | 3GPP TS 29.244 8.2.62 | |
| | Framed-Route | | 3GPP TS 29.244 8.2.109 | |
| | Framed-Ipv6-Route | | 3GPP TS 29.244 8.2.111 | |
| | MAC address | | 3GPP TS 29.244 8.2.93 | |
| | C-TAG | | 3GPP TS 29.244 8.2.94 | |
| | S-TAG | | 3GPP TS 29.244 8.2.95 | |
| | Logical Port | | TR-459 6.9.2 | |
| | PPPoE Session ID | | TR-459 6.9.5 | |
| | L2TP tunnel (Grouped) | L2TP tunnel endpoint | | TR-459 6.9.10 |
| | | L2TP Session ID | | TR-459 6.9.11 |
| | Create URR (Grouped) | URR ID | | 3GPP TS 29.244 8.2.54 |
| Measurement Method | | 3GPP TS 29.244 8.2.40 | | |
| Reporting Triggers | | 3GPP TS 29.244 8.2.19 | | |
| Measurement Period | | 3GPP TS 29.244 8.2.42 | | |
| Volume Threshold | | 3GPP TS 29.244 8.2.13 | | |
| Volume Quota | | 3GPP TS 29.244 8.2.50 | | |
| PDN Type | | | 3GPP TS 29.244 8.2.79 | |
| BBF SGRP ID | | | TR-459 6.9.15 | |

Table 10: 3GPP PFCP Session Establishment Response Messages IEs Applicable to TR-459

| PFCP message type | PFCP Grouped IEs and IEs | Source Reference |
|-------------------|--------------------------|------------------|
|-------------------|--------------------------|------------------|

| | | | |
|--|------------------------------------|---------------------|-----------------------|
| PFCP Session Establishment Response | Node ID | | 3GPP TS 29.244 8.2.38 |
| | Cause | | 3GPP TS 29.244 8.2.1 |
| | Offending IE | | 3GPP TS 29.244 8.2.22 |
| | BBF Error Code | | TR-459 6.9.22 |
| | UP F-SEID | | 3GPP TS 29.244 8.2.37 |
| | Created Traffic Endpoint (Grouped) | Traffic Endpoint ID | 3GPP TS 29.244 8.2.92 |
| | | Local F-TEID | 3GPP TS 29.244 8.2.3 |

Table 11: 3GPP PFCP Session Modification Request Messages IEs Applicable to TR-459

| PFCP message type | PFCP Grouped IEs and IEs | | Source Reference | |
|--|-----------------------------------|---|--------------------------------|-----------------------|
| PFCP Session Modification Request | CP F-SEID | | 3GPP TS 29.244 8.2.1 | |
| | Remove PDR (Grouped) | PDR ID | 3GPP TS 29.244 8.2.36 | |
| | Remove FAR (Grouped) | FAR ID | 3GPP TS 29.244 8.2.74 | |
| | Remove Traffic Endpoint (Grouped) | Traffic Endpoint ID | 3GPP TS 29.244 8.2.92 | |
| | PPP LCP Connectivity (Grouped) | Traffic Endpoint ID | 3GPP TS 29.244 8.2.92 | |
| | | Verification Timers | TR-459 6.9.7 | |
| | | PPP LCP Magic Number | TR-459 6.9.8 | |
| | BBF ACL (Grouped) | BBF ACL Name BBF ACL Direction BBF ACL Family | TR-459 6.9.13 TR-459 6.9.14 | |
| | Create PDR (Grouped) | Precedence | | 3GPP TS 29.244 8.2.65 |
| | | PDR ID | | 3GPP TS 29.244 8.2.36 |
| | | PDI (grouped) | Source Interface | 3GPP TS 29.244 8.2.2 |
| | | | Traffic Endpoint ID | 3GPP TS 29.244 8.2.92 |
| | | | SDF Filter | 3GPP TS 29.244 8.2.5 |
| | | | Ethernet Packet Filter | 3GPP TS 29.244 8.2.98 |
| | | | L2TP type | TR-459 6.9.12 |
| | | Outer Header Removal | | 3GPP TS 29.244 8.2.64 |
| | | BBF Outer Header Removal | | TR-459 6.9.4 |
| | FAR ID | | 3GPP TS 29.244 8.2.74 | |
| | Create FAR (grouped) | FAR ID | | 3GPP TS 29.244 8.2.74 |
| | | Apply Action | | 3GPP TS 29.244 8.2.26 |
| Forwarding Parameters (grouped) | | Destination interface | 3GPP TS 29.244 8.2.24 | |
| | | Outer header creation | 3GPP TS 29.244 8.2.56 | |
| | | Linked Traffic Endpoint ID | 3GPP TS 29.244 8.2.92 | |

| | | | |
|-----------------------------------|--|---------------------------|------------------------|
| | | BBF Outer Header Creation | TR-459 6.9.3 |
| | | Network instance | 3GPP TS 29.244 8.2.4 |
| | | MTU | TR-459 6.9.9 |
| Create Traffic Endpoint (Grouped) | Traffic Endpoint ID | | 3GPP TS 29.244 8.2.92 |
| | Local F-TEID | | 3GPP TS 29.244 8.2.3 |
| | Network instance | | 3GPP TS 29.244 8.2.4 |
| | UE IP address | | 3GPP TS 29.244 8.2.62 |
| | Framed-Route | | 3GPP TS 29.244 8.2.109 |
| | Framed-Ipv6-Route | | 3GPP TS 29.244 8.2.111 |
| | MAC address | | 3GPP TS 29.244 8.2.93 |
| | C-TAG | | 3GPP TS 29.244 8.2.94 |
| | S-TAG | | 3GPP TS 29.244 8.2.95 |
| | Logical Port | | TR-459 6.9.2 |
| | PPPoE Session ID | | TR-459 6.9.5 |
| | L2TP tunnel (Grouped) | L2TP tunnel endpoint | |
| L2TP Session ID | | TR-459 6.9.11 | |
| Update PDR (Grouped) | Precedence | | 3GPP TS 29.244 8.2.65 |
| | PDR ID | | 3GPP TS 29.244 8.2.36 |
| | SDF Filter | | 3GPP TS 29.244 8.2.5 |
| | Ethernet Packet Filter | | 3GPP TS 29.244 8.2.98 |
| Update FAR (grouped) | FAR ID | | 3GPP TS 29.244 8.2.74 |
| | Apply Action | | 3GPP TS 29.244 8.2.26 |
| | Update Forwarding Parameters (grouped) | MTU | TR-459 6.9.9 |
| Update Traffic Endpoint (grouped) | Traffic Endpoint ID | | 3GPP TS 29.244 8.2.92 |
| | UE IP address | | 3GPP TS 29.244 8.2.62 |
| | Framed-Route | | 3GPP TS 29.244 8.2.109 |
| | Framed-Ipv6-Route | | 3GPP TS 29.244 8.2.111 |
| Create URR (Grouped) | URR ID | | 3GPP TS 29.244 8.2.54 |
| | Measurement Method | | 3GPP TS 29.244 8.2.40 |
| | Reporting Triggers | | 3GPP TS 29.244 8.2.19 |
| | Measurement Period | | 3GPP TS 29.244 8.2.42 |

| | | | |
|--|------------|------------------|---------------------------|
| | | Volume Threshold | 3GPP TS 29.244 8.2.13 |
| | | Volume Quota | 3GPP TS 29.244 8.2.50 |
| | Query URR | | 3GPP TS 29.244 7.5.4.10-1 |
| | Remove URR | | 3GPP TS 29.244 7.5.4.8 |

Table 12: 3GPP PFCP Session Modification Response Messages IEs Applicable to TR-459

| PFCP message type | PFCP Grouped IEs and IEs | | Source Reference |
|---|------------------------------------|-----------------------|-----------------------|
| PFCP Session Modification Response | Cause | | 3GPP TS 29.244 8.2.1 |
| | Offending IE | | 3GPP TS 29.244 8.2.22 |
| | UP F-SEID | | 3GPP TS 29.244 8.2.37 |
| | BBF Error Code | | TR-459 6.9.22 |
| | Created Traffic Endpoint (grouped) | Traffic Endpoint ID | 3GPP TS 29.244 8.2.92 |
| | Usage Report | BBF Event Time Stamp* | TR-459.2 6.5.9 |

Note*: This IE shall provide the timestamp when the collection of the information in this report was performed by the DBNG-UP.
 - If DBNG-UP does not provide this information, the DBNG-CP can use the timestamp when the PFCP message was received from the DBNG-UP.

Table 13: 3GPP PFCP Session Deletion Request Messages IEs Applicable to TR-459

| PFCP message type | PFCP Grouped IEs and IEs | | Source Reference |
|--------------------------------------|--------------------------|--|------------------|
| PFCP Session Deletion Request | | | |

Table 14: 3GPP PFCP Session Deletion Response Messages IEs Applicable to TR-459

| PFCP message type | PFCP Grouped IEs and IEs | | Source Reference |
|---------------------------------------|--|-----------------------|------------------------|
| PFCP Session Deletion Response | Cause | | 3GPP TS 29.244 8.2.1 |
| | Offending IE | | 3GPP TS 29.244 8.2.22 |
| | Usage Report | BBF Event Time Stamp* | TR-459.2 6.5.9 |
| | Additional Usage Reports Information** | | 3GPP TS 29.244 8.2.9.1 |

Note*: This IE shall provide the timestamp when the collection of the information in this report was performed by the DBNG-UP.
 If DBNG-UP does not provide this information, the DBNG-CP can use the timestamp when the PFCP message was received from the DBNG-UP.

Note**: Cause is set to “More Usage Report to send” when this IE is used.

Table 15: 3GPP PFCP Session Report Request IEs Applicable to TR-459

| PFCP message type | PFCP Grouped IEs and IEs | P | Source Reference | |
|-----------------------------|--------------------------|------------------------|-----------------------|-----------------------|
| PFCP Session Report Request | Report Type | M | 3GPP TS 29.244 8.2.21 | |
| | Usage Report | URR ID | M | 3GPP TS 29.244 8.2.54 |
| | | UR-SEQN | M | 3GPP TS 29.244 8.2.71 |
| | | Usage Report Trigger | M | 3GPP TS 29.244 8.2.41 |
| | | Volume Measurement | C | 3GPP TS 29.244 8.2.44 |
| | | BBF Event Time Stamp* | O | TR-459.2 (IE 32792) |
| PFCP SRReq-Flags** | C | 3GPP TS 29.244 8.2.119 | | |

[R-170] Below are the mandatory octet bits that MUST be supported for the PFCP IE Usage Report Trigger:
Octet 5:

Bit 1 – PERIO (Periodic Reporting): when set to “1”, this indicates a periodic report.

Bit 8 – IMMER (Immediate Report): when set to “1”, this indicates an immediate report reported on CP function demand.

Octet 6:

Bit 4 – TERM (Termination Report): when set to “1”, this indicates a usage report being reported (in a PFCP Session Deletion Response) for a URR due to the termination of the PFCP session, or a usage report being reported (in a PFCP Session Modification Response) for a URR due to the removal of the URR or dissociated from the last PDR.

Note*: This IE shall provide the timestamp when the collection of the information in this report was performed by the DBNG-UP.

If DBNG-UP does not provide this information, the DBNG-CP can use the timestamp when the PFCP message was received from the DBNG-UP.

Note**: Only the final Session Report Request would have PFCP SRReq-Flags IE with the PSDBU flag set to 1. Please refer to 3GPP TS 29.244 [28] section 8.2.119 for more information on IE PFCP SRReq-Flags.

Table 16: 3GPP PFCP Session Report Response IEs Applicable to TR-459

| PFCP message type | PFCP Grouped IEs and IEs | P | Source Reference |
|------------------------------|--------------------------|---|-----------------------|
| PFCP Session Report Response | Cause* | M | 3GPP TS 29.244 8.2.1 |
| | Offending IE | C | 3GPP TS 29.244 8.2.22 |

Note*: Cause can be set to “More Usage Report to send”.

6.1.3 PFCP information elements

Information Elements are encoded as TLVs. Each PFCP session may use individual IEs or grouped IEs (IE that contains other IEs) for DBNG-CP and DBNG-UP communication.

TR-459 [23] re-uses IEs defined in 3GPP TS 29.244 [28]. These IEs are all listed in the applicability tables in sections 6.1.1 and 6.1.2. Some of these IEs are dependent on UP Function Feature flags in 3GPP, but mandatory for the DBNG application.

All of the following 3GPP UP Function Feature flags MUST be set:

- FTUP for GTP-U tunneling
- FRRT for framed routes
- PDIU for traffic endpoint support

Other 3GPP UP Function Feature Flags, or CP Function Feature Flags as defined in 3GPP TS 29.244 [28] MAY be supported by the DBNG-UP and DBNG-CP, but they are all considered optional.

Some 3GPP IEs are supported only if a specific BBF UP Function Feature flag is set as listed in Table 33. Some IEs may contain a wide array of flags (e.g., to specify reporting triggers), many of which are not applicable to the DBNG application. Table 17 lists the flags that MUST be supported per IE.

Table 17 summarizes the list of 3GPP IEs from TS 29.244 [28] that MUST be supported by a DBNG-UP for TR-459 [23] compliance.

Table 17: 3GPP PFCP IEs that must be supported for TR-459

| IE Type value (Decimal) | Information elements |
|-------------------------|--|
| 0 | Reserved |
| 1 | Create PDR |
| 2 | PDI |
| 3 | Create FAR |
| 4 | Forwarding Parameters |
| 6 | Create URR |
| 7 | Create QER |
| 9 | Update PDR |
| 10 | Update FAR |
| 11 | Update Forwarding Parameters |
| 13 | Update URR |
| 14 | Update QER |
| 15 | Remove PDR |
| 16 | Remove FAR |
| 17 | Remove URR |
| 18 | Remove QER |
| 19 | Cause |
| 20 | Source Interface |
| 21 | F-TEID |
| 22 | Network Instance |
| 23 | SDF Filter |
| 26 | MBR |
| 27 | GBR |
| 28 | QER Correlation ID |
| 29 | Precedence |
| 31 | Volume Threshold |
| 33 | Monitoring Time |
| 34 | Subsequent Volume Threshold |
| 37 | *Reporting Triggers |
| 39 | Report Type |
| 40 | Offending IE |
| 41 | Forwarding Policy |
| 42 | Destination Interface |
| 43 | UP Function Features |
| 44 | Apply Action |
| 51 | **Load Control Information |
| 52 | **Sequence Number |
| 53 | **Metric |
| 54 | **Overload Control Information |
| 55 | **Timer |
| 56 | PDR ID |
| 57 | F-SEID |
| 60 | Node ID |
| 62 | *Measurement Method |
| 63 | *Usage Report Trigger |
| 64 | Measurement Period |
| 66 | Volume Measurement |
| 71 | Quota Holding Time |
| 73 | Volume Quota |
| 77 | Query URR |
| 78 | Usage Report (Session Modification Response) |
| 79 | Usage Report (Session |

| | |
|-----|---------------------------------------|
| | Deletion Response) |
| 80 | Usage Report (Session Report Request) |
| 81 | URR ID |
| 82 | Linked URR ID |
| 84 | Outer Header Creation |
| 89 | CP Function Features |
| 90 | Usage Information |
| 93 | UE IP Address |
| 94 | Packet Rate |
| 95 | Outer Header Removal |
| 96 | Recovery Time Stamp |
| 100 | *Measurement Information |
| 104 | UR-SEQN |
| 106 | Activate Predefined Rules |
| 107 | Deactivate Predefined Rules |
| 108 | FAR ID |
| 109 | QER ID |
| 110 | **OCI Flags |
| 111 | PFCP Association Release Request |
| 112 | Graceful Release Period |
| 113 | PDN Type |
| 114 | Failed Rule ID |
| 117 | User Plane Inactivity Timer |
| 121 | Subsequent Volume Quota |
| 125 | Query URR Reference |
| 126 | Additional Usage Reports Information |
| 127 | Create Traffic Endpoint |
| 128 | Created Traffic Endpoint |
| 129 | Update Traffic Endpoint |
| 130 | Remove Traffic Endpoint |
| 131 | Traffic Endpoint ID |
| 132 | Ethernet Packet Filter |
| 133 | MAC address |
| 134 | C-TAG |
| 135 | S-TAG |
| 136 | Ethertype |
| 147 | Additional Monitoring Time |
| 153 | Framed-Route |
| 155 | Framed-Ipv6-Route |

Notes:

- IEs with “*” have selected FLAGS, identified in Table 18, which MUST be supported.
- IEs with “**” are conditional based on CP function feature flag, specified in Table 19.

The following flags within the 3GPP PFCP IEs MUST be supported.

Table 18: 3GPP PFCP IE flags that must be supported for TR-459

| IE Type value (Decimal) | Information elements | Mandatory support flags | 3GPP UP function feature flag | BBF UP function feature flag |
|-------------------------|----------------------|-------------------------|-------------------------------|------------------------------|
| 37 | Reporting Triggers | PERIO | | |
| | | VOLTH | | Volume-threshold-reporting |

| | | | | |
|-----|-------------------------|-------|------|----------------------------|
| | | QUHTI | | Volume-quota |
| | | LIUSA | | |
| | | VOLQU | | Volume-quota |
| | | IPMJL | IPTV | |
| | | QUVTI | | Volume-quota |
| 62 | Measurement Method | VOLUM | | |
| 63 | Usage Report Trigger | PERIO | | |
| | | VOLTH | | Volume-threshold-reporting |
| | | QUHTI | | Volume-quota |
| | | IMMER | | |
| | | VOLQU | | Volume-quota |
| | | LIUSA | | |
| | | TERMR | | |
| | | IPMJL | IPTV | |
| | | QUVTI | | Volume-quota |
| 100 | Measurement Information | MNOP | MNOP | |

The following 3GPP PFCP IEs are to be supported based on the 3GPP CP function feature flag.

Table 19: 3GPP PFCP IE that must be supported for TR-459 and are conditional based on the 3GPP CP function features flag

| IE Type value (Decimal) | Information elements | 3GPP CP function features flag |
|-------------------------|------------------------------|--------------------------------|
| 51 | Load Control Information | LOAD |
| 53 | Metric | LOAD, OVRL |
| 54 | Overload Control Information | OVRL |
| 55 | Timer | LOAD, OVRL |
| 110 | OCI Flags | LOAD, OVRL |

6.2 PDR Matching Rules

A DBNG CUPS system follows the basic packet matching model as described in clause 5.2 of 3GPP TS 29.244 [28]. This section defines extensions to that forwarding model that are for example required due to protocol extensions defined in this document (e.g., SGRP) or due to BNG specific use cases that are not present in 29.244 [28](e.g., the CPRi interface).

6.2.1 Per-session, per-logical-port, and default CPRi matching precedence

A DBNG CUPS system follows the basic packet matching model as described in clause 5.2 of 3GPP TS 29.244 [28]. A few well-defined exceptions to that model are defined in this section.

An important aspect in this model is that PDR matching rules between PFCP sessions in general do not overlap, with some well-defined exceptions. This is because there is a two-step lookup mechanism where first the full PFCP session is matched directly and subsequently a more detailed PDR within that PFCP session is matched using a precedence list.

The DBNG adds an exception to this rule for the purpose of CPR tunnels, which will overlap in the uplink direction:

1. The default CPR tunnel's PDRs will overlap with any per-logical-port tunnel PDR and with any per-session CPR PDR.
2. The per-logical-port tunnel PDR will overlap with the per-session CPR PDR of any session on that logical port.

This overlap is a fundamental property of the DBNG solution and is handled using the following fall-through logic. In general, it matches on the most specific first (i.e., session level) to the least specific (i.e., default CPRi).

- Upon receiving a packet, a DBNG-UP will first try to match any non-CPR PFCP session and any of its PDRs as defined in 3GPP TS 29.244 [28]. If a session and PDR match is found, the associated action is applied. If the session is found, but no PDR match is found, the packet is dropped.
 - In scope of this document all PFCP session identifying attributes are assumed to be part of the 'traffic endpoint' IE, and never part of the 'PDI' IE. For example, a logical port, MAC address, and PPPoE session ID are session identifying attributes, but an SDF filter is not.
- Else, if no session match is found, the DBNG-UP will try to match any per-logical-port CPR session by virtue of the logical port. If a per-logical-port CPR session is found and any of its PDRs match, the associated action is applied. If a per-logical-port CPR session is found, but no PDR match is found, the packet is dropped.
- Else, the DBNG-UP will try to match the packet against the PDR rules of the default CPR session if such a session is signaled. If any PDR matches, the associated action is applied.
- Else, the packet is dropped.

It is important in the above logic that as soon as a packet matches a session based on the traffic endpoint, there is no fall-through and the packet either MUST match any PDR and follow the rules of the linked FAR, or is dropped otherwise.

Beneath we give a few DHCP packet matching examples for IPoE, supposing the following PFCP sessions are installed on a DBNG-UP. We focus on the traffic endpoint parameters and assume appropriate Ethernet Filter/SDF rules for IPoE CPRi Tunnel rules are present for each session.

1. An IPoE PFCP session on logical port 'A' with S-Tag 10, C-Tag 20, and MAC 00-00-5E-00-53-01.
2. A per-logical port CPRi PFCP session on logical port 'A', with an additional Ethernet Packet Filter matching S-tag range 1 to 1024.

3. A per-logical port CPRi PFCP session on logical port 'B', which drops all traffic.
4. A default CPRi PFCP session.

The DBNG-UP will handle an incoming DHCP CPRi packet as follows, based on the set header fields:

- Logical port 'A', S-tag 10, C-tag 20, MAC 00-00-5E-00-53-01: Matches the IPoE session (1) and will be forwarded using the per-session CPRi tunnel to the DBNG-CP.
- Logical port 'A', S-tag 10, C-tag 20, MAC 00-00-5E-00-53-02: Matches the per-logical-port CPRi session for 'A' (2), and within that session matches the S-Tag range. Will be forwarded using the per-logical-port CPRi tunnel to the DBNG-CP.
- Logical port 'A', S-tag 2000, C-tag 20, MAC 00-00-5E-00-53-03: Matches the per-logical-port CPRi session for 'A' (2), but does not match the S-Tag range. Because no other PDR within the PFCP session is defined, the packet is dropped.
- Logical port 'B', S-tag 10, C-tag 20, MAC 00-00-5E-00-53-04: Matches the per-logical-port CPRi session for 'B' (3), and will be dropped as defined.
- Logical port 'C', S-tag 10, C-tag 20, MAC 00-00-5E-00-53-05: Matches the default CPRi session (4) and will be forwarded using the default CPRi Tunnel.

6.2.2 ARP and NS rules

When signaling an IPv4 or IPv6 gateway address in an SGRP, the DBNG-CP will add a PDR to an IPoE session with the following content in the PDI IE:

- For an IPv4 gateway: an Ethernet Packet Filter IE including only an Ethertype IE set to 0x0806
- For an IPv6 gateway: an Ethernet Packet Filter IE including only an Ethertype IE set to 0x86DD, and an SDF 'permit out 58 from any to any ICMP types 135'.

This PDR will point to a FAR including a BBF Apply Action IE set to 'Reply-To-Request'.

When a DBNG-UP receives an ARP Request packet or an ICMPv6 NS packet, it will first look up the PFCP session based on the session key as identified by the Traffic Endpoint IEs.

- If the matching PFCP session related to a CPRi tunnel (i.e., it has no PDN type IE set), the packet is either forwarded to the DBNG-CP or dropped depending on the PDRs for that session.
- If the matching PFCP session is a subscriber session (i.e., it has the PDN Type IE set), and it queries for an IPv4 or IPv6 gateway that is part of a linked SGRP the DBNG-UP will generate an ARP response or ICMPv6 NA for this SGRP, if the PDR mentioned above is present AND the SGRP state is active. Else, the packet is dropped.
- If the matching PFCP session is a subscriber session but the queried IPv4 or IPv6 address is not present in an SGRP, it is up to DBNG-UP policy whether to reply to this packet, even if no explicit ARP/NS matching rules are defined in the PFCP session.

When an ARP Response or ICMPv6 NA message is generated by the DBNG-UP, the advertised MAC address MUST be the SGRP MAC address if a MAC address was signaled for the SGRP linked to the UP. If no MAC address is signaled in the SGRP, the DBNG-UP uses a MAC address that is associated with the logical port, this MAC address MUST be the same as what is sent in CPRi NSH (see section *NSH header information*).

6.2.3 Unicast Traffic for Local DBNG-UP IP addresses

For any unicast protocol (e.g., ICMP(v6) echo) the UP can handle traffic locally if the following conditions are met:

- The packet matches a PFCP session PDR rule for which the action is 'forward', and the destination interface is 'core'.

Note: when a DBNG-CP explicitly asks to redirect this packet to the DBNG-CP using the CPRi interface (Destination interface = cp-function) the UP MUST NOT handle this traffic locally.

- The packet's destination IP is a local UP IP address in context of the Network Instance indicated by the FAR rule, including any /32 (IPv4) or /128 (IPv6) addresses signaled using SGRPs.
Note: when no Network Instance is present in an SGRP or FAR, a default Network Instance is assumed.

6.2.4 Priority Between Network Redirect and Per-session Rules

In case of for example DHCP relay there is a potential overlap between a generic redirection rule that captures all DHCP(v6) packets from a server, where only a 'Network Instance' is specified, and per-session rules that specify both a Network Instance and a UE IP Address.

A DBNG-UP applies the following fall-through logic for an IP packet coming from Source Interface Core, to align as much as possible with 3GPP TS 29.244 [28] clause 5.2:

- First it matches any subscriber PFCP session (with a PDN Type) based on its Traffic Endpoint IE (typically Network Instance + IP address). If any such session is found, it is matched against its PDR rules which are applied as such. If no PDR is matched the packet is dropped.
Note: if a DBNG-CP wants to intercept unicast control-plane messages such as DHCP Ack or DHCPv6 Reply, it MUST also install a per-session DHCP redirect rule in the downstream direction.
- If no subscriber PFCP session matches, it matches a PFCP redirect session (no PDN type) based on the Network Instance defined in that session and follows the PDR rules defined in that session. If no such session exists, the packet is dropped.

Note: To simplify the DBNG-UP lookup logic at most one redirect PFCP session MUST exist per Network Instance. If multiple rules are required for the same network instance (for example, DHCP, DHCPv6 and L2TP Tunnel) they are to be combined in one session.

Note: The absence of a Network Instance IE implies a match against a default network instance, not all network instances.

6.3 PFCP Connectivity Requirements

For SCi reliability, a retransmission strategy is employed which is inherent to the protocol used for that interface (see TS 29.244 [28] clause 6.4): lost request messages are sent N times at T intervals. N is a "retry-count", T is a "timeout".

For Heartbeat Requests, N and T should be configurable separately from other message types.

In this section the following definitions will be used:

- T_h : A timer value between consecutive Heartbeat Request messages. This is an integer value in units of seconds.
- N_h : The number of times a Heartbeat Request is retried (including the first message) when no response is received before concluding the peer is unresponsive. This is an integer value higher than 1.
- T_r : A timer value after which a message request (other than Heartbeat Request) is retried when there is no response. This is an integer value in unit of milliseconds.
- N_r : The number of times a message request (other than Heartbeat Request) is retried (including the first message) when no response is received before concluding the peer is unresponsive. This is an integer value equal or higher than 1.

The heartbeat mechanism is conceived to detect a reboot of the PFCP peer via the Recovery Time Stamp IE (see 3GPP TS 23.007 [24] clause 19A).

However, a series of missing Heartbeat Response to a Heartbeat Request retried N_h times spaced apart T_h may also be consequence of a relatively long connectivity fault at PFCP level.

- By “**long PFCP connectivity fault**” it is intended a fault detected via heartbeat mechanism that lasts longer than or equal to $N_h * T_h$.
- By “**short PFCP connectivity fault**” it is intended a fault, detected via the loss of any type of PFCP messages, that lasts less than $N_h * T_h$.

A long PFCP connectivity fault is considered as a non-recoverable failure which tears down the subscriber sessions and the CPR tunnels: all in-progress PFCP messages are subject to failure and request retransmission counters are reset.

If a short PFCP connectivity fault occurs, a number of PFCP messages may not be delivered to the destination. Although the loss of the PFCP messages to setup new sessions is not worrying (as it may be recovered by timers' expiration and retry mechanisms inherent to CPEs), the loss of other types of PFCP messages may cause the missing release of resources on DBNG-UP and DBNG-CP, with no possibility to automatically recover these resources from their freezing. For example, an undelivered PFCP Session Report Request, which would have informed DBNG-CP of an expired LCP keepalive, will not update the DBNG-CP about the CPE disconnection and will keep, both on DBNG-UP and DBNG-CP, reserved resources associated with the disconnected CPE. Likewise, an undelivered PFCP Session Deletion Request will fail to signal to DBNG-UP that it can free up some resources.

There aren't mechanisms to recover from the resources freezing caused by the loss of certain types of PFCP requests: therefore, this has to be avoided. So achieve this, the parameters N_h and T_h are to be configurable within some ranges and the Operators need to choose the values of N_r , T_r , N_h and T_h such that some conditions are met (see Note after [R-172]), in order to limit the consequences of the loss of those types of PFCP requests. The suggested settings compel the DBNG-UP and the DBNG-CP to keep retrying the PFCP requests until they are acknowledged or until the heartbeat mechanism detects a long PFCP connectivity fault.

Note: In the current issue of this specification (Issue 2), no distinction among the types of PFCP requests is made to identify the messages for which the total retrying interval might be shortened as their loss can be tolerated. This is left FFS.

The actions to be taken by the DBNG-UP and the DBNG-CP in case of a long connectivity fault are reported in [R-173] to [R-175].

In the following, some requirements concerning the DBNG heartbeat mechanism in addition to those specified in TS 29.244 [28] and TS 23.007 are introduced.

[R-171] PFCP Heartbeat Requests sent by DBNG-UP or DBNG-CP SHOULD be repeated N_h times every T_h seconds if not acknowledged by the PFCP Heartbeat Response. T_h SHOULD span at least in the range [5, 60] seconds, with a granularity of 1 second. N_h SHOULD be an integer spanning at least in the range [2, 5].

[R-172] The retransmission parameters T_r and N_r for PFCP request messages other than Heartbeat Requests MUST be configurable independently from the parameters N_h and T_h related to Heartbeat Requests, on both DBNG-UP or DBNG-CP.

Note: For the current issue of this specification (Issue 2), the condition

$$T_r \times (N_r - 1) = T_h \times N_h$$

is recommended with T_r at least one order of magnitude smaller than T_h (for example $T_r = 0.1 \times T_h$). It is also suggested to use the same settings for N_r , T_r , N_h , T_h on both DBNG-UP and DBNG-CP.

The procedure to derive the best settings is the following:

- Choose T_h and N_h on the basis of the maximum time the DBNG-UP is allowed to run without control.
- Choose T_r one order of magnitude smaller than T_h .
- Derive a value for N_r from the formula above.

Note: In principle, N_r might be configured larger than the value derived from the formula. However, retransmissions beyond the heartbeat failure detection are futile, while an excessively large value of N_r might have impacts on the resources to be allocated.

With T_r and N_r values so chosen, PFCP message requests are retried until a response is received or heartbeat expires. It is understood that, if the heartbeat expires, all in-progress PFCP messages are not retried any longer, even if they have been retried less than N_r times.

[R-173] If the DBNG-UP detects a DBNG-CP restart from a Recovery Time Stamp change or detects a long PFCP connectivity fault with the DBNG-CP,

- the DBNG-UP MUST reset its local Recovery Time Stamp for the DBNG-CP;
- the DBNG-UP MUST delete the session contexts and release the relative allocated resources;
- the DBNG-UP MUST delete all the SGRPs and IP prefixes which the DBNG-UP was programmed to handle, and MUST withdraw promptly the routing advertisements relative to those prefixes.

[R-174] If the DBNG-CP detects a DBNG-UP restart from a Recovery Time Stamp change or detects a long PFCP connectivity fault with the DBNG-UP,

- in the case the DBNG-CP supports a Recovery Time Stamp per DBNG-UP, it MUST reset the Recovery Time Stamp for this DBNG-UP
- the DBNG-CP MUST carry out, limited to the sessions associated to non-resilient SGRPs programmed on that specific DBNG-UP, the operations that follow the Session Deletion procedure towards the external systems;
- the DBNG-CP MUST delete the session contexts associated to non-resilient SGRPs programmed on that specific DBNG-UP and release the relative allocated resources;
- the DBNG-CP MUST release all the non-resilient SGRPs and the IP prefixes associated with SGRPs which were programmed on that specific DBNG-UP;

Note: It is implementation specific whether the CP acts immediately or delays this action.

- the DBNG-CP MUST retain resilient SGRPs, its IP prefixes and associated session contexts which were programmed on the failed DBNG-UP, because such SGRPs and sessions could become active on other DBNG-UP(s);
- the DBNG-CP MUST carry out the switchover, limited to the active resilient SGRPs programmed on that specific DBNG-UP, onto other DBNG-UPs.

Note: the operations that follow the Session Deletion procedure towards external systems may consist in sending AAA Accounting Stop messages, notifications to LEA, etc.

Deleting session contexts and releasing the relative allocated resources on DBNG-CP implies freeing up resources such as IP addresses, cleaning session states from subscriber database, etc.

[R-175] In order to make an efficient use of the heartbeat messages, the DBNG-CP and the DBNG-UP MAY avoid sending a new PFCP Heartbeat Request to the peer if, within an interval equal to T_h from the previous PFCP Heartbeat Request, they receive a Heartbeat Response from the peer.

6.4 General PFCP information exchanges for a subscriber session

For the MS-BNG use cases, this document separates PDRs into two categories.

- A. PDRs to match on subscriber control packets
 - Typically require a minimum of two PDRs
 1. To redirect control packets from access to the DBNG-CP through the CPR Interface.
 2. To redirect control packets from DBNG-CP back to access through the CPR Interface. Control packets can include: DHCP, PPP discovery packets, PPP (LCP, Auth, IPCP, IPv6CP), L2TP (control), and router solicits.
- B. PDRs to match on subscriber data packets
 - Again, typically require a minimum of two PDRs
 1. To forward traffic upstream by matching on data packets arriving from the access and forwarding the packets to the network interface.
 2. To forward traffic downstream by matching on IP packets from the network interface and forwarding the packets back to the subscriber.

Therefore, a typical subscriber session would require at least 4 PDRs.

6.4.1 General PFCP rules for control packet redirection

For redirecting control packets from the DBNG-UP to the DBNG-CP, the following grouped IEs are typically used:

- PDR – Identifies the rule.
- PDI – A grouped IE to specify the matching criteria using the source interface and the traffic endpoint. Filter rules are sometimes used to match on more specific sub-flow. Details are in section 6.4.4.
- FAR – Specify the forwarding action and the destination for the redirected control packet. The control messages are encapsulated for tunneling.
- For more information on:
 - Traffic endpoint see section 6.8.1.5
 - Filter see section 6.4.4.

A typical template is shown in Table 20 below for control packet redirection from the DBNG-UP to the DBNG-CP through the CPR Interface.

Table 20: Example of a PDR for Control Packet Redirection from DBNG-UP to DBNG-CP

| Direction | PDR | FAR |
|---------------------------------------|--|--|
| Control packet from the RG to DBNG-CP | PDR ID PDI: Source Interface Traffic-Endpoint Filter FAR ID | FAR ID Apply Action Forwarding Parameters: Destination Interface Outer Header Creation |

For redirecting control packets from the DBNG-CP to the DBNG-UP, the following list of grouped IEs are typically used:

- PDR – Identifies the rule.
- Outer header removal – Removes the tunnel encapsulation from the control packet.
- PDI – A grouped IE to specify the matching criteria using the source interface and the traffic endpoint.
- FAR – Specify the forwarding action, the destination, and the traffic endpoint for the control packet.
- For more information on:
 - Traffic endpoint see section 6.8.1.5

From the attributes above, a typical template is shown in Table 21 below for control packet redirection from the DBNG-CP to the DBNG-UP through the CPR interface.

Table 21: Example of a PDR for Control Packet redirection from DBNG-CP to DBNG-UP

| Direction | PDR | FAR |
|---------------------------------------|--|---|
| Control packet from DBNG-CP to the RG | PDR ID Precedence Outer Header Removal PDI: Source Interface Traffic-Endpoint FAR ID | FAR ID Apply Action Forwarding Parameters: Destination Interface Linked Traffic Endpoint ID |

6.4.2 General PFCP rules for data packet forwarding

For the upstream direction, data packets from the subscriber are routed through the network interface. PFCP utilizes a list of IEs to program the subscriber data forwarding. The following is a list of grouped IEs typically used for wireline:

- PDR – Identifies the rule.
- PDI – A grouped IE to specify the matching criteria using a combination of source interface and traffic endpoint. Filter rules are sometimes used to match on more specific sub-flow. Details are in section 6.4.4.
- FAR – Specify the forwarding action and the destination for the data packet.
- For more information on:
 - Traffic endpoint see section 6.8.1.5
 - BBF Outer Header removal see section 6.9.4.

Below in Table 22, a typical template for upstream data packet forwarding is shown. Traffic is forwarded from the subscriber through the DBNG-UP to the network core.

Table 22: Example of a PDR for upstream data packet forwarding through the DBNG-UP

| Direction | PDR | FAR |
|-----------|---|---|
| Upstream | PDR ID BBF Outer Header Removal PDI: Source Interface Traffic-Endpoint Filter FAR ID | FAR ID Apply Action Forwarding Parameters: Destination Interface Network-Instance |

***Bolded text indicates BBF PFCP extension**

For the downstream direction, IP packets are routed from the network to the DBNG-UP and are forwarded back to the subscriber, the following list of grouped IEs are typically used:

- PDR – Identifies the rule.
- PDI – A grouped IE to specify the matching criteria using the source interface and the traffic endpoint.
- FAR – Specify the forwarding action, the destination, and the traffic endpoint for the data packet.
- For more information on:
 - Traffic endpoint see section 6.8.1.5
 - BBF Outer Header creation see section 6.9.3

Below in Table 23, a typical template for downstream data packet forwarding is shown. Traffic is forwarded from the network core through the DBNG-UP and then to the subscriber.

Table 23: Example of a PDR for downstream data packet forwarding through the DBNG-UP

| Direction | PDR | FAR |
|------------|--|---|
| Downstream | PDR ID PDI: Source Interface Traffic-Endpoint FAR ID | FAR ID Apply Action Forwarding Parameters: Destination Interface BBF Outer Header Creation Linked Traffic Endpoint ID |

***Bolded text indicates BBF PFCP extension**

6.4.3 General PFCP rules for Server Control Packet Redirection

For redirecting external DHCP or DHCPv6 server originated control packets from the DBNG-UP to the DBNG-CP when the server is reachable through the DBNG-UP, the following grouped IEs are used:

- PDR – Identifies the rule.
- PDI – A grouped IE to specify the matching criteria using the source interface and the traffic endpoint. Filter rules are used to match on downstream DHCP and DHCPv6 control packets. Details are in section 6.4.4.

- FAR – Specify the forwarding action and the destination for the redirected control packet. The control messages are encapsulated for tunneling.
- For more information on:
 - Traffic endpoint see section 6.8.1.5
 - Filter see section 6.4.4.

Table 24 and Table 25 below specify server control packet redirection from the DBNG-UP to the DBNG-CP through the CPR Interface for DHCPv4 and DHCPv6.

Table 24: PDR for DHCPv4 Server Control Packet Redirection from DBNG-UP to DBNG-CP

| Direction | PDR | FAR |
|---|--|--|
| Control packet from the External DHCP Server to DBNG-CP | PDR ID PDI: Source Interface = Core Network Instance = "xyz" SDF Filter = protocol=udp, srcPort=67,dstPort=67 and 68 FAR ID | FAR ID Apply Action = forward Forwarding Parameters: Destination Interface= CP-function Outer Header Creation = Include GTP-U 5/1 bit default- TEID=xxxx, IPv4/v6=CP-IP- address> |

Note that, for the SDF Filter, the srclp may optionally be specified for each external DHCP Server but is not mandatory. The destination IP address will either be the UP IP address for the V-interface (GIADDR) or the RG's assigned IP address (for the case of DHCP renew or request), and, thus, it is not practical to specify the destIp in the SDF. Note that, by default, there is a server control packet redirect tunnel per network instance in which DHCP relay is configured.

Table 25: PDR for DHCPv6 Server Control Packet Redirection from DBNG-UP to DBNG-CP

| Direction | PDR | FAR |
|---|--|--|
| Control packet from the External DHCPv6 Server to DBNG-CP | PDR ID PDI: Source Interface = Core Network Instance = "xyz" SDF Filter = protocol=udp, srcPort=547,dstPort=547 FAR ID | FAR ID Apply Action = forward Forwarding Parameters: Destination Interface= CP-function Outer Header Creation = Include GTP-U 5/1 bit default- TEID=xxxx, IPv4/v6=CP-IP- address> |

Note that, for the SDF Filter, the srclpv6 may optionally be specified for each external DHCPv6 Server but is not mandatory. The destination IPv6 address will be the DBNG relay global IPv6 address (e.g., the UP IPv6 address for the V-interface). Note that, by default, there is a server control packet redirect tunnel per network instance in which DHCPv6 relay is configured.

For redirecting upstream DHCP and DHCPv6 control packets from the DBNG-CP to the DBNG-UP, the following list of grouped IEs are used:

- PDR – Identifies the rule.
- Outer header removal – Removes the tunnel encapsulation from the control packet.
- PDI – A grouped IE to specify the matching criteria using the source interface and the traffic endpoint.
- FAR – Specify the forwarding action, the destination, and the traffic endpoint for the control packet.
- For more information on:
 - Traffic endpoint see section 6.8.1.5

Table 26 below specifies server control packet redirection from the DBNG-CP to the DBNG-UP through the CPR Interface for DHCPv4 and DHCPv6.

Table 26: PDR for DHCP Server Control Packet Redirection from DBNG-CP to DBNG-UP

| Direction | PDR | FAR |
|-----------|-----|-----|
|-----------|-----|-----|

| | | |
|--|---|--|
| Control packet from DBNG-CP to the external DHCP/DHCPv6 server via DBNG-UP | PDR ID Precedence Outer Header Removal = Remove GTP-U PDI: Source Interface=CP-function Local F-TEID = Choose FAR ID | FAR ID Apply Action Forwarding Parameters: Destination Interface = Core Network Instance = "xyz" |
|--|---|--|

The CPR rules specified above create one PFCP session for each Network Instance. Based on the BBF UP function feature flag "NSH header insertion on CPR tunnel" in Table 65, an NSH header may optionally be included to carry the network instance, allowing for a single, server CPRi for all network instances.

As an example, to allow the forwarding of DHCP relay packets towards the network interface, the BBF UP feature list should include the following:

- IPoE
- Issue 2 minimum feature set
- Server Control Packet Redirection

6.4.4 General Information PFCP Filter IEs

Filters are required when a sub-flow of a traffic endpoint needs to be further separated. Matching on a sub-flow, can be done at Layer 2, Layer 3, or both. For Layer 2, IE Ethernet Packet Filter is used and for Layer 3, IE Service Data Filter (SDF) is used. Both types of filters are well defined in 3GPP TS 29.244 [28]. To cover the wireline case, further extensions are required on Ethernet Packet Filter IE, see section 6.8.1.2

6.4.5 General information on NSH header

When the DBNG-CP and DBNG-UP establish a PFCP session, a unique GTP tunnel ID (TEID) is assigned per subscriber session to tunnel the control packets between the DBNG-CP and DBNG-UP. The DBNG-CP must also be able to return on the tunnel control packets back to the subscriber's exact logical port. Both the TEID and logical port are key information for the PFCP session establishment procedure where the DBNG-CP programs bi-directional traffic forwarding rules on the DBNG-UP.

One of the NSH header purposes is to allow the DBNG-CP to learn the subscriber logical port. The NSH header encapsulation contains the logical port information which is mandatory for each subscriber's initial control message packet redirected to the DBNG-CP. Afterwards, the logical port information obtained by the DBNG-CP can be used to program the control packet forwarding rule on the DBNG-UP to redirect the control traffic from DBNG-CP to the subscriber.

For both per-logical-port tunnels and dedicated subscriber tunnels, the NSH header on subsequent control packets would not be strictly necessary for programming bidirectional forwarding rules: the logical port required for forwarding is actually already programmed on the DBNG-UP as a traffic endpoint. However, nothing prevents the DBNG-CP from programming the DBNG-UP to insert the NSH header in those packets too.

6.4.5.1 GTP Tunnel Endpoint ID

In per subscriber dedicated CPR tunnel, every control packet exchanged between the DBNG-CP and DBNG-UP is encapsulated with a GTP Tunnel Endpoint ID (TEID). The TEID is uniquely assigned to each individual subscriber, providing the full subscriber context on the DBNG-CP.

6.4.5.2 NSH header insertion option on CPRi

While GTP TEID provides a way to uniquely identify a subscriber control session in DBNG-CP, it MAY also be required identifying a set of subscribers in order to deterministically process their control packets. NSH metadata provides an aggregate identification for a group of subscribers sharing a common logical port, so that DBNG-CP could determine the affinity of control packets based on NSH header.

A BBF UP function feature flag named “Issue 2 minimal feature set” may be set however to indicate that the NSH header would be included on every control packet redirected in the CPR tunnel. In this case, the resulting protocol stack will include the GTP header followed by the NSH header and then by the full Ethernet control packet sent by the subscriber.

The inclusion of the NSH header in a subscriber dedicated CPR is up to the discretion of the DBNG-CP. Without the use of the NSH header, the DBNG-CP may use the logical port for subscriber context lookup inside of the TEID. In the case where NSH header is required, the DBNG-CP utilizes BBF Outer Header Creation IE to indicate that the NSH header is required.

6.5 PFCP use case and information exchanges

This section describes the information exchange required to cover different MS-BNG use cases. IEs that are extended by BBF are highlighted in **bold**.

6.5.1 Overview of default CPR and per logical port CPR tunnel

6.5.1.1 Default Control Packet Redirect Rule(s)

Upon initialization, the DBNG-CP is agnostic of the DBNG-UP physical layout. The DBNG-CP programs a default PFCP rule that redirects all unknown subscriber packets through a dedicated CPR interface tunnel also known as the default tunnel. Unknown subscribers are subscriber without any prior PFCP session on the DBNG-CP. The DBNG-CP can either create a subscriber session immediately or further delay the subscriber session creation.

Immediate subscriber PFCP session creation

Immediately after the first subscriber redirected control packet to the DBNG-CP, the DBNG-CP may initialize a new PFCP session for the subscriber. It may also choose to not respond to the subscriber and leave the subscriber on the default tunnel for various reasons such as mitigating malicious users. The immediate PFCP session establishment allows the DBNG to control per subscriber session at the earliest possible such as control message throttling, filtering, and QoS treatment.

Benefits

- Agnostic of DBNG-UP physical layout.
- Only one common CPR tunnel is required.
- Ability to control subscriber session as early as possible (both control and data packets).

Costs

- May consume PFCP sessions prior to authentication.

Note: The DBNG-CP should have mechanism to mitigate DDoS attack with malicious RG control messages in order to avoid unnecessary consumption of resources on DBNG UP.

6.5.1.2 Per logical port Control Packet Redirect Rule(s)

An alternative to default CPR tunnel is to pre-configure in the DBNG-CP a list of DBNG-UP logical ports. This model allows for both immediate subscriber PFCP session creation and delayed session creation. Note: Methodology of pre-configuration of the logical ports on DBNG-CP and management of the port list on the DBNG-CP is out of scope for this WT.

Immediate subscriber PFCP session creation

Similar to the default CPR tunnel, the DBNG-CP would initialize a PFCP session for each of the defined logical port after PFCP association. In this case, the per logical port CPR tunnel would only have an upstream rule. The per logical port CPR tunnel redirects every subscriber initial control packet to the DBNG-CP. Afterwards, the DBNG-CP may initialize a new PFCP session for the subscriber. It may also choose to not respond to the subscriber and leave the subscriber on the per logical port CPR tunnel for various reasons such as mitigating malicious users.

Delayed subscriber PFCP session creation

In this case, the per logical port PFCP session is a bi-directional rule to forward subscriber control packets per logical port. All unknown subscriber on the same logical port will reuse the CPR tunnel. Prior to an address assignment to the subscriber such as DHCP ACK or IPCP ACK, the DBNG-CP MUST create a PFCP session for the subscriber. In case where a PPPoE session that must be tunneled over an L2TP tunnel, the DBNG-CP MUST create the PFCP session for the subscriber before the identification or the creation of the L2TP tunnel. From this point forwards, the subscriber would use the dedicated CPR tunnel signaled in the dedicated PFCP session.

The DBNG-CP can choose between delayed and immediate subscriber PFCP session creation on a per-session basis, for example based on protocol, on a given UP.

Benefits

- Dedicated PFCP sessions are consumed only after the address assignment or the L2TP tunnel identification.

Costs

- The DBNG-CP configuration of the DBNG-UP logical port is vendor specific.
- Additional PFCP sessions are consumed, as there are per logical port sessions as well as subscriber PFCP sessions.
- Lack of control on the specific subscriber control packet exchange during the delayed period.

6.5.1.3 Use case: Default control packet redirection for immediate session Creation

As shown in the call flow diagram in section 4.5.2, the DBNG-CP and DBNG-UP form an association using PFCP Association Setup messages. During the association setup, the DBNG-UP informs the DBNG-CP of DBNG function(s) support. Based on this information, the DBNG-UP is instructed to program a default forwarding rule to redirect all control packets from unknown subscribers to the DBNG-CP.

6.5.1.3.1 Immediate Session PFCP control packet redirection rule

Control packet redirection PFCP rules follow the general description highlighted in section 6.4.1. Filters rules would be required to match a control packet for example includes: DHCPv4 or PPPoE packet. A new extension is required to tunnel control packet along with attached meta data to the DBNG-CP:

- **BBF Outer Header Creation:** A BBF IE extension to insert the logical port information as a Network Service Header (NSH) to the subscriber control packet when redirected to the DBNG-CP. Details of this new extension is in section 6.9.3.

6.5.1.4 Use case: Default control packet redirection for delayed session creation.

As shown in the call flow diagram in section 4.5.2, the DBNG-CP and DBNG-UP form an association using PFCP Association Setup messages. During the association setup, the DBNG-UP informs the DBNG-CP of DBNG function(s) support. Based on this information, the DBNG-UP is instructed to program default forwarding rules to redirect all control packets from unknown subscribers to the DBNG-CP in upstream direction. A control packet redirection rule is required to forward control packets from DBNG-CP to the RGs in downstream direction in case the subscriber session creation is delayed. To indicate the support of DBNG function(s) a new extension is required:

- **BBF UP Function Features:** A BBF IE extension for capability flag. Please see new extension in section 6.7.1

6.5.1.4.1 Delayed Session PFCP control packet redirection rule

Control packet redirection PFCP rules follow the general description highlighted in section 6.4.1. Filter rules would be required to match a control packet such as a DHCPv4 or a PPPoE packet in upstream and downstream directions. A default two-way CPR tunnel is required for each access interface.

Upstream CPR:

- PDI rule: Match on control packet received by the DBNG-UP, the traffic-endpoint IE, and filter
 - Traffic-Endpoint IE content is the “logical port”. The traffic endpoint MUST not include MAC address. In case of a per-logical port tunnel, the “C-Tag Range” and “S-Tag Range” IEs maybe used within an “Ethernet Packet Filter” to restrict traffic to a specific VLAN range.
 - Filter defines either “SDF Filter” or “Ethernet Packet Filter” based on IPoE or PPPoE case.
- FAR rule: DBNG-UP forwards the matched control packet to the DBNG-CP
 - Outer Header Creation IE contains “GTPu-UDP-IPv4/v6”

Table 27: Example of a PDR for Control Packet Redirection from DBNG-UP to DBNG-CP

| Direction | PDR | FAR |
|---------------------------------------|--|---|
| Control packet from the RG to DBNG-CP | PDR ID PDI: Source Interface Traffic-Endpoint Filter FAR ID | FAR ID Apply Action Forwarding Parameters: Destination Interface BBF Outer Header Creation Outer Header Creation |

Downstream CPR:

- PDI rule: Match the control packet with default GTPu Tunnel ID and the traffic-endpoint.
 - Traffic Endpoint IE content is GTPu-Tunnel-ID
 - Outer Header Removal IE, DBNG-UP removes the GTP-u Outer header which contains “GTPu-UDP-IPv4/v6”
- FAR rule: DBNG-UP forwards the control packet to the RG using Linked Traffic Endpoint ID.
 - The Linked Traffic Endpoint ID is the traffic endpoint of corresponding upstream PDR for that access interface.

Table 28: Example of a PDR for Common Control Packet redirection from DBNG-CP to DBNG-UP

| Direction | PDR | FAR |
|---------------------------------------|--|---|
| Control packet from DBNG-CP to the RG | PDR ID Precedence Outer Header Removal PDI: Source Interface Traffic-Endpoint FAR ID | FAR ID Apply Action Forwarding Parameters: Destination Interface Linked Traffic Endpoint ID |

6.5.2 Use case: IPoE

For IPoE using DHCPv4, DHCPv6, or Neighbor Discovery (for SLAAC). The data forwarding rules are split between control packet forwarding and data traffic forwarding. Control packets, DHCP, DHCPv6, and router solicit packets must be redirected through the CPR Interface to DBNG-CP for address assignment. And subscriber IPoE data traffic is forwarded through the User Plane. For DHCPv4 or DHCPv6 relay in which the external server is connected to the DBNG-UP via the A10 interface, data forwarding rules extend to control packet exchanges between DBNG-CP and external DHCP or DHCPv6 server via DBNG-UP.

6.5.2.1 PFCP Control Packet redirection rule

Control packet redirection PFCP rules follow the general description highlighted in section 6.4.1. Filter rules are required to match on DHCP, DHCPv6, or ICMPv6 control packets only. Further extensions are required on Traffic Endpoint support IPoE use case:

- **Traffic-Endpoint IE extensions required:** Ethernet header information such as C-Tag, S-Tag, and the logical port where the control packet is coming from. Detailed information of the extension is in section 6.8.1.5.

6.5.2.2 PFCP Data Packet Forwarding rule

Data packet forwarding PFCP rules follow the general description highlighted in section 6.4.2. Below are further extensions required to support IPoE use case for data forwarding:

- **BBF Outer Header Creation:** BBF extended IE to construct the Ethernet header for packet forwarding to the subscriber. Details of the extension are in section 6.9.3.
- **BBF Outer Header Removal:** BBF extended IE to remove the Ethernet header from data packets before forwarding to the network interface. Details of the extension are in section 6.9.4.
- **Traffic-Endpoint IE extensions required:** Ethernet header information such as C-Tag, S-Tag, and the logical port where the control packet is coming from. Detailed information of the extension is in section 6.8.1.5.

6.5.3 Use case: PPPoE

For PPPoE, data forwarding rules are split between control packet redirection and data packet forwarding. Only control packets such as PPPoE discovery, PPP LCP, PPP PAP/ CHAP, PPP IPCP or PPP IPV6CP should be redirected to the CPR Interface while data traffic should be routed through the network interface.

6.5.3.1 PFCP Control Packet redirection rule

PFCP Control packet redirection PFCP rules follow the general description highlighted in section 6.4.1. Filter rules would match on PPP control message such as discovery, LCP, and NCP. For PPPoEv6, DHCPv6 and Router Solicit messages should also be redirected to the DBNG-CP. Further extensions are required to support PPPoE use case:

- **Ethernet packet filter IE extension required:** Specifies PPP attributes such as the PPP protocol type. Details of the extension are in section 6.8.1.3.
- **Traffic-Endpoint IE extensions required:** Ethernet header information such as C-Tag, S-Tag, PPPoE session ID, and the logical port where the control packet is coming from. Detailed information of the extension is in section 6.8.1.5.

Note: Including PPPoE session ID in Traffic-Endpoint IE is optional as PPPoE session ID may not be generated on reception of PADI at the DBNG-CP (In the case where the PPPoE session ID is available and included in the Traffic-Endpoint, the PADR will utilize the default CPR interface). An alternative is to include PPPoE session ID using Update Traffic-endpoint IE at a later time.

6.5.3.2 PFCP Data Packet Forwarding rule

Data packet forwarding PFCP rules follow the general description highlighted in section 6.4.2. Further extensions required to support PPPoE use case for data forwarding:

- **BBF Outer Header Creation:** BBF extended IE to construct the PPPoE header for packet forwarding to the subscriber. Details of the extension are in section 6.9.3.
- **BBF Outer Header Removal:** BBF extended IE to remove the PPPoE header from data packets before forwarding to the network interface. Details of the extension are in section 6.9.4.
- **Ethernet packet filter IE extension required:** Specified PPP attributes such as the PPP protocol type. Details of the extension are in section 6.8.1.3.
- **Traffic-Endpoint IE extensions required:** Ethernet header information such as C-Tag, S-Tag, PPPoE session ID, and the logical port where the control packet is coming from. Detailed information of the extension is in section 6.8.1.5.
- **PPP LCP Connectivity IE:** An IE to specify the LCP echo keepalive properties.
- **MTU:** An IE used to enforce the MRU specified by the CPE.

6.5.4 Use Case: L2TP LAC

The call flow for LAC subscriber will initially follow the same PPPoE procedure as specified in section 6.5.3, authentication will inform the DBNG-CP that the subscriber requires an L2TP tunnel to the LNS. A PFCP session dedicated to the L2TP tunnel is established, in order to redirect the traffic exchanged between LAC and LNS. Afterward when the L2TP tunnel and session are established, the PPPoE session would be tunneled to the LNS. Therefore, the PFCP sessions are as follows:

1. An initial PFCP session to allow PPPoE procedure follows the control packet forwarding rule listed in section 6.4.1. After authentication, the DBNG-CP identifies that the subscriber requires an LNS connection. The LAC will need to determine:
 - If a new L2TP tunnel is required for the L2TP session, section 6.5.4.1 outline the PFCP information exchange to allow for L2TP tunnel setup.
 - Or if an existing L2TP tunnel can be reused, then move to step 2)
2. The subscriber PFCP session is modified to forward both PPP control and data packets to the LNS, while the PPPoE control packets coming from the subscriber, such as PADT, are still redirected to the LAC. See section 6.5.4.2.

6.5.4.1 PFCP session for L2TP tunnel setup

Each L2TP tunnel setup would require an individual PFCP session to forward L2TP control messages between DBNG-CP and DBNG-UP. The PFCP rules to redirect L2TP tunnel setup control messages follow the general description highlighted in section 6.4.1, instead of forwarding control message to access ports, L2TP control messages are forwarded to the network core instead. L2TP control messages from specific tunnels are redirected from the network core to the DBNG-CP. Further extensions are required on Traffic Endpoint to support L2TP tunnel signaling:

- **Traffic-Endpoint IE extensions required:** For the DBNG-UP to DBNG-CP direction, matching L2TP control messages based on tunnel ID and allowing DBNG-UP to select the IP address for the L2TP tunnel. Detailed information of the extension is in section 6.8.1.5.
- **L2TP type:** Used to match on L2TP control messages.

6.5.4.2 PFCP update for L2TP session

The PFCP rule in section 6.5.3.1 will continue to redirect PPPoE discovery control packets such as PADT to the local LAC-CP. Other PPP control packets, previously redirected to the LAC DBNG-CP, are updated to be forwarded to the LNS instead, specified in section 6.5.4.1. Both PPP control and data packet forwarding PFCP rules to the LNS follow the general description highlighted in section 6.2. Further extensions are required on Traffic Endpoint to forward PPP packets to the L2TP tunnel:

- **BBF Outer Header Creation:** BBF extended IE to construct the Ethernet downstream and creating the L2TP header upstream. Details of the extension are in section 6.9.3.
- **BBF Outer Header Removal:** BBF extended IE to remove the Ethernet header in the upstream direction and removing the L2TP header in the downstream direction. Details of the extension are in section 6.9.4.
- **Ethernet packet filter IE extension required:** Specifies PPP attributes such as the PPP protocol types. Details of the extension are in section 6.8.1.3.
- **Traffic-Endpoint IE extensions required:** For upstream, Ethernet header information such as C-Tag, S-Tag, and the logical port where the control packet is coming from. And for downstream, match on specific session within a L2TP tunnel. Detailed information of the extension is in section 6.8.1.5.
- **L2TP type:** Identify to only match on L2TP data messages.

Note: The SRC IP of packet going towards LNS shall be obtained from the L2TP Tunnel Endpoint IE present in the downstream PDR and for this purpose the traffic endpoints must be linked using the Linked Traffic Endpoint ID IE.

6.5.5 Use Case: L2TP LNS

After DBNG-UP and DBNG-CP association, a default PFCP rule is installed to redirect new L2TP control message for tunnel setup from DBNG-UP to DBNG-CP. Upon completion of PPP authentication, a PFCP session can be installed for control message and data packets forwarding. Authentication can be sped up with standard L2TP LCP and authentication proxy. Therefore, the PFCP sessions are as follows:

1. A PFCP session for each L2TP tunnel, described below in section 6.5.5.1
2. Individual PFCP session for each subscriber L2TP session, described in section 6.5.5.2 and 6.5.5.3

6.5.5.1 PFCP session for L2TP tunnel setup

Individual PFCP session is used to redirect L2TP control packets from the DBNG-UP to DBNG-CP. The PFCP rules to redirect L2TP tunnel setup control messages follow the general description highlighted in section 6.4.1. Further extensions are required on Traffic Endpoint to support L2TP tunnel control signaling:

- **Traffic-Endpoint IE extensions required:** For the DBNG-UP to DBNG-CP direction, matching L2TP control messages based on tunnel ID and tunnel IP address. Detailed information of the extension is in section 6.8.1.5.
- **L2TP type:** Used to match on L2TP control messages.

6.5.5.2 PFCP Control Packet redirection rule

L2TP data packets are split into two categories, one is for PPP control and the other is subscriber data traffic. Control packet redirection PFCP rules follow the general description highlighted in section 6.4.1. L2TP data packet containing PPP control message must be redirected to the DBNG-CP for processing. Further extensions are required on Traffic Endpoint and Ethernet Filter to support L2TP use case:

- **Ethernet packet filter IE extension required:** Redirects PPP control messages within the L2TP data packet only. Details of the extension are in section 6.8.1.3
- **Traffic-Endpoint IE extensions required:** To match on specific session within a L2TP tunnel. Detailed information of the extension is in section 6.8.1.5.
- **L2TP type:** Used to match on L2TP data messages.

6.5.5.3 PFCP Data Packet Forwarding rule

This rule is programmed after authentication. Data packet forwarding PFCP rules follows the general description highlighted in section 6.4.2. Further extensions are required to support LNS use case for data forwarding:

- **BBF Outer Header Creation:** BBF extended IE to construct both L2TP and PPP header downstream. Details of the extension are in section 6.9.3.
- **BBF Outer Header Removal:** BBF extended IE to remove both L2TP and PPP header in the upstream direction. Details of the extension are in section 6.9.4.
- **Traffic-Endpoint IE extensions required:** for upstream to match on specific L2TP tunnel ID and the session ID. For downstream match on the subscriber IP address information. Detailed information of the extension is in section 6.8.1.5.
- **L2TP type:** Used to match on L2TP data messages.
- **PPP LCP Connectivity IE:** An IE to specify the LCP echo hello properties. Note, upon failure, a PFCP session report is sent from the DBNG-UP to the DBNG-CP informing of the inactivity.
- **MTU:** Used to enforce the MRU specified by the CPE.

Note: The SRC IP of packet going towards LAC shall be obtained from the L2TP Tunnel Endpoint IE present in the upstream PDR and for this purpose the traffic endpoints must be linked using the Linked Traffic Endpoint ID IE.

6.5.6 Use Case: TWAG

In all layer 2 and layer 3 models, the DBNG-UP can optionally decide the TEID for the DBNG-CP to use. There are two ways for DBNG-UP to select TEID range and inform the DBNG-CP.

- 1) During association the DBNG-UP can pass the range of TEID for DBNG-CP to use
- 2) During PFCP session establishment, the DBNG-CP requests a TEID from the DBNG-UP to use.

Both options are covered in section 6.5.6.1. Afterwards, the PFCP session will update the DBNG-UP (or if there is no PFCP session, on establishing one) to redirect control traffic to the DBNG-CP and to forward data traffic to the EPC core through a GTP encapsulation.

6.5.6.1 DBNG-UP TEID assignment (optional)

The DBNG-CP establishes a PFCP session with the DBNG-UP and requests for a TEID. The PDR will request for TEIDs from the DBNG-UP. The DBNG-UP will respond with a list of TEIDs for the DBNG-CP in a session response message.

6.5.6.2 PFCP Control Packet redirection rule

For upstream: From step 8 of the layer 2 trigger initial attached call flow in section 4.5.30.1 and step 12 of the layer 3 trigger initial attached call flow in section 4.5.30.3, the PDN GW would have assigned the DBNG-CP an IP address and/or IPv6 SLAAC prefix for the subscriber. DHCP and RS control packets from the subscriber access interface are redirected to the CPR Interface following the general description highlighted in section 6.4.1. Further extensions are required on Traffic Endpoint to support the TWAG use case:

- **Traffic-Endpoint IE extensions required:** Ethernet header information such as C-Tag, S-Tag, and the logical port where the control packet is coming from. Detailed information of the extension is in section 6.8.1.5.

6.5.6.3 PFCP Data Packet Forwarding rule

Data packet forwarding PFCP rules follow the general description highlighted in section 6.4.2. Further extensions are required to support TWAG use case for data forwarding:

- **BBF Outer Header Removal:** BBF extended IE to remove the Ethernet header from data packets before forwarding to the EPC. Details of the extension are in section 6.9.4
- **Traffic-Endpoint IE extensions required:** Ethernet header information such as C-Tag, S-Tag, and the logical port where the control packet is coming from. Detailed information of the extension is in section 6.8.1.5.

6.5.7 Subscriber Group and Subscriber Prefix Use Case

As explained in section 4.4.2, the use of SGRP allows the same restoration fate to be assigned to a group of subscriber sessions. In this section, the “BBF SGRP” IE is specified.

The IPv4 default gateway or IPv6 default gateway is signaled as part of the SGRP Prefix list, with the prefix length value set to individual host address length (32 or 128).

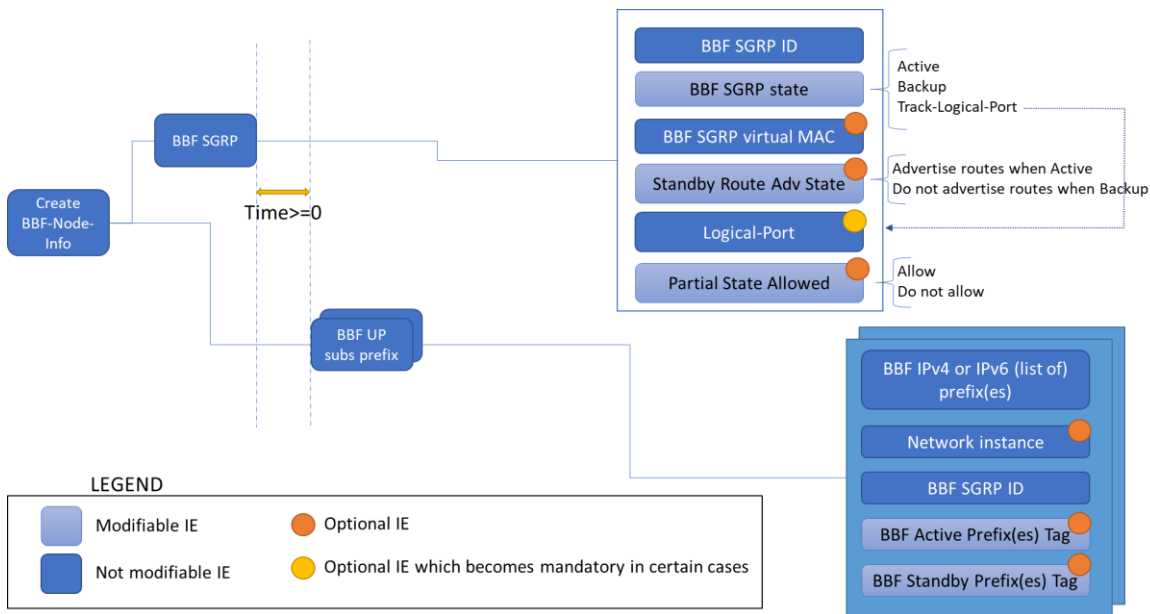


Figure 80: Subscriber Group and Subscriber Prefix

As shown in Figure 80, BBF SGRP must be created by the DBNG-CP before the subscriber prefixes that are associated with it.

This can be accomplished in two ways:

1. With different PFCP Association Update Request messages including a BBF-Node-Info Create IE, where the message that creates the SGRP must be sent before the message that populates the SGRP with prefixes.
2. With a single PFCP Association Update Request message including a BBF-Node-Info Create IE that contains both the SGRP and the list of prefixes for this SGRP tagged with this SGRP ID.

The BBF SGRP may map to one or more address pools defined on the DBNG-CP. In other words, a SGRP essentially contains a list of prefixes belonging to one or more pools.

A SGRP can also exist with no associated Subscriber Prefix (e.g., when assigning /32 from AAA server).

Note: It is good practice to map an SGRP to a single pool and populate a pool with prefixes that have similar characteristics in terms of resilience, network instance, etc"

BBF SGRP can be associated with a Logical Port. This Logical Port is used by the DBNG-UP when DBNG-CP has set the SGRP Track-Logical-Port State as defined in Section 4.4.7 and 6.7.6.

Multiple SGRPs can be associated with the same Logical Port.

In deployments, a DBNG-UP will most likely have more than one SGRP assigned to it.

6.6 BBF PFCP Information Element Summary

BBF PFCP IE type values specified in Table 29 through Table 33 are IE applicability tables for each broadband use case. "Yes" indicates that the IE indicated in column 2 MUST be included for that particular use case. "No" indicates that the IE indicated in column 2 MUST NOT be included for that particular use case. "Opt" means the IE can provide more specific matching information for the PFCP session and MAY be included for that particular use case.

CP UL: control packets uplink, from access to DBNG-CP

CP DL control packets downlink, from DBNG-CP to access

CP UL to N: control packets uplink to network, for LAC use case

CP DL from N: control packets downlink from network, for LAC use case

DP UL: data packets uplink, from access to network

DP DL: data packets downlink, from network to access

Table 29: BBF extended Information Elements for Node Association, Control Packet Redirect Sessions

| IE Type value (Decimal) | Information Elements | Section | Use Case | | | |
|-------------------------|------------------------------|-----------|-----------------------|-------------|------------------|-------|
| | | | CP and UP Association | Default CPR | Per logical port | |
| | | | | CP UL | CP UL | CP DL |
| 32768 | BBF UP Function Features | 6.9.1 | Yes | No | No | No |
| 32769 | Logical Port | 6.9.2 | No | No | No | No |
| 32770 | BBF Outer Header Creation | 6.9.3 | No | Yes | Opt | No |
| 32771 | BBF Outer Header Removal | 6.9.4 | No | No | No | No |
| 32772 | PPPoE Session ID | 6.9.5 | No | No | No | No |
| 32773 | PPP protocol | 6.9.6 | No | Opt | Opt | No |
| 32774 | Verification Timers | 6.9.7 | No | No | No | No |
| 32775 | PPP LCP Magic Number | 6.9.8 | No | No | No | No |
| 32776 | MTU | 6.9.9 | No | No | No | No |
| 32777 | L2TP tunnel endpoint | 6.9.10 | No | Opt | Opt | No |
| 32778 | L2TP session ID | 6.9.11 | No | No | No | No |
| 32779 | L2TP type | 6.9.12 | No | Opt | Opt | No |
| 32780 | PPP LCP connectivity | 6.8.4 | No | No | No | No |
| 32781 | L2TP Tunnel | 6.8.6 | No | Opt | Opt | No |
| 32793 | BBF-node-info-create | 6.7.3.1 | No | No | No | No |
| 32794 | BBF-node-info-modify | 6.7.3.2 | No | No | No | No |
| 32795 | BBF-node-info-delete | 6.7.3.3 | No | No | No | No |
| 32796 | BBF Logical Port Report | 6.7.5.2 | No | No | No | No |
| 32797 | BBF SGRP Notification Report | 6.7.5.3 | No | No | No | No |
| 32798 | BBF Network Instance Report | 6.7.5.4 | No | No | No | No |
| 32799 | BBF SGRP Error | 6.7.5.3.1 | No | No | No | No |
| 32800 | BBF SGRP | 6.7.6 | No | No | No | No |
| 32801 | BBF UP Subscriber Prefix | 6.7.7 | No | No | No | No |
| 32819 | BBF Prefix Error | 6.7.5.3.2 | No | No | No | No |
| 32802 | BBF ACL | 6.8.7 | No | No | No | No |
| 32803 | BBF Direction | 6.9.13 | No | No | No | No |
| 32804 | BBF Family | 6.9.14 | No | No | No | No |
| 32806 | BBF SGRP Identifier | 6.9.15 | No | No | No | No |
| 32807 | BBF SGRP State | 6.9.16 | No | No | No | No |
| 32808 | BBF SGRP Flags | 6.9.17 | No | No | No | No |
| 32809 | BBF | 6.9.18 | No | No | No | No |

| IE Type value (Decimal) | Information Elements | Section | Use Case | | | |
|-------------------------|----------------------------------|---------|-----------------------|-------------|------------------|-------|
| | | | CP and UP Association | Default CPR | Per logical port | |
| | | | | CP UL | CP UL | CP DL |
| | Operational Condition | | | | | |
| 32810 | BBF IPv4 Prefix | 6.9.19 | No | No | No | No |
| 32811 | BBF IPv6 Prefix | 6.9.20 | No | No | No | No |
| 32812 | BBF Prefix Tag | 6.9.21 | No | No | No | No |
| 32813 | BBF Error Code | 6.9.22 | No | No | No | No |
| 32814 | BBF Error Message | 6.9.23 | No | No | No | No |
| 32815 | BBF Maximum ACL Chain Length | 6.9.24 | No | No | No | No |
| 32816 | BBF Forwarding Capacity | 6.9.25 | No | No | No | No |
| 32817 | BBF Connectivity Status | 6.9.26 | No | No | No | No |
| 32818 | Vendor-specific Node Report Type | 6.9.27 | No | No | No | No |
| 32820 | BBF C-Tag Range | 6.9.28 | No | No | Yes | No |
| 32821 | BBF S-Tag Range | 6.9.29 | No | No | Yes | No |

Table 30: BBF extended Information Elements for IPoE/ TWAG

| IE Type value (Decimal) | Information Elements | Section | Use Case | | | | | | | | |
|-------------------------|------------------------------|-----------|----------|-------|-------|-------|-------|-------|-------|-------|-----|
| | | | IPoE | | | | TWAG | | | | |
| | | | CP UL | CP DL | UP UL | UP DL | CP UL | CP DL | UP UL | UP DL | |
| 32768 | BBF UP Function Features | 6.9.1 | No | No | No | No | No | No | No | No | No |
| 32769 | Logical Port | 6.9.2 | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 32770 | BBF Outer Header Creation | 6.9.3 | OPT | No | No | Yes | OPT | No | No | No | Yes |
| 32771 | BBF Outer Header Removal | 6.9.4 | No | No | Yes | No | No | No | Yes | No | No |
| 32772 | PPPoE Session ID | 6.9.5 | No | No | No | No | No | No | No | No | No |
| 32773 | PPP protocol | 6.9.6 | No | No | No | No | No | No | No | No | No |
| 32774 | Verification Timers | 6.9.7 | No | No | No | No | No | No | No | No | No |
| 32775 | PPP LCP Magic Number | 6.9.8 | No | No | No | No | No | No | No | No | No |
| 32776 | MTU | 6.9.9 | No | No | No | No | No | No | No | No | No |
| 32777 | L2TP tunnel endpoint | 6.9.10 | No | No | No | No | No | No | No | No | No |
| 32778 | L2TP session ID | 6.9.11 | No | No | No | No | No | No | No | No | No |
| 32779 | L2TP type | 6.9.12 | No | No | No | No | No | No | No | No | No |
| 32780 | PPP LCP connectivity | 6.8.4 | No | No | No | No | No | No | No | No | No |
| 32781 | L2TP Tunnel | 6.8.6 | No | No | No | No | No | No | No | No | No |
| 32793 | BBF-node-info-create | 6.7.3.1 | No | No | No | No | No | No | No | No | No |
| 32794 | BBF-node-info-modify | 6.7.3.2 | No | No | No | No | No | No | No | No | No |
| 32795 | BBF-node-info-delete | 6.7.3.3 | No | No | No | No | No | No | No | No | No |
| 32796 | BBF Logical Port Report | 6.7.5.2 | No | No | No | No | No | No | No | No | No |
| 32797 | BBF SGRP Notification Report | 6.7.5.3 | No | No | No | No | No | No | No | No | No |
| 32798 | BBF Network Instance Report | 6.7.5.4 | No | No | No | No | No | No | No | No | No |
| 32799 | BBF SGRP Error | 6.7.5.3.1 | No | No | No | No | No | No | No | No | No |
| 32800 | BBF SGRP | 6.7.6 | No | No | No | No | No | No | No | No | No |
| 32801 | BBF UP Subscriber Prefix | 6.7.7 | No | No | No | No | No | No | No | No | No |
| 32819 | BBF Prefix Error | 6.7.5.3.2 | No | No | No | No | No | No | No | No | No |
| 32802 | BBF ACL | 6.8.7 | No | No | Opt | Opt | No | No | No | No | Opt |
| 32803 | BBF Direction | 6.9.13 | No | No | Opt | Opt | No | No | No | No | Opt |
| 32804 | BBF Family | 6.9.14 | No | No | Opt | Opt | No | No | No | No | Opt |
| 32806 | BBF SGRP Identifier | 6.9.15 | No | No | No | No | Yes | No | No | No | No |
| 32807 | BBF SGRP State | 6.9.16 | No | No | No | No | No | No | No | No | No |
| 32808 | BBF SGRP Flags | 6.9.17 | No | No | No | No | No | No | No | No | No |
| 32809 | BBF | 6.9.18 | No | No | No | No | No | No | No | No | No |

| IE Type value (Decimal) | Information Elements | Section | Use Case | | | | | | | | |
|-------------------------|----------------------------------|---------|----------|-------|-------|-------|-------|-------|-------|-------|----|
| | | | IPoE | | | | TWAG | | | | |
| | | | CP UL | CP DL | UP UL | UP DL | CP UL | CP DL | UP UL | UP DL | |
| | Operational Condition | | | | | | | | | | |
| 32810 | BBF IPv4 Prefix | 6.9.19 | No | No | No | No | No | No | No | No | No |
| 32811 | BBF IPv6 Prefix | 6.9.20 | No | No | No | No | No | No | No | No | No |
| 32812 | BBF Prefix Tag | 6.9.21 | No | No | No | No | No | No | No | No | No |
| 32813 | BBF Error Code | 6.9.22 | No | No | No | No | No | No | No | No | No |
| 32814 | BBF Error Message | 6.9.23 | No | No | No | No | No | No | No | No | No |
| 32815 | BBF Maximum ACL Chain Length | 6.9.24 | No | No | No | No | No | No | No | No | No |
| 32816 | BBF Forwarding Capacity | 6.9.25 | No | No | No | No | No | No | No | No | No |
| 32817 | BBF Connectivity Status | 6.9.26 | No | No | No | No | No | No | No | No | No |
| 32818 | Vendor-specific Node Report Type | 6.9.27 | No | No | No | No | No | No | No | No | No |
| 32820 | BBF C-Tag Range | 6.9.28 | No | No | No | No | No | No | No | No | No |
| 32821 | BBF S-Tag Range | 6.9.29 | No | No | No | No | No | No | No | No | No |

Table 31: BBF extended Information Elements for PPPoE

| IE Type value (Decimal) | Information Elements | Section | Use Case | | | | |
|-------------------------|------------------------------|-----------|----------|-------|-------|-------|---------|
| | | | PPPoE | | | | |
| | | | CP UL | CP DL | UP UL | UP DL | Session |
| 32768 | BBF UP Function Features | 6.9.1 | No | No | No | No | No |
| 32769 | Logical Port | 6.9.2 | Yes | Yes | Yes | Yes | No |
| 32770 | BBF Outer Header Creation | 6.9.3 | OPT | No | No | Yes | No |
| 32771 | BBF Outer Header Removal | 6.9.4 | No | No | Yes | No | No |
| 32772 | PPPoE Session ID | 6.9.5 | Yes | No | Yes | No | No |
| 32773 | PPP protocol | 6.9.6 | Yes | No | Yes | No | No |
| 32774 | Verification Timers | 6.9.7 | No | No | No | No | Yes |
| 32775 | PPP LCP Magic Number | 6.9.8 | No | No | No | No | Yes |
| 32776 | MTU | 6.9.9 | No | No | No | Yes | No |
| 32777 | L2TP tunnel endpoint | 6.9.10 | No | No | No | No | No |
| 32778 | L2TP session ID | 6.9.11 | No | No | No | No | No |
| 32779 | L2TP type | 6.9.12 | No | No | No | No | No |
| 32780 | PPP LCP connectivity | 6.8.4 | No | No | No | No | Yes |
| 32781 | L2TP Tunnel | 6.8.6 | No | No | No | No | No |
| 32793 | BBF-node-info-create | 6.7.3.1 | No | No | No | No | No |
| 32794 | BBF-node-info-modify | 6.7.3.2 | No | No | No | No | No |
| 32795 | BBF-node-info-delete | 6.7.3.3 | No | No | No | No | No |
| 32796 | BBF Logical Port Report | 6.7.5.2 | No | No | No | No | No |
| 32797 | BBF SGRP Notification Report | 6.7.5.3 | No | No | No | No | No |
| 32798 | BBF Network Instance Report | 6.7.5.4 | No | No | No | No | No |
| 32799 | BBF SGRP Error | 6.7.5.3.1 | No | No | No | No | No |
| 32800 | BBF SGRP | 6.7.6 | No | No | No | No | No |
| 32801 | BBF UP Subscriber Prefix | 6.7.7 | No | No | No | No | No |
| 32819 | BBF Prefix Error | 6.7.5.3.2 | No | No | No | No | No |
| 32802 | BBF ACL | 6.8.7 | No | No | Opt | Opt | No |
| 32803 | BBF Direction | 6.9.13 | No | No | Opt | Opt | No |
| 32804 | BBF Family | 6.9.14 | No | No | Opt | Opt | No |
| 32806 | BBF SGRP Identifier | 6.9.15 | No | No | No | No | Yes |
| 32807 | BBF SGRP State | 6.9.16 | No | No | No | No | No |
| 32808 | BBF SGRP Flags | 6.9.17 | No | No | No | No | No |
| 32809 | BBF | 6.9.18 | No | No | No | No | No |

| IE Type value (Decimal) | Information Elements | Section | Use Case | | | | | |
|-------------------------|----------------------------------|---------|----------|-------|-------|-------|---------|----|
| | | | PPPoE | | | | | |
| | | | CP UL | CP DL | UP UL | UP DL | Session | |
| | Operational Condition | | | | | | | |
| 32810 | BBF IPv4 Prefix | 6.9.19 | No | No | No | No | No | No |
| 32811 | BBF IPv6 Prefix | 6.9.20 | No | No | No | No | No | No |
| 32812 | BBF Prefix Tag | 6.9.21 | No | No | No | No | No | No |
| 32813 | BBF Error Code | 6.9.22 | No | No | No | No | No | No |
| 32814 | BBF Error Message | 6.9.23 | No | No | No | No | No | No |
| 32815 | BBF Maximum ACL Chain Length | 6.9.24 | No | No | No | No | No | No |
| 32816 | BBF Forwarding Capacity | 6.9.25 | No | No | No | No | No | No |
| 32817 | BBF Connectivity Status | 6.9.26 | No | No | No | No | No | No |
| 32818 | Vendor-specific Node Report Type | 6.9.27 | No | No | No | No | No | No |
| 32820 | BBF C-Tag Range | 6.9.28 | No | No | No | No | No | No |
| 32821 | BBF S-Tag Range | 6.9.29 | No | No | No | No | No | No |

Table 32: BBF extended Information Elements for L2TP LAC

| IE Type value (Decimal) | Information Elements | Section | Use Case | | | | | |
|-------------------------|------------------------------|-----------|------------|--------------|------------------------|-------|-------|-------|
| | | | LAC tunnel | | LAC subscriber session | | | |
| | | | CP UL to N | CP DL from N | CP UL | CP DL | UP UL | UP DL |
| 32768 | BBF UP Function Features | 6.9.1 | No | No | No | No | No | No |
| 32769 | Logical Port | 6.9.2 | No | No | Yes | Yes | Yes | Yes |
| 32770 | BBF Outer Header Creation | 6.9.3 | No | No | Opt | No | Yes | Yes |
| 32771 | BBF Outer Header Removal | 6.9.4 | No | No | No | No | Yes | No |
| 32772 | PPPoE Session ID | 6.9.5 | No | No | Yes | No | Yes | No |
| 32773 | PPP protocol | 6.9.6 | No | No | Yes | No | Yes | No |
| 32774 | Verification Timers | 6.9.7 | No | No | No | No | No | No |
| 32775 | PPP LCP Magic Number | 6.9.8 | No | No | No | No | No | No |
| 32776 | MTU | 6.9.9 | No | No | No | No | No | No |
| 32777 | L2TP tunnel endpoint | 6.9.10 | No | Yes | No | No | No | Yes |
| 32778 | L2TP session ID | 6.9.11 | No | No | No | No | No | Yes |
| 32779 | L2TP type | 6.9.12 | No | Yes | No | No | No | Yes |
| 32780 | PPP LCP connectivity | 6.8.4 | No | No | No | No | No | No |
| 32781 | L2TP Tunnel | 6.8.6 | No | Yes | No | No | No | Yes |
| 32793 | BBF-node-info-create | 6.7.3.1 | No | No | No | No | No | No |
| 32794 | BBF-node-info-modify | 6.7.3.2 | No | No | No | No | No | No |
| 32795 | BBF-node-info-delete | 6.7.3.3 | No | No | No | No | No | No |
| 32796 | BBF Logical Port Report | 6.7.5.2 | No | No | No | No | No | No |
| 32797 | BBF SGRP Notification Report | 6.7.5.3 | No | No | No | No | No | No |
| 32798 | BBF Network Instance Report | 6.7.5.4 | No | No | No | No | No | No |
| 32799 | BBF SGRP Error | 6.7.5.3.1 | No | No | No | No | No | No |
| 32800 | BBF SGRP | 6.7.6 | No | No | No | No | No | No |
| 32801 | BBF UP Subscriber Prefix | 6.7.7 | No | No | No | No | No | No |
| 32819 | BBF Prefix Error | 6.7.5.3.2 | No | No | No | No | No | No |
| 32802 | BBF ACL | 6.8.7 | No | No | No | No | Opt | Opt |
| 32803 | BBF Direction | 6.9.13 | No | No | No | No | Opt | Opt |
| 32804 | BBF Family | 6.9.14 | No | No | No | No | Opt | Opt |
| 32806 | BBF SGRP Identifier | 6.9.15 | No | No | No | No | No | No |
| 32807 | BBF SGRP State | 6.9.16 | No | No | No | No | No | No |
| 32808 | BBF SGRP Flags | 6.9.17 | No | No | No | No | No | No |
| 32809 | BBF | 6.9.18 | No | No | No | No | No | No |

| IE Type value (Decimal) | Information Elements | Section | Use Case | | | | | |
|-------------------------|----------------------------------|---------|------------|--------------|------------------------|-------|-------|-------|
| | | | LAC tunnel | | LAC subscriber session | | | |
| | | | CP UL to N | CP DL from N | CP UL | CP DL | UP UL | UP DL |
| | Operational Condition | | | | | | | |
| 32810 | BBF IPv4 Prefix | 6.9.19 | No | No | No | No | No | No |
| 32811 | BBF IPv6 Prefix | 6.9.20 | No | No | No | No | No | No |
| 32812 | BBF Prefix Tag | 6.9.21 | No | No | No | No | No | No |
| 32813 | BBF Error Code | 6.9.22 | No | No | No | No | No | No |
| 32814 | BBF Error Message | 6.9.23 | No | No | No | No | No | No |
| 32815 | BBF Maximum ACL Chain Length | 6.9.24 | No | No | No | No | No | No |
| 32816 | BBF Forwarding Capacity | 6.9.25 | No | No | No | No | No | No |
| 32817 | BBF Connectivity Status | 6.9.26 | No | No | No | No | No | No |
| 32818 | Vendor-specific Node Report Type | 6.9.27 | No | No | No | No | No | No |
| 32820 | BBF C-Tag Range | 6.9.28 | No | No | No | No | No | No |
| 32821 | BBF S-Tag Range | 6.9.29 | No | No | No | No | No | No |

Table 33: BBF extended Information Elements for L2TP LNS

| IE Type value (Decimal) | Information Elements | Section | Use Case | | | | | | |
|-------------------------|------------------------------|-----------|------------|-------|------------------------|-------|-------|-------|---------|
| | | | LNS tunnel | | LNS subscriber session | | | | |
| | | | CP UL | CP DL | CP UL | CP DL | UP UL | UP DL | Session |
| 32768 | BBF UP Function Features | 6.9.1 | No | No | No | No | No | No | No |
| 32769 | Logical Port | 6.9.2 | No | No | No | No | No | No | No |
| 32770 | BBF Outer Header Creation | 6.9.3 | No | No | No | No | No | Yes | No |
| 32771 | BBF Outer Header Removal | 6.9.4 | No | No | No | No | Yes | No | No |
| 32772 | PPPoE Session ID | 6.9.5 | No | No | No | No | No | No | No |
| 32773 | PPP protocol | 6.9.6 | No | No | Yes | No | Yes | No | No |
| 32774 | Verification Timers | 6.9.7 | No | No | No | No | No | No | Yes |
| 32775 | PPP LCP Magic Number | 6.9.8 | No | No | No | No | No | No | Yes |
| 32776 | MTU | 6.9.9 | No | No | No | No | Yes | No | No |
| 32777 | L2TP tunnel endpoint | 6.9.10 | Yes | No | No | No | Yes | No | No |
| 32778 | L2TP session ID | 6.9.11 | No | No | No | No | Yes | No | No |
| 32779 | L2TP type | 6.9.12 | Yes | No | No | No | Yes | No | No |
| 32780 | PPP LCP connectivity | 6.8.4 | No | No | No | No | No | No | Yes |
| 32781 | L2TP Tunnel | 6.8.6 | Yes | No | No | No | Yes | No | No |
| 32793 | BBF-node-info-create | 6.7.3.1 | No | No | No | No | No | No | No |
| 32794 | BBF-node-info-modify | 6.7.3.2 | No | No | No | No | No | No | No |
| 32795 | BBF-node-info-delete | 6.7.3.3 | No | No | No | No | No | No | No |
| 32796 | BBF Logical Port Report | 6.7.5.2 | No | No | No | No | No | No | No |
| 32797 | BBF SGRP Notification Report | 6.7.5.3 | No | No | No | No | No | No | No |
| 32798 | BBF Network Instance Report | 6.7.5.4 | No | No | No | No | No | No | No |
| 32799 | BBF SGRP Error | 6.7.5.3.1 | No | No | No | No | No | No | No |
| 32800 | BBF SGRP | 6.7.6 | No | No | No | No | No | No | No |
| 32801 | BBF UP Subscriber Prefix | 6.7.7 | No | No | No | No | No | No | No |
| 32819 | BBF Prefix Error | 6.7.5.3.2 | No | No | No | No | No | No | No |
| 32802 | BBF ACL | 6.8.7 | No | No | No | No | Opt | Opt | No |
| 32803 | BBF Direction | 6.9.13 | No | No | No | No | Opt | Opt | No |
| 32804 | BBF Family | 6.9.14 | No | No | No | No | Opt | Opt | No |
| 32806 | BBF SGRP Identifier | 6.9.15 | No | No | No | No | No | No | Yes |
| 32807 | BBF SGRP State | 6.9.16 | No | No | No | No | No | No | No |
| 32808 | BBF SGRP Flags | 6.9.17 | No | No | No | No | No | No | No |
| 32809 | BBF Operational | 6.9.18 | No | No | No | No | No | No | No |

| IE Type value (Decimal) | Information Elements | Section | Use Case | | | | | | |
|-------------------------|----------------------------------|---------|------------|-------|------------------------|-------|-------|-------|---------|
| | | | LNS tunnel | | LNS subscriber session | | | | |
| | | | CP UL | CP DL | CP UL | CP DL | UP UL | UP DL | Session |
| | Condition | | | | | | | | |
| 32810 | BBF IPv4 Prefix | 6.9.19 | No | No | No | No | No | No | No |
| 32811 | BBF IPv6 Prefix | 6.9.20 | No | No | No | No | No | No | No |
| 32812 | BBF Prefix Tag | 6.9.21 | No | No | No | No | No | No | No |
| 32813 | BBF Error Code | 6.9.22 | No | No | No | No | No | No | No |
| 32814 | BBF Error Message | 6.9.23 | No | No | No | No | No | No | No |
| 32815 | BBF Maximum ACL Chain Length | 6.9.24 | No | No | No | No | No | No | No |
| 32816 | BBF Forwarding Capacity | 6.9.25 | No | No | No | No | No | No | No |
| 32817 | BBF Connectivity Status | 6.9.26 | No | No | No | No | No | No | No |
| 32818 | Vendor-specific Node Report Type | 6.9.27 | No | No | No | No | No | No | No |
| 32820 | BBF C-Tag Range | 6.9.28 | No | No | No | No | No | No | No |
| 32821 | BBF S-Tag Range | 6.9.29 | No | No | No | No | No | No | No |

Below tables shows which BBF IEs are processed according to each BBF UP Function Feature flag. However, it must be noted even though some BBF UP Function Feature flag are a superset of others, the functionality is limited to only the BBF UP Function Feature flag. For example, the UP signals only the LNS function, although LNS is a super set of PPPoE, the DBNG-UP would be expected to only support the LNS function.

Table 34: BBF UP Function Features and required BBF extended PFCP IEs

| BBF UP Function Features | IE Type Value | IE name |
|--------------------------|---------------|----------------------|
| PPPoE | 32772 | PPPoE Session ID |
| | 32773 | PPP protocol |
| | 32776 | MTU |
| IPoE | N/A | N/A |
| LAC | 32772 | PPPoE Session ID |
| | 32773 | PPP protocol |
| | 32777 | L2TP tunnel endpoint |
| | 32778 | L2TP session ID |
| | 32779 | L2TP type |
| | 32781 | L2TP Tunnel |
| LNS | 32772 | PPPoE Session ID |
| | 32773 | PPP protocol |
| | 32776 | MTU |
| | 32777 | L2TP tunnel endpoint |
| | 32778 | L2TP session ID |
| | 32779 | L2TP type |
| LCP keepalive offload | 32774 | Verification Timers |
| | 32775 | PPP LCP Magic Number |
| | 32780 | PPP LCP connectivity |

| BBF UP Function Features | IE Type Value | IE name |
|-----------------------------|----------------------------------|------------------------------|
| Issue 2 minimum feature set | 32793 | BBF-node-info-create |
| | 32794 | BBF-node-info-modify |
| | 32795 | BBF-node-info-delete |
| | 32796 | BBF Logical Port Report |
| | 32797 | BBF SGRP Notification Report |
| | 32798 | BBF Network Instance Report |
| | 32799 | BBF SGRP Error |
| | 32800 | BBF SGRP |
| | 32801 | BBF UP Subscriber Prefix |
| | 32819 | BBF Prefix Error |
| | 32805 | BBF Node Report Type |
| | 32806 | BBF SGRP Identifier |
| | 32807 | BBF SGRP State |
| | 32808 | BBF SGRP Flags |
| | 32809 | BBF Operational Condition |
| | 32810 | BBF IPv4 Prefix |
| | 32811 | BBF IPv6 Prefix |
| | 32812 | BBF Prefix Tag |
| | 32813 | BBF Error Code |
| | 32814 | BBF Error Message |
| | 32815 | Maximum ACL Chain Length |
| 32816 | BBF Forwarding Capacity | |
| 32817 | BBF Connectivity Status | |
| 32818 | Vendor-specific Node Report Type | |
| 32820 | C-Tag Range | |
| 32821 | S-Tag Range | |
| Volume quota and thresholds | *N/A | *N/A |
| ACL | 32802 | BBF ACL |
| | 32803 | BBF Direction |
| | 32804 | BBF Family |
| Relay | *N/A | *N/A |

*Requires only standard 3GPP PFCP IEs

The following table provides a reference guide for 3GPP PFCP IEs and its applicability to this TR use cases.

Table 35: 3GPP PFCP Information Element Types and applicability

| IE Type value (Decimal) | 3GPP Information Elements | Use Case | | | | | | | | |
|-------------------------|------------------------------|-----------------------|-------------|-------|------|------------|------------------------|------------|------------------------|------|
| | | CP and UP Association | Default CPR | PPPoE | IPoE | LAC tunnel | LAC subscriber session | LNS tunnel | LNS subscriber session | TWAG |
| 1 | Create PDR | No | yes | yes | yes | yes | yes | yes | yes | yes |
| 2 | PDI | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 3 | Create FAR | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 4 | Forwarding Parameters | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 8 | Created PDR | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 9 | Update PDR | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 10 | Update FAR | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 11 | Update Forwarding Parameters | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 15 | Remove PDR | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 16 | Remove FAR | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 19 | Cause | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 20 | Source Interface | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 21 | F-TEID | No | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 22 | Network Instance | No | ? | No | No | Yes | Yes | Yes | Yes | No |
| 23 | SDF Filter | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 29 | Precedence | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 42 | Destination Interface | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 43 | UP Function Features | Yes | No | No | No | No | No | No | No | No |
| 44 | Apply Action | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 56 | PDR ID | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 57 | F-SEID | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 60 | Node ID | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |

| IE Type value (Decimal) | 3GPP Information Elements | Use Case | | | | | | | | |
|-------------------------|---------------------------|-----------------------|-------------|-------|------|------------|------------------------|------------|------------------------|------|
| | | CP and UP Association | Default CPR | PPPoE | IPoE | LAC tunnel | LAC subscriber session | LNS tunnel | LNS subscriber session | TWAG |
| 84 | Outer Header Creation | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 89 | CP Function Features | Yes | No | No | No | No | No | No | No | No |
| 93 | UE IP Address | No | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 95 | Outer Header Removal | No | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 96 | Recovery Time Stamp | Yes | No | No | No | No | No | No | No | No |
| 108 | FAR ID | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 113 | PDN Type | No | No | Yes* | Yes* | No | Yes* | No | Yes** | Yes* |
| 127 | Create Traffic Endpoint | No | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 128 | Created Traffic Endpoint | No | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 129 | Update Traffic Endpoint | No | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 130 | Remove Traffic Endpoint | No | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 131 | Traffic Endpoint ID | No | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 132 | Ethernet Packet Filter | No | Yes | Yes | No | Yes | Yes | Yes | Yes | No |
| 133 | MAC address | No | No | Yes | Yes | Yes | Yes | No | No | Yes |
| 134 | C-TAG | No | No | Yes | Yes | Yes | Yes | No | No | Yes |
| 135 | S-TAG | No | No | Yes | Yes | Yes | Yes | No | No | Yes |
| 136 | Ethertype | No | Yes | Yes | No | Yes | Yes | Yes | Yes | No |
| 153 | Framed-Route | No | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 155 | Framed-IPv6-Route | No | No | Yes | Yes | Yes | Yes | Yes | Yes | Yes |

*: Must be Ethernet

** : Must be IPv4v6, which indicates the LNS session may support any of IPv4 only, IPv6 only or dual stack connectivity.

6.7 PFCP Grouped IE extensions for Node Related Messages

This section highlights extensions required on grouped IEs such as Association setup and Association update messages, to cover DBNG use cases. PFCP IEs in this section includes a presence requirement, indicated by the letter “P”:

- Mandatory (M) indicates that this IE MUST be present if used.
- Conditional (C) indicates that this IE is conditional based on use case specified.

6.7.1 PFCP Association Setup Request

In the case where the DBNG-UP initiates the Association Setup procedure, it sends an Association Setup Request which includes the “BBF UP Function Features” IE to notify to the DBNG-CP the BBF supported features.

Table 36: BBF extended Information Element(s) in a PFCP Association Setup Request

| Information element | P | Condition / Comment | IE Type | IE Originating Entity |
|---------------------------------|---|--|--|-----------------------|
| BBF UP Function Features | C | This IE MUST be present if the DBNG-UP function sends this message, and the DBNG-UP function supports at least one DBNG-UP feature defined in this IE. When present, this IE MUST indicate the features the BBF UP Function supports. | BBF UP Function Features Details in section 6.9.1 | DBNG-UP |
| Maximum ACL Chain Length | O | This IE MUST be present if the DBNG-UP supports ACL Chaining | Maximum ACL Chain Length Details in section 6.9.24 | DBNG-UP |

6.7.2 PFCP Association Setup Response

In the case where the DBNG-CP initiates the association setup, the DBNG-UP would respond with this IE “BBF UP Function Features” to notify the DBNG-CP of its functional support.

Table 37: BBF extended Information Element(s) in a PFCP Association Setup Response

| Information element | P | Condition / Comment | IE-Type | IE Originating Entity |
|---------------------------------|---|--|--|-----------------------|
| BBF UP Function Features | C | This IE MUST be present if the DBNG-UP function sends this message, and the DBNG-UP function supports at least one DBNG-UP feature defined in this IE. When present, this IE MUST indicate the features the BBF UP Function supports. | BBF UP Function Features Details in section 6.9.1 | DBNG-UP |
| Maximum ACL Chain Length | O | This IE MUST be present if the DBNG-UP supports ACL Chaining | Maximum ACL Chain Length Details in section 6.9.24 | DBNG-UP |

6.7.3 PFCP Association Update Request

In the case of a change in the supported BBF features, the DBNG-UP sends an Association Update Request which includes the “BBF UP Function Features” IE to notify to the DBNG-CP the features change.

In case the DBNG-CP needs to create or delete or modify an SGRP or a prefix on a UP, the DBNG-CP sends an Association Update Request which includes respectively the BBF-node-info-create IE, the BBF-node-info-delete IE, the BBF-node-info-modify IE.

Table 38: BBF extended Information Element(s) in a PFCP Association Update Request

| Information element | P | Condition / Comment | IE Type | IE Originating Entity |
|---------------------------------|---|---|--|-----------------------|
| BBF UP Function Features | C | If present, this IE shall indicate the supported Features when the sending node is the UP function. | BBF UP Function Features Details in section 6.9.1 | DBNG-UP |
| Create BBF-Node-Info | C | This IE shall be included if node information is to be created on the UP. This IE MUST NOT be included when the DBNG-UP sends this message. Several IEs with the same IE type may be present. | BBF-node-info-create Details in section 6.7.3.1 | DBNG-CP |
| Modify BBF-Node-info | C | This IE shall be included if node information is to be updated on the UP. This IE MUST NOT be included when the DBNG-UP sends this message. Several IEs with the same IE type may be present. | BBF-node-info-modify Details in section 6.7.3.2 | DBNG-CP |
| Remove BBF-Node-info | C | This IE shall be included if node information is to be deleted on the UP. This IE MUST NOT be included when the DBNG-UP sends this message. Several IEs with the same IE type may be present. | BBF-node-info-delete Details in section 6.7.3.3 | DBNG-CP |
| Maximum ACL Chain Length | O | This IE MUST be present if the DBNG-UP supports ACL Chaining and the maximum chain length changed. | Maximum ACL Chain Length Details in section 6.9.24 | DBNG-UP |

6.7.3.1 BBF-node-info create

Table 39: Information Elements in an BBF-node-info-create

| Octet 1 and 2 | | BBF-node-info-create IE Type= 32793* (decimal) | |
|--------------------------|---|---|--|
| Octets 3 and 4 | | Length = n | |
| Octets 5 and 6 | | Enterprise ID 3561 | |
| Information elements | P | Condition / Comment | IE Type |
| BBF SGRP | C | This IE shall be included to create a subscriber group. Several IEs with the same IE type may be present to provision a list of subscriber groups | BBF SGRP Details in section 6.7.6 |
| BBF UP subscriber prefix | C | This is grouped IE and shall be included to create subscriber prefix(es) Several IEs with the same IE type may be present to provision a list of subscriber prefixes | BBF UP subscriber prefix Details in section 6.7.7 |

* - 32781 was the last IE Type value used in Issue 1. IE Type values 32782-32792 are being used in TR-459.2 and TR-459.3 and hence this specification starts at 32793 from this section.

6.7.3.2 BBF-node-info modify

Table 40: Information Elements in an BBF-node-info-modify

| Octet 1 and 2 | | BBF-node-info-modify IE Type= 32794(decimal) | |
|--------------------------|---|---|---|
| Octets 3 and 4 | | Length = n | |
| Octets 5 and 6 | | Enterprise ID 3561 | |
| Information elements | P | Condition / Comment | IE Type |
| BBF SGRP | C | This IE shall be included to modify subscriber group parameters Several IEs with the same IE type may be present to modify a list of SGRP | BBF SGRP Details in section 6.7.6 |
| BBF UP Subscriber prefix | C | This IE shall be included to modify subscriber prefix parameter(s) Several IEs with the same IE type may be present to change a list of subscriber prefixes. | BBF UP Subscriber prefix Details in section 6.7.7 |

6.7.3.3 BBF-node-info delete

Table 41: Information Elements in an BBF-node-info-delete

| Octet 1 and 2 | | BBF-node-info-delete IE Type= 32795 (decimal) | |
|--------------------------|---|--|---|
| Octets 3 and 4 | | Length = n | |
| Octets 5 and 6 | | Enterprise ID 3561 | |
| Information elements | P | Condition / Comment | IE Type |
| BBF SGRP | C | This IE shall be included to delete a subscriber group. Several IEs with the same IE type may be present to delete a list of SGRPs | BBF SGRP Details in section 6.7.6 |
| BBF UP subscriber prefix | C | This IE shall be present if the CP informs the UP about the prefix to be removed. Several IEs with the same IE type may be present to delete a list of subscriber prefixes | BBF UP subscriber prefix Details in section 6.7.7 |

6.7.4 PFCP Association Update Response

In case the DBNG-CP created or deleted or modified an SGRP or a prefix on a DBNG-UP by sending an Association Update Request, the DBNG-UP replies with an Association Update Response which includes a BBF SGRP Notification IE to report the operational condition of SGRP.

Table 42: BBF extended Information Element(s) in a PFCP Association Update Response

| Information element | P | Condition / Comment | IE-Type | IE Originating Entity |
|---------------------------------|---|---|---|-----------------------|
| BBF UP Function Features | C | If present, this IE shall indicate the supported Features when the sending node is the UP function. | BBF UP Function Features Details in section 6.9.1 | DBNG-UP |
| BBF SGRP Notification | C | This IE MUST be present in response to SGRP create/modify/delete message, to update the status of SGRP on DBNG-UP to DBNG-CP. | BBF SGRP Notification Details in section 6.7.5.3 | DBNG-UP |

6.7.5 PFCP Node Report Request

In order to accomplish wireline case needs, the DBNG-UP may send to the DBNG-CP a special "PFCP Node Report Request", which is distinguished by the presence of the "Vendor Specific Report type" bit.

Section 6.7.5 and 6.7.5.1 define how a "PFCP Node Report Request" specific for the wireline use shall be formulated.

The PFCP Node Report Request shall be sent by the DBNG-UP function to report information to the DBNG-CP function that is not specific to a PFCP session.

The 3GPP PFCP Node Report Request is enhanced with a new BBF Node Report IE.

Table 43: Information Elements in PFCP Node Report Request

| Information elements | P | Condition / Comment | IE Type |
|----------------------------------|---|--|---|
| Vendor-Specific Node Report Type | C | This IE shall be present if the Node Report Type IE indicates a Vendor-Specific Report. When present, this IE shall indicate the type of the Vendor-Specific Report. When Present, the Node Report Type IE MUST NOT set any type bit other than VSR (Vendor-Specific Report) | Vendor-Specific Node Report Type Details in section 6.9.27 |
| BBF Logical Port Report | C | This IE shall be present if the LPR flag is set in the Vendor-Specific Node Report Type IE is set. This report indicates status of a logical port configured on the DBNG-UP to anchor subscriber sessions. Several IEs within the same IE type may be present to report status of multiple logical ports. | BBF Logical Port Report Details in section 6.7.5.2 |
| BBF SGRP Notification Report | C | This IE shall be present if the SGR flag is set in the Vendor-Specific Node Report Type IE is set. This report indicates status of an SGRP and is typically sent when status changes occurred. Several IEs within the same IE type may be present to report status of multiple SGRPs. | BBF SGRP Notification Report Details in section 6.7.5.3 |
| BBF Network Instance Report | C | This IE shall be present if the NIR flag is set in the Vendor-Specific Node Report Type IE is set. This report indicates status of a Network instance on the DBNG-UP. Several IEs within the same IE type may be present to report status of multiple Network Instance. | BBF Network Instance Report Details in section 6.7.5.4 |

6.7.5.1 3GPP Node Report Type

For TR-459i2, the 3GPP Node Report Type IE (Type =101) shall indicate the "Vendor Specific Report type"

For TR-459i2, the Node Report Type IE shall be encoded as shown in Figure 81.




Figure 81: 3GPP Node Report Type

When VSR bit is set to 1, it indicates the Vendor Specific Report is expected. When VSR bit is set, other bits defined by 3GPP shall not be set and should be sent as a separate report.

6.7.5.2 BBF Logical Port Report IE

This report indicates status of a logical port configured on the DBNG-UP to anchor subscriber sessions. This report is a purely informational report and does not pose any requirements or restrictions on the DBNG-CP. For example, sending a usage load of 100% discourages but does not prohibit a DBNG-CP from using that logical port in a subsequent PFCP Session Establishment Request message. Therefore, sending this report does not require any function capability exchange before it can be sent. A DBNG-CP that does not support or is not interested in this Report will simply ignore the IE as per default 3GPP TS 29.244 [28] behavior.

The DBNG-UP SHOULD have a throttling mechanism to control the rate of these notifications.

As per [R-129] in case of DBNG-UP resiliency, the DBNG-UP immediately generates a report upon forwarding capability changes, unless throttling applies.

Otherwise, a DBNG-UP is free when to send a report. The report may be generated immediately upon any status change, but it is also allowed to apply thresholds or hold off timers before sending a status change, or even to send the report periodically.

This report IE can be included in a BBF Vendor-Specific PFCP Node Report Request, according to the definition given in section 6.7.3.3.

Table 44: Information Elements in BBF Logical Port Report

| | | | |
|-----------------------|--|--|---|
| Octet 1 and 2 | BBF Logical Port Report IE Type = 32796 (decimal) | | |
| Octets 3 and 4 | Length = n | | |
| Octets 5 and 6 | Enterprise ID 3561 | | |
| Information elements | P | Condition / Comment | IE Type |
| Logical Port | M | This IE shall contain an opaque value to indicate the logical port on which Ethernet traffic is received | Logical Port Details in section 6.9.2 |
| Forwarding Capability | M | This IE shall contain the forwarding capability of the Logical Port. 0% indicates no forwarding is possible and 100% indicates forwarding at full capacity is possible. Any value between 0% and 100% indicates partial failure and forwarding is possible but not at full capacity, for example failure of a single LAG member. | BBF Forwarding Capability Details in section 6.9.25 |

| | | | |
|-------------|---|---|--|
| | | Partial failure indicates the remaining forwarding capacity as a percentage of the maximum capacity. For example, if an all-active LAG with 5 equal-capacity LAG members is used, and one LAG member fails, a Forwarding Capability of 80% is expected. | |
| MAC Address | O | If present, this IE shall contain the MAC address associated with the logical port in the SOUR(Source) flag/field | MAC Address (3GPP 29.244 IE Type=133) |
| Usage | O | The current load on the port as a percentage value, with 100% indicating the port is fully loaded and can no longer accept more sessions. The exact meaning of the usage value is opaque to the control plane. It is local policy on the UP as to the information used to fill out this IE. | Metric (3GPP TS 29.244 V17.0.0 IE Type=53) |

6.7.5.3 BBF SGRP Notification Report IE

Section 6.7.5 specifies the new BBF grouped IE which is used to send notification of SGRP from DBNG-UP to DBNG-CP. This IE MUST be included in response to create, modify or delete SGRP messages described in section 4.4. This report IE can be included in a BBF Vendor-Specific PFCP Node Report Request, according to the definition given in section 6.7.5 to report the operational condition of the SGRP in DBNG-UP which includes status changes, errors, failures etc.

As shown in Table 45, the BBF SGRP Notification Report is a grouped IE that may nest other grouped IEs which are BBF specific. In the particular case of SGRP errors, it must include the BBF SGRP Error IE described in section 6.7.5.3.1.

In case of prefix errors, the BBF Prefix Error must be included as well: this is described in BBF Error Code Section 6.9.22.

Table 45: Information Elements in BBF SGRP Notification Report

| Octet 1 and 2 | BBF SGRP Notification Report IE Type = 32797 (decimal) | | |
|--------------------------------|---|--|--|
| Octets 3 and 4 | Length = n | | |
| Octets 5 and 6 | Enterprise ID 3561 | | |
| Information elements | P | Condition / Comment | IE Type |
| BBF SGRP ID | M | This IE shall be used to denote the SGRP ID on the DBNG-UP/ DBNG-CP Several IEs with the same IE type may be present to denote a list of SGRP ID(s) and their status | BBF SGRP ID Details in section 6.9.15 |
| BBF SGRP Operational Condition | M | This IE shall be used to denote the SGRP operational condition on DBNG-UP. This IE MUST be present for a SGRP. This indicates possible conditions: <ul style="list-style-type: none"> Up: indicates successful operation for SGRP on DBNG-UP, which includes forwarding, routing, core/access network and dynamic subscriber interface/context operations. Down: indicates SGRP is not operational due to a failure. Possible failures: access/core network, forwarding, routing, subscriber context or any reason internal to DBNG-UP. Not-ready: indicates that the DBNG-UP is not-ready for state changes and DBNG-CP must wait for operational condition of SGRP to return to either up or down value, to make any changes for SGRP. | BBF SGRP operational condition Details in section 6.9.18 |
| BBF SGRP | C | This Grouped IE MUST be included if there is an error in programming | BBF SGRP Error |

| | | | |
|------------------|---|---|---|
| Error | | SGRP on DBNG-UP This IE MUST be included if the BBF SGRP Operational Condition indicates 'Down' | Details in section 6.7.5.3.1 |
| BBF Prefix Error | C | This Grouped IE MUST be included if there is an error in programming the prefix on DBNG-UP. Several IEs with the same IE type may be present to indicate errors for different prefixes. In this case each BBF Prefix Error IE MUST contain either a BBF IPv4 prefix or BBF IPv6 prefix IE. When the Error does not contain a BBF IPv4 prefix IE or BBF IPv6 prefix IE, programming of all prefixes signaled in the Request message failed with the same reason. | BBF Prefix Error Details in section 6.7.5.3.2 |

6.7.5.3.1 BBF SGRP Error

BBF SGRP Error IE is included when DBNG-UP experiences a programming error for SGRP. This IE MUST be included by the DBNG-UP in the response of an SGRP creation/modification message if the programming is not successful, or in a node report sent by the DBNG-UP during normal operation if an error is experienced on the SGRP. The IE is defined in Table 46.

Table 46: Information Elements in BBF SGRP Error

| Octet 1 and 2 | BBF SGRP Error IE Type = 32799 (decimal) | | |
|------------------------|--|---|--|
| Octets 3 and 4 | Length = n | | |
| Octets 5 and 6 | Enterprise ID 3561 | | |
| Information elements | P | Condition / Comment | IE Type |
| BBF SGRP Error Code | C | This IE indicates the error code related to SGRP programming on DBNG-UP | BBF Error Code Details in section 6.9.22 |
| BBF SGRP Error Message | O | This IE Must be included in case SGRP Error Code is set as Other | BBF Error Message Details in section 6.9.23 |

6.7.5.3.2 BBF Prefix Error

BBF Prefix Error IE is included when DBNG-UP experiences a programming error for BBF prefix. This IE MUST be included during prefix create/modification/deletion if an error is experienced by DBNG-UP. The IE is defined in table below.

Table 47: Information Elements in BBF Prefix Error

| Octet 1 and 2 | BBF Prefix Error IE Type = 32819 | | |
|-----------------------|----------------------------------|---|--------------------------------------|
| Octets 3 and 4 | Length = n | | |
| Octets 5 and 6 | Enterprise ID 3561 | | |
| Information elements | P | Condition / Comment | IE Type |
| BBF Prefix Error Code | M | This IE indicates the error code related to Prefix programming on DBNG-UP | BBF Error Code Details in section |

| | | | |
|---------------------------------|----------|--|---|
| | | | 6.9.22 |
| BBF Prefix Error Message | O | This IE Must be included in case Prefix Error Code is set as Other | BBF Error Message Details in section 6.9.23 |
| BBF IPv4 prefix | C | If present, this IE shall identify the IPv4 prefix for which an error occurred. When this is included, the BBF IPv6 prefix IE MUST NOT be included. | BBF IPv4 prefix Details in section 6.9.19 |
| BBF IPv6 prefix | C | If present, this IE shall identify the IPv6 prefix for which an error occurred. When this is included, the BBF IPv4 prefix IE MUST NOT be included. | BBF IPv6 prefix Details in section 6.9.20 |
| Network Instance | C | If present, this IE shall identify the Network Instance for the prefix(es) above. If this is not present, and either the BBF IPv4 Prefix IE or the BBF IPv6 Prefix IE are present, the Network Instance will be the default network instance. | Network Instance (3GPP TS 29.244 IE Type=22) |

6.7.5.4 BBF Network Instance Report IE

This report indicates status of a Network instance configured on the DBNG-UP. This report is a purely informational report and does not pose any requirements or restrictions on the DBNG-CP. Therefore, sending this report does not require any function capability exchange before it can be sent. A DBNG-CP that does not support or is not interested in this Report will simply ignore the IE as per default 3GPP 29.244 [28] behavior.

The DBNG-UP SHOULD have a throttling mechanism to control the rate of these notifications.

As per [R-129] in case of DBNG-UP resiliency, the DBNG-UP immediately generates a report upon status changes, unless throttling applies. The DBNG-UP can also send the report outside of an immediate change, for example periodically.

This report IE can be included in a BBF Vendor-Specific PFCP Node Report Request, according to the definition given in section 6.7.5.

Table 48: Information Elements in BBF Network Instance Report

| | | | |
|----------------------|--|--|---|
| Octet 1 and 2 | BBF Network Instance Report IE Type = 32798 (decimal) | | |
| Octets 3 and 4 | Length = n | | |
| Octets 5 and 6 | Enterprise ID 3561 | | |
| Information elements | P | Condition / Comment | IE Type |
| Network Instance | M | This IE identifies the Network Instance to which the report applies. | Network Instance (3GPP TS 29.244 IE Type=22) |
| Connectivity Status | M | This IE shall contain the connectivity status of the network instance. This status reflects whether the network instance has connectivity to the rest of the network or is isolated. To determine this, the DBNG-UP can for example use one of following technologies: | BBF Connectivity Status Details in section 6.9.26 |

| | | | |
|--|--|--|--|
| | | <ul style="list-style-type: none"> Maintain a set of BFD [43][44] sessions to known upstream interconnect routers. When all BFD [43][44] sessions fail the network instance can be considered isolated, otherwise it is connected. Monitor BGP neighbors, if neighbors are Idle the network instance can be considered isolated, otherwise it is connected. Monitor the A10 network interfaces, if all are down, the network instance can be considered isolated, otherwise it is connected. In an IPSEC VPN monitor the presence of any Security Association, if no SA is present the network instance can be considered isolated, otherwise it is connected. | |
|--|--|--|--|

6.7.6 BBF SGRP

Table 49 specifies the “BBF SGRP” Grouped IE in line with Figure 80.

Table 49: BBF SGRP

| Octet 1 and 2 | | BBF SGRP IE Type = 32800 (decimal) | |
|-------------------------|----------|---|--|
| Octets 3 and 4 | | Length = n | |
| Octets 5 and 6 | | Enterprise ID 3561 | |
| Information elements | P | Condition / Comment | IE Type |
| BBF SGRP ID | M | This IE uniquely identifies the Subscriber Group (SGRP) | BBF SGRP ID Details in section 6.9.15 |
| BBF SGRP State | C | This IE specifies the state of the SGRP. The Status can be: Active or Backup or Track Logical Port | BBF SGRP state Details in section 6.9.16 |
| BBF virtual MAC | O | The IE specifies the V interface MAC for the SGRP ID | MAC address |
| BBF Logical Port | O | The IE specifies the Logical Port Id for the SGRP. This is mandatory when State is set to “Track Logical Port”. Several IEs with the same IE type may be presented to provision multiple Logical Ports | Logical Port Details in section 6.9.2 |
| BBF SGRP Flags | C | This IE is included if any of the following flags is set: <ul style="list-style-type: none"> - Route Advertisement State: Specifies if, when the SGRP is in the backup state, the DBNG-UP should advertise prefixes associated with this SGRP ID, using the tags signaled for that prefix in the BBF UP Subscriber prefix IE. When not set, the DBNG-UP does not advertise the prefixes associated with this SGRP ID when its state is backup. - Partial State Allowed: Specifies if, when the SGRP is in backup state, the DBNG-UP is allowed to use partial implementation state for sessions associated with this SGRP ID. When not set, the DBNG-UP must install full session state when the SGRP is in backup state. | BBF SGRP Flags Details in section 6.9.17 |

The following rules apply:

- BBF SGRP ID - must be always present
- BBF SGRP State - must be always present only in either BBF-node-info-create IE or BBF-node-info-modify IE and can assume the following values:
 - Active
 - Backup
 - Track-logical port: in this case the two DBNG-UPs, which are assumed to be linked by an active-backup connection on the access side, will decide by themselves which one of the two has to handle the sessions on the basis of their relevant logical-ports state
 - In this case, both UPs will receive a message with “BBF SGRP state” set to “Track-Logical-Port”, even if the relevant Logical-Port field, which must be present in the same message, will be populated with a value specific of the UP.
 - Each UP will auto-determine if it is active or backup and adjust advertisements accordingly to the rules received for active/backup state.
 - The UP which recognizes that is active will also handle the keepalives (if the offloading is requested)
 - The UP which recognizes that is backup will also align with the Partial State Allowed flag.
- BBF SGRP virtual MAC - optional. In case the UP resiliency is not supported, it may be not present, but, in case UP resiliency is supported, it should be populated to handle smooth switchovers. When present, the DBNG-CP MUST populate the virtual MAC in the ‘Source MAC address value’ field of the MAC Address IE as defined in 3GPP TS 29.244 [28] clause 8.2.93.
- BBF Logical-Port
 - Mandatory only if the BBF SGRP state is set to «Track-Logical-Port». When «Track-Logical-Port» is set, only a single Logical port is allowed.
 - Otherwise, logical port is optional. Multiple Logical-Ports can optionally be associated with the SGRP, therefore the BBF Logical-Port optional field may be used by the “BBF-Node-info Create” IE to create a new SGRP on a DBNG-UP that may allow logical port(s) to achieve a certain resiliency state. One example where this is useful is in the case where the DBNG-UP needs to establish resiliency through signaling on the V interface and needs the logical ports information prior to subscriber session establishment.

When Logical Port(s) are associated to the SGRP in the BBF SGRP IE, subsequent subscriber session setups should associate to these logical port(s) within the SGRP, otherwise session setup should be rejected by the DBNG-UP.

Multiple Logical Port IEs may be present to indicate multiple logical ports belong to the SGRP in the “BBF-Node-info Create” IE.

When this field is not sent, logical port(s) are associated to the SGRP based on the subscriber session establishments.

Provisioning of multiple Logical-ports under the same SGRP is used to map the SGRP logical ports for which there are no active subscriber sessions. In this use case, the SGRP CP-driven switchover is performed to switch a bulk of logical-ports.

However, it is not recommended to set the state of the SGRP to “Track Logical Port” in case the SGRP has multiple Logical-Ports associated (which is typically done in case the Operator wants to achieve an UP-triggered switchover).

- BBF SGRP Flags - optional: if not present, all flags are assumed to be off, i.e., set to 0.
 - Route Advertisement State behaves as explained in Table 49.
 - Partial State Allowed behaves as explained in Table 49 and section 4.4.10.1. If DBNG-UP does not support partial state and “Partial State Allowed” is enabled, the DBNG-UP shall signal that with a specific error in the PFCP Response and should not accept the message sent by CP.

6.7.6.1 SGRP Modify

SGRP modification can be used to do one of the following actions:

- changing of the “BBF SGRP state”
- changing of the “Backup Route Advertisement State”
- changing of the “Partial State Allowed”.

The other fields of “BBF SGRP” cannot be modified: therefore, if a change is required for those, the SGRP must be deleted and re-created on the user planes that had been assigned that SGRP either as active or as backup.

6.7.6.2 SGRP Delete

The following IEs MUST be ignored, if present, in the BBF-node-info delete message:

- BBF SGRP state
- BBF virtual MAC
- BBF Logical Port
- BBF SGRP Flags

SGRP deletion can be used to delete a SGRP from a DBNG-UP when it's no longer needed. It is important in this scenario that all final accounting statistics are retrieved by the DBNG-CP, therefore the DBNG-CP behaves as follows:

1. SGRP is active on the DBNG-UP or track-logical-port is enabled – In this case, SGRP deletion on the DBNG-UP MUST be preceded by individual subscriber termination and state deletion, using the PFCP Session Deletion procedure.
2. SGRP is in backup on the DBNG-UP – In this case, the SGRP deletion MAY be executed without preliminary actions by the DBNG-CP. The CP then locally removes all PFCP session state associated to that UP and SGRP without using the PFCP Session Deletion procedure. Because during a local SGRP delete the DBNG-UP and DBNG-CP are temporarily out of sync this is not the recommended model. The DBNG-CP SHOULD still precede the SGRP deletion with individual subscriber termination and state deletion where possible.

If the DBNG-UP receives an SGRP delete and there are still PFCP sessions or UP Subscriber Prefixes linked to that SGRP, it MUST locally delete all these session and prefixes without relying on a subsequent PFCP deletion message. The DBNG-UP immediately answers the SGRP delete with the “BBF Operational Condition IE” set to “Not-ready”, it does not delay this message until the local SGRP delete is completed. In case the DBNG-CP subsequently creates the same SGRP while the local delete is ongoing, the DBNG-UP SHOULD reject this with the “BBF Error Code” IE set to “SGRP programming error because of ongoing delete”.

If the DBNG-UP receives an SGRP delete and there are no PFCP sessions or UP Subscriber prefixes linked to that SGRP, it answers the SGRP delete with the “BBF Operational Condition IE” set to “Up”.

SGRP deletion on both active and backup DBNG-UP is required in order to:

- change “BBF Virtual MAC” of a SGRP- in which case it would be followed by an SGRP Create with the new Virtual MAC
- change the “BBF SGRP ID” of a SGRP - in which case SGRP deletion would be followed by an SGRP Create with the new ID
- change “BBF Logical Port” of the SGRP on that DBNG-UP - in which case SGRP deletion would be followed by an SGRP Create with the new Logical Port

6.7.7 BBF UP subscriber prefix

Table 50 specifies the “BBF UP subscriber prefix” IE in line with Figure 80.

Table 50: BBF UP Subscriber Prefix

| Octet 1 and 2 | | BBF UP Subscriber Prefix IE Type = 32801 (decimal) | |
|------------------------------|----------|---|---|
| Octets 3 and 4 | | Length = n | |
| Octets 5 and 6 | | Enterprise ID 3561 | |
| Information elements | P | Condition / Comment | IE Type |
| BBF SGRP ID | M | This IE must be present when BBF IP prefix is specified | BBF SGRP ID Details in section 6.9.15 |
| BBF IPv4 prefix | C | If present, this IE shall describe an IPv4 prefix assigned to the UP. If not present, at least one BBF IPv6 Prefix IE MUST be present. Several IEs with the same IE type may be present to provision a list of framed routes for example. | BBF IPv4 prefix Details in section 6.9.19 |
| BBF IPv6 prefix | C | If present, this IE shall describe an IPv6 prefix assigned to the UP. If not present, at least one BBF IPv4 Prefix IE MUST be present. Several IEs with the same IE type may be present to provision a list of framed routes for example. | BBF IPv6 prefix Details in section 6.9.20 |
| Network Instance | O | If present, this IE shall identify the Network Instance for the prefix(es) above. If this is not present, the Network Instance will be the default network instance. | Network Instance |
| BBF Active Prefix Tag | O | The IE may be present to associate the Active SGRP's Prefix(es) to a Tag. Only one Active Prefix Tag may be present, and its Usage field must be set to 'Active SGRP'. If not present, it is assumed the DBNG-UP knows how to advertise the prefix(es), when it is active, | BBF Prefix Tag Details in section 6.9.21 |
| BBF Backup Prefix Tag | O | The IE may be present to associate the Backup SGRP's Prefix(es) to a Tag. Only one Backup Prefix Tag may be present, and its Usage field must be set to 'Backup SGRP'. If not present, it is assumed the DBNG-UP knows how to advertise the prefix(es), when it is backup, taking also into account the «Standby Route Advertisement State» of the SGRP | BBF Prefix Tag Details in section 6.9.21 |

Note:

When requiring tags and network instance for prefix(es), different grouped IEs must be used

In case the SGRP state is set as Active or Standby, the UP, to decide whether to advertise the prefixes with "BBF Active Prefix Tag" or "BBF Backup Prefix Tag", shall consider the state assigned by the CP.

In case the SGRP state is set to Track-logical-port, the UP, to decide whether to advertise the prefixes with "BBF Active Prefix Tag" or "BBF Backup Prefix Tag", shall consider the operational condition of the SGRP, as defined in 4.4.10 and section 6.7.6.

If the Operational Condition is Up and the BBF Active prefix tag is present, then the UP must use the BBF Active Prefix Tag.

If the Operational Condition is Down and the BBF back prefix tag is present, then the UP must use the BBF Backup Prefix Tag.

If the Operational Condition is Not-ready, then the DBNG-UP may choose not to advertise based on conditions."

6.7.7.1 BBF UP Subscriber Prefix Modify

BBF UP Subscriber Prefix modify can be used to do the following actions:

- changing "BBF Active Prefix Tag"
- changing "BBF Backup Prefix Tag"

The change of both tags consists in overwriting the previous tag. If a tag has to be withdrawn, the prefix must be deleted and recreated without tag (see section 6.9.21). If a tag has to be added, the prefix can be updated using a "BBF UP Subscriber Prefix modify" where the tag is already not present.

The other fields of "BBF UP subscriber prefix" cannot be modified: therefore, if a change is required for those, the subscriber prefix must be deleted and re-created on the UPs that had been assigned that SGRP either as active or as backup.

In the "BBF UP Subscriber Prefix modify" message the complete list of prefixes must be sent even when only a subset of the prefixes in the list has to be updated, so that the DBNG-UP can successfully process the Modify message.

6.7.7.2 BBF UP Subscriber Prefix Deletion

When DBNG-CP needs to delete one or more prefixes, the BBF-Node-Info-Delete will contain the prefixes to be deleted. This operation must be always executed on all the UPs where the SGRP is either active or in backup state.

"BBF UP Subscriber Prefix delete" is required in order to:

- change the "Network Instance" for a given prefix, in which case the UP Subscriber Prefix "Deletion" will be followed by a Subscriber Prefix "Create" with the new Network instance.
- withdraw "BBF Active Prefix Tag" or "BBF Backup Prefix Tag", or both

The subscriber sessions associated with the deleted prefix are assumed to be terminated by the DBNG-CP prior to changing the network instance by subscriber prefix deletion/re-creation.

However, if the CP sends a Subscriber Prefix Deletion message to a UP before terminating the individual subscribers, the DBNG-UP should accept it: in this case, it should immediately withdraw the relevant route advertisement towards the core, while affected subscribers are terminated one-by-one by the DBNG-CP.

The BBF SGRP ID and the (list of) prefix(es) to be deleted must be always present in the BBF UP subscriber prefix deletion message; the Network Instance may be omitted, but in this case the "default Network Instance" is meant; the following IEs, if present, MUST be ignored:

- BBF Active Prefix Tag
- BBF Backup Prefix Tag.

6.8 PFCP Grouped IE extensions for Session Related Messages

This section highlights extensions required on grouped IEs such as PDR, PDI, FAR, to cover DBNG use cases. PFCP IEs in this section includes a presence requirement, indicated by the letter “P”:

- Mandatory (M) indicates that this IE MUST be present if used.
- Conditional (C) indicates that this IE is conditional based on use case specified.

6.8.1 PFCP Session Establishment Request

In the case where PPP LCP connectivity is required for the subscriber session, “PPP LCP Connectivity” IE is included in the PFCP session establishment request message.

In the case where an ACL is required to be applied for the subscriber session, “BBF ACL” grouped IE is included in the PFCP session establishment request message.

Table 51: BBF extended Information Element(s) in a PFCP Session Establishment Request

| Information element | P | Condition / Comment | IE Type |
|-----------------------------|---|---|---|
| PPP LCP connectivity | C | This IE MUST be present if periodic LCP echo hello is required. | PPP LCP connectivity Details in section 6.8.4 |
| BBF ACL | C | ACL/filter name. This IE MUST be present if an ACL or an ACL Chain needs to be attached to Subscriber session. Several IEs within the same IE type may be present for different family and direction combinations. | BBF ACL Details in section 6.8.7 |

6.8.1.1 Create PDR

The create PDR grouped IE should include BBF Outer Header Removal to remove various wireline encapsulations.

Table 52: BBF extended Create PDR IE(s) within PFCP Session Establishment Request

| Information element | P | Condition / Comment | IE Type |
|---------------------------------|---|---|---|
| Octet 1 and 2 | | Create PDR IE Type = 1(decimal) | |
| Octets 3 and 4 | | Length = n | |
| BBF Outer Header Removal | C | This IE MUST be present if the DBNG-UP function is required to remove header(s) from the packets matching this PDR. | BBF Outer Header Removal Details in section 6.9.3 |

6.8.1.2 PDI

The PDI IE should include L2TP type IE in the case of supporting L2TP subscribers.

Table 53: BBF extended PDI IE within PFCP Session Establishment Request

| Octet 1 and 2 | | PDI IE Type = 2 (decimal) | |
|---------------------|---|---|---|
| Octets 3 and 4 | | Length = n | |
| Information element | P | Condition / Comment | IE Type |
| L2TP type | C | This IE MUST be present if identification of the L2TP type control or data is required. | L2TP type Details in section 6.9.12 |

6.8.1.3 Ethernet Packet Filter

3GPP TS 29.244 [28] already defines a grouped IE named "Ethernet Packet Filter" which contains a list of IEs to specify properties of the filter. To cover the wireline case, further extensions are required to allow the matching of sub-flows of a traffic endpoint. Table 54 defines these extensions.

Table 54: BBF extended Ethernet Packet Filter IE(s) within PFCP Session Establishment Request

| Octet 1 and 2 | | Ethernet Packet Filter IE Type = 132 (decimal) | |
|------------------------|---|---|---|
| Octets 3 and 4 | | Length = n | |
| Information element | P | Condition / Comment | IE Type |
| PPP Protocol | C | If present, this IE MUST identify the PPP protocol to match for the incoming packet. (See section 6.9.6 for IE details) Multiple IEs may be present | PPP Protocol Details in section 6.9.6 |
| BBF C-Tag Range | C | If present, this IE MUST identify a range of C-tags to match for the incoming packet. If the logical port supports only single tagged vlans and vlan filtering is required, this IE MUST be present. Multiple IEs may be present to indicate multiple ranges. In this case the ranges MUST NOT overlap. | BBF C-Tag Range Details in section 6.9.28 |
| BBF S-Tag Range | C | If present, this IE MUST identify a range of S-tags to match for the incoming packet. If the logical port supports only single tagged vlans, this IE MUST not be present. Multiple IEs may be present to indicate multiple ranges. In this case the ranges MUST NOT overlap. | BBF S-Tag Range Details in section 6.9.29 |

Please note the following:

- "C-Tag Range" and "S-Tag Range" IEs do not redefine "C-Tag" and "S-Tag" IEs in 3GPP TS 29.244 [28], but they define additional IEs.
- When both S-Tag Range and C-Tag Range are specified. The Filter will create all combinations of the specified S-Tag and C-Tag Ranges. i.e., S-Tag contains range [A and B] and C-Tag contains range [X and Y], the system will match on (A,X), (A,Y), (B,X), and (B,Y).
- Multiple Ethernet Packet Filter IEs are used to separate and limit the ranges. i.e., using the same example above, to achieve range (A,X) and (B,Y). The first Ethernet Packet Filter IE would contain S-Tag range A and C-Tag range X only and the second one would contain only S-Tag range B and C-Tag range Y.

6.8.1.4 Forwarding Parameters

The Forwarding Parameters IE in FAR can include the BBF Outer Header Creation IE and the MTU IE. The BBF Outer Header Creation is used to encapsulate the subscriber data packet in various wireline encapsulations. The MTU IE is primarily used in the case of PPPoE.

Table 55: BBF extended Forwarding Parameters IE in FAR

| Octet 1 and 2 | Forwarding Parameters IE Type = 4 (decimal) | | |
|----------------------------------|---|--|--|
| Octets 3 and 4 | Length = n | | |
| Information elements | P | Condition / Comment | IE Type |
| BBF Outer Header Creation | C | This IE MUST be present if the DBNG-UP function is required to add outer header(s) to the outgoing packet. | BBF Outer Header Creation Details in section 6.9.3 |
| MTU | C | This IE MUST be present to enforce an MTU on outgoing packets. In the case of PPPoE, this may be based on negotiated MRU value | MTU Details in section 6.9.9 |

6.8.1.5 Create Traffic Endpoint

3GPP TS 29.244 [28] already defines a grouped IE named "Create Traffic Endpoint", which contains a list of IEs to specify properties of the endpoint, such as the GTP tunnel TEID or the subscriber IP address. To cover the wireline case, further extensions are required to describe the endpoint. Moreover, the use of some IEs already defined by 3GPP are to be constrained. Table 56 defines the extensions and constraints needed for the wireline case.

As defined in Section 6.2, each subscriber PFCP session utilizes four PDR rules: two PDR rules to forward control packets bi-directionally and other two PDR rules to forward data traffic bi-directionally. The PDR utilizes a combination of Traffic Endpoint and Filter rules to differentiate between control or data packets.

Traffic Endpoint contains combination of IE which includes port, VLAN tag(s), source MAC address, and if available an IP address to uniquely identify a subscriber service. The filters assist in the differentiation between control traffic and data traffic within a Traffic Endpoint. The filters allow packet match on Ethernet and/or IP information, for example, Ethertype and IP source ports.

The source MAC address within a Traffic Endpoint IE is a key identifier for two purposes:

1. Differentiate individual subscribers
2. Broadband service anti-spoof filter

The destination MAC address within a Traffic Endpoint MUST NOT be used.

If the DBNG-CP needs to override the default gateway MAC or PPPoE Access Concentrator MAC, this must be done by setting a virtual MAC on the SGRP.

For routed broadband service, the RG is a unique device with a single MAC address. For bridge broadband service, devices do not typically utilize contiguous source MAC addresses. For that reason, signaling a range of source/destination MAC addresses MUST NOT be used in the Traffic Endpoint.

Table 56: BBF extended Create Traffic Endpoint IE(s) within PFCP Session Establishment Request

| Octet 1 and 2 | | Create Traffic Endpoint IE Type = 127(decimal) | |
|--|---|--|---|
| Octets 3 and 4 | | Length = n | |
| Information elements | P | Condition / Comment | IE Type |
| The following IEs are defined in 3GPP TS 29.244 [28] but are not defined to be part of the Traffic Endpoint IE in that specification. | | | |
| MAC address | C | If present, this IE MUST be used to identify the MAC address of the subscriber where the SOUR bit MUST be set. All other bits MUST NOT be set (see NOTE1 below and 3GPP TS 29.244 [28] for IE details) | MAC address |
| C-Tag | C | If present, this IE MUST be used to identify the customer VLAN Tag of the traffic endpoint (see 3GPP TS 29.244 [28] for IE details) | C-Tag |
| S-Tag | C | If present, this IE MUST be used identify the service VLAN Tag of the traffic endpoint (see 3GPP TS 29.244 [28] for IE details) | S-Tag |
| BBF Extended IEs below | | | |
| Logical Port | C | If present, this IE MUST be used to provide an opaque byte string obtained from the NSH header to indicate the logical port for the subscriber. (see section 6.9.2 for IE details) | Logical port Details in section 6.9.2 |
| PPPoE Session ID | C | If present, this IE MUST be used to identify the PPPoE session ID of the subscriber. (see section 6.9.5 for IE details) | PPPoE Session ID Details in section 6.9.5 |
| L2TP tunnel | C | If present, this IE MUST be present if a L2TP tunnel is required. (see section 6.8.6 for IE details) | L2TP Tunnel Details in section 6.8.6 |
| NOTE1: The MAC address IE contains 4 fields: SOUR, DEST, USOU, and UDES. The SOUR MUST be set; DEST, USOU, and UDES MUST NOT be set. | | | |

6.8.2 PFCP Session Establishment Response

In the case of programming failures on DBNG-UP, the BBF Failure Cause IE MAY be included in the response to indicate the type of failure.

Table 57: BBF extended Information Element(s) in a PFCP Session Establishment Response

| Information element | P | Condition / Comment | IE Type |
|----------------------|---|--|---|
| BBF Failure Cause IE | O | This IE MAY be present in case of programming failure encountered on DBNG-UP | BBF Error Code IE Defined in sec 6.9.22 |

6.8.3 PFCP Session Modification Request

In the case where PPP LCP connectivity is required for the subscriber session, “PPP LCP Connectivity” IE is included in the PFCP session modification message.

In the case where an ACL is required to be changed for the subscriber session, “BBF ACL” grouped IE is included in the PFCP session modification request message.

Table 58: BBF extended Information Element(s) in a PFCP Session Modification Request

| Information element | P | Condition / Comment | IE Type |
|-----------------------------|---|--|---|
| PPP LCP connectivity | C | This IE MUST be present if periodic LCP echo hello is required and any of the LCP Connectivity parameters has changed. | PPP LCP connectivity Details in section 6.8.4 |
| BBF ACL | C | ACL/filter name. This Shall be present if an ACL or an ACL Chain needs to be attached for Subscriber session or if an ACL or ACL chain needs to be removed. Several IEs within the same IE type may be present for different family and direction combinations. | BBF ACL Details in section 6.8.7 |

6.8.3.1 Update Forwarding Parameters

The Update forwarding Parameters IE in FAR can include the MTU IE. The MTU IE is primarily used in the case of PPPoE and can for example be included when a new MRU is negotiated during a session lifetime.

Table 59: BBF extended Update Forwarding Parameters IE(s) in FAR

| Octet 1 and 2 | Update Forwarding Parameters IE Type = 11 (decimal) | | |
|---------------------|---|--|--|
| Octets 3 and 4 | Length = n | | |
| Information element | P | Condition / Comment | IE Type |
| MTU | C | This IE MUST be present to enforce an MTU on outgoing packets. In the case of PPPoE, this may be based on negotiated MRU value. This IE MUST only be provided if it is changed. | MTU Details in section 6.9.9 |

6.8.3.2 Ethernet Packet Filter

Table 60: BBF extended Ethernet Packet Filter IE(s) within PFCP Session Modification Request

| Octet 1 and 2 | Ethernet Packet Filter IE Type = 132 (decimal) | | |
|---------------------|--|---|---|
| Octets 3 and 4 | Length = n | | |
| Information element | P | Condition / Comment | IE Type |
| C-Tag Range | C | If present, this IE MUST identify a range of C-tags to match for the incoming packet. Multiple IEs may be present to indicate multiple ranges. In this case the ranges MUST NOT overlap. | C-Tag Range Details in section 6.9.28 |
| S-Tag Range | C | If present, this IE MUST identify a range of S-tags to match for the incoming packet. Multiple IEs may be present to indicate multiple ranges. In this case the ranges MUST NOT overlap. | S-Tag Range Details in section 6.9.29 |

As per 3GPP TS 29.244 [28] if any modification to a PDI, including its nested IEs such as S-Tag and C-Tag ranges, is required, the full PDI and thus the full list of ranges must be sent again.

The modified S-Tag and C-Tag ranges takes effect immediately once applied to the DBNG-UP.

It is recommended that the DBNG-CP should not modify the S-Tag and C-Tag ranges for a logical port when subscribers are still active on the logical port. If there are ongoing control message transactions on the per-logical port CPRi, within the modified S-Tag and C-Tag range, then these transactions may be impacted by the new filter rules.

6.8.4 PFCP Session Modification Response

In the case of programming failures on DBNG-UP, the BBF Failure Cause IE MAY be included in the response to indicate the type of failure.

Table 61: PFCP Session Modification Response

| Information element | P | Condition / Comment | IE Type |
|-----------------------------|---|--|--|
| BBF Failure Cause IE | O | This IE MAY be present in case of programming failure encountered on DBNG-UP | BBF Error Code IE Defined in sec 6.9.22 |

6.8.5 PPP LCP Connectivity

The table below is a new BBF specified grouped IE used to specify PPP LCP connectivity check parameters.

In the case where LCP echo hello is offloaded on the DBNG-UP, the LCP echo hellos MUST not be redirected to the DBNG-CP.

Table 62: PPP LCP Connectivity

| Octet 1 and 2 | | PPP LCP Connectivity IE Type = 32780 | |
|-----------------------------|---|---|---|
| Octets 3 and 4 | | Length = n | |
| Octets 5 and 6 | | Enterprise ID 3561 | |
| Information elements | P | Condition / Comment | IE Type |
| Traffic Endpoint ID | M | Identifies the context on which connectivity verification must be performed. | Traffic Endpoint ID |
| Verification Timers | C | This IE MUST be used to indicate the frequency and number of retries for verification messages. If this IE is not present, no periodic verification is started. | Verification Timers Details in section 6.9.7 |
| PPP LCP Magic Number | C | If present this IE MUST be used to indicate which magic number to use when generating keepalives and which magic number to verify incoming keepalives against. | PPP LCP Magic Number Details in section 6.9.8 |

6.8.6 L2TP Tunnel

In the case where L2TP tunnel required for the subscriber session, this IE is included.

Table 63: L2TP Tunnel

| Octet 1 and 2 | L2TP Tunnel IE Type = 32781 | | |
|----------------------|-----------------------------|---|--|
| Octets 3 and 4 | Length = n | | |
| Octets 5 and 6 | Enterprise ID 3561 | | |
| Information elements | P | Condition / Comment | IE Type |
| L2TP tunnel endpoint | C | If present, this IE MUST be used to identify the L2TP tunnel ID and IP information. | L2TP Tunnel Endpoint Details in section 6.9.10 |
| L2TP Session ID | C | If present, this IE MUST be used to identify the L2TP session ID. | L2TP session ID Details in section 6.9.11 |

6.8.7 BBF ACL IE

Table 64 defines BBF ACL Grouped IE. This is used to manage ACL operations for a subscriber session. If ACL is preconfigured on DBNG-UP, DBNG-CP uses this IE to program ACL on subscriber session. Subsequent Session modification messages containing this IE can modify already applied ACL(s) on the subscriber session.

Table 64: BBF ACL

| Octet 1 and 2 | BBF ACL IE Type = 32802 (decimal) | | |
|----------------------|--|---|--|
| Octets 3 and 4 | Length = n | | |
| Octets 5 and 6 | Enterprise ID 3561 | | |
| Information elements | P | Condition / Comment | IE Type |
| BBF ACL Name | C | This IE specifies the name of the ACL to be applied to subscriber. This IE is MUST for attaching ACL to the subscriber session. Several IEs may be present to attach multiple ACLs to the subscriber session. ACL attachment follows the order in which Names are received on DBNG-UP. DBNG-CP MUST ensure the order for proper attachment of ACLs on DBNG-UP. If this IE is not present then the corresponding ACL will be removed based on direction and family | Activate Predefined Rules (3GPP TS29.244 sec 8.2.72) |
| BBF ACL Direction | M | This IE specifies ACL Direction for a specific subscriber session IN, OUT | BBF Direction Details in section 6.9.13 |
| BBF ACL Family | M | This IE specifies the Address-family for ACL, possible values: I2-eth, IPv4, IPv6, IPv46 ³⁴ | BBF Family Details in section 6.9.14 |

³Expects the DBNG-UP to have same named ACL applied to both IPv4 and IPv6 families

⁴If family "Ipv46" is used, no other ACL for family "Ipv4" or family "Ipv6" may be present for the same direction and all subsequent ACL operations should be done with same family (i.e., IPv46) and same direction value for entire session lifecycle.

6.9 BBF PFCP IE extensions

This section highlights required IE extensions to cover MS-BNG use cases. Please refer to 3GPP TS 29.244 [28] section 8.1.1 for further information on vendor specific extensions. Below is the 3GPP specified IE format for Vendor Specific IE.

Figure 82 depicts the format of a vendor-specific Information Element, which content is not specified, and the IE Type value MUST be within the range of 32768 to 65535. From m to (m+4) are octets which can be defined by BBF for future uses.

| Octets | Bits | | | | | | | |
|------------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = xxx (decimal) | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID | | | | | | | |
| 7 to (n+4) | IE specific data or content of a grouped IE | | | | | | | |
| m to (m+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 82: 3GPP Vendor-Specific Information Element Format Reference

Many IEs have spare values reserved for future extensions; some examples include:

- Spare bits in octets used as bitfields
- Spare octets at the end of a variable-length IE, usually denoted with a “These octet(s) is/are present only if explicitly specified” statement.
- Unused/undefined enumeration values.

A sending PFCP entity complying to this specification MUST NOT include spare octets or set unused enumeration values, it MUST also set any spare bits in bitfields to ‘0’. A receiving PFCP entity MUST ignore any unknown spare octets, enumeration values, or bits.

For more details on forward and backward compatibility, see section 6.10.

6.9.1 BBF UP Function Features

The BBF UP Function Features IE indicates the features supported by the DBNG-UP function and MUST be coded as depicted in Figure 83.

| Octets | Bits | | | | | | | |
|-------------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32768 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 to 8 | Supported-Features | | | | | | | |
| 9 to 10 | Additional Supported-Features 1 | | | | | | | |
| 11 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 83: BBF UP Function Features

The BBF UP Function Features IE takes the form of a bitmask where each bit set indicates that the corresponding feature is supported. Undefined bits MUST be set to zero by senders and MUST be ignored by the receiver.

Supported-Features

When present, MUST be encoded as per Table 65 which specifies the features defined on the DBNG-UP.

Table 65: BBF UP Function Features

| Feature Octet / Bit | Feature | Description |
|---------------------|-----------------------------|---|
| 7/1 | PPPoE | <p>Informes the DBNG-CP that the DBNG-UP supports PPPoE.</p> <p>Setting this flag also implies the DBNG-UP ability to support immediate session creation model for PPPoE sessions.</p> |
| 7/2 | IPoE | <p>Informes the DBNG-CP that the DBNG-UP supports IPoE.</p> <p>Setting this flag also implies the DBNG-UP ability to support immediate session creation model for IPoE sessions.</p> |
| 7/3 | LAC | <p>Informes the DBNG-CP that the DBNG-UP supports LAC.</p> <p>Setting this flag also implies the DBNG-UP ability to support immediate session creation model for LAC sessions.</p> |
| 7/4 | LNS | <p>Informes the DBNG-CP that the DBNG-UP supports LAC.</p> <p>Setting this flag also implies the DBNG-UP ability to support immediate session creation model for LNS sessions.</p> |
| 7/5 | LCP keepalive offload | <p>Informes the DBNG-CP that the DBNG-UP supports PPP LCP echo.</p> |
| 8/1 | Issue 2 Minimum Feature Set | <p>This bit informs the DBNG-CP that the following features are supported:</p> <ul style="list-style-type: none"> - Per Logical Port CPR tunnel: Indicates the DBNG-UP additional ability to support bi-directional forwarding rules per logical port CPR tunnel. This also includes S-Tag and C-Tag range filtering. Note: Delayed PFCP session setup utilizes per logical port CPR tunnel. - NSH header insertion: Informes the DBNG-CP that the DBNG-UP supports NSH header insertion on any CPR tunnel - SGRP support: as defined in this document. Note this includes: SGRP Subscriber prefix and SGRP resiliency. <p>Additionally, one of the following feature flags MUST also be set:</p> <ul style="list-style-type: none"> - IPoE - PPPoE - LAC - LNS <p>This bit must always be set by a DBNG-UP compliant to this version of the document</p> |
| 8/2 | Volume quota and thresholds | <p>Informes the DBNG-CP that the DBNG-UP supports monitoring subscriber volume quota and thresholds.</p> <p>This bit can be set only when "Issue 2 Minimum Feature Set" bit is also set</p> |
| 8/3 | ACL | <p>Informes the DBNG-CP that the DBNG-UP supports ACL.</p> <p>Maximum ACL Chain Length IE MAY be included along with BBF Function Feature IE if ACL chaining is supported.</p> <p>If Maximum ACL Chain Length IE is not present, then DBNG-UP does not support ACL Chaining.</p> <p>This bit can be set only when "Issue 2 Minimum Feature Set" bit is also set</p> |

| | | |
|-----|-----------------------------------|--|
| 8/4 | Server Control Packet Redirection | <p>This bit indicates the UP supports a redirect tunnel per Network Instance in addition to access CPRi rules. For more details see section 6.4.3. This can for example be used to redirect packets to and from a DHCP server in the DHCP relay case.</p> <p>This bit can be set only when "Issue 2 Minimum Feature Set" bit is also set</p> |
|-----|-----------------------------------|--|

Note:

- For IPoE, PPPoE, LAC, and LNS, both IPv4 and IPv6 are supported.
- Bit 7/6 is defined in TR-459.3, bits 7/7 and 7/8 are defined in TR-459.2
- Once any BBF UP Function Features bit has been set, it cannot be unset in a subsequent PFCP Association Update Request or Response.

In addition to the above, there are some key UP features that do not require any feature flag because 1) it is covered by 3GPP feature flags and 2) are architectural requirements:

- Accounting call flows and use case specified in section 4.5.46, the required feature 3GPP flags are specified in Table 17 and Table 18

6.9.2 Logical Port

The Logical Port IE MUST be encoded as Figure 84.




Figure 84: Logical Port

Logical-port-id

Encode a logical-port-id, which is an opaque byte string to indicate the logical port on which Ethernet traffic is received. This should not contain S-Tag or C-Tag values, which are signaled more explicitly in PFCP described in 3GPP TS 29.244 [28].

6.9.3 BBF Outer Header Creation

The BBF Outer Header Creation indicates a header is to be added to the packet before forwarding and MUST be coded as depicted in Figure 85. The BBF Outer Header also contains parameters to construct the header. This IE can be used in combination with the 3GPP TS 29.244 [28] Outer Header Creation IE (Type 84). If both are present, both headers are added, the 'Outer Header Creation IE' will be the top header and "BBF Outer Header" will be the next (lower) header.

Note: Some examples of where both headers will be used are during GTP tunneling of control packets and L2TP tunneling of data packets.

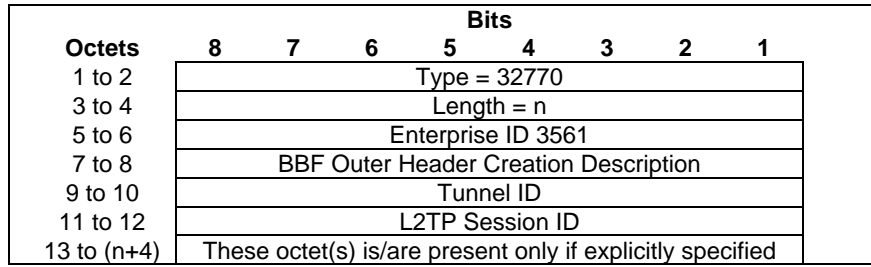


Figure 85: BBF Outer Header Creation

BBF Outer Header Creation Description

When present, MUST be encoded as specified in Table 66. It takes the form of a bitmask where each bit indicates the outer header to be created in the outgoing packet. Undefined bits MUST be zeroed on transmission and MUST be ignored by the receiver.

Table 66: BBF Outer Header Creation Description

| Octet/bit | Outer Header to be created in the outgoing packet |
|-----------|---|
| 7/1 | CPR-NSH |
| 7/2 | Traffic-Endpoint |
| 7/3 | L2TP |
| 7/4 | PPP |

Note: At least one bit of the Outer Header Creation Description field MUST be set to 1.

- CPR-NSH: NSH header insertion is controlled by the DBNG-CP. Based on the BBF UP function feature flag “Issue 2 Minimum Feature Set”, the DBNG-CP may set this bit when establishing per-logical-port or per-session CPR tunnels. An NSH header defined in RFC 8300 [47] will be attached to the control packet, which contains meta data for the logical port. The NSH encapsulated control packet tunneled over a GTP-u tunnel utilizing the IE Outer Header Creation. For more information on NSH, look at section 6.9.3.1.
- Traffic-Endpoint: The header creation for the packet will be based on the IE Linked Traffic Endpoint within the same FAR. Further detail:
 - This is used for layer 2 traffic forwarding. The traffic endpoint specified must contain a logical port and a MAC address. Optionally, the traffic endpoint can also contain S-Tag, C-Tag and PPPoE Session ID. Note: The Linked Traffic Endpoint always refer to traffic endpoint that flows in the opposite direction, therefore the source and destination MAC address must always be swapped when reconstructing the Ethernet header.
- L2TP – Creates the L2TP header with indicated tunnel ID and session ID. For LAC, this is used to encapsulate the PPP packet with a L2TP header before forwarding to LNS. For LNS, this is used in combination with PPP to encapsulate the IP packet with both PPP and L2TP before forwarding to the LAC.
- PPP – Creates the PPP data packet header.

Tunnel ID

Indicates the L2TP Tunnel ID. This field is only present if the L2TP bit in BBF Outer Header Creation Description is set

Session ID

Indicates the L2TP Session ID. This field is only present if the L2TP bit in BBF Outer Header Creation Description is set

6.9.3.1 NSH header information

In order to signal the logical port and the local DBNG-UP port MAC, Ethernet packets will be encapsulated using an MD-type 2 NSH header as defined in RFC 8300 [47]. This header is shown in Figure 86.

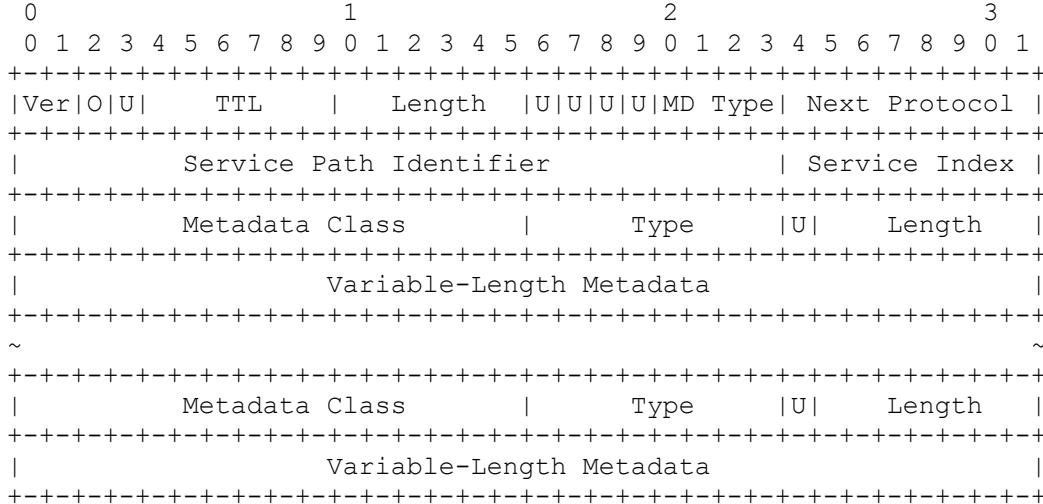


Figure 86: NSH header information

Where OAM MUST be set to 0, TTL MUST be set to 1, MD Type MUST be set to 2 and Next Protocol MUST be set to 0x3 (Ethernet). Service Path Identifier is initially not used and MUST be set always to 0, Service Index MUST be set to 255. Version and (NSH) length are per RFC 8300 [47].

Specific Metadata types and information:

- Metadata Class = 0x0200
- Logical port (Type=0, Length=N): an opaque byte string identifying an access context on a DBNG-UP (e.g., port, lag, Ethernet tunnel, ...)
- MAC (Type=1, Length=6): The local DBNG-UP MAC associated with the logical port
- Network Instance (Type=2, Length=N): an opaque byte string identifying the Network Instance on which this control packet arrived on the DBNG-UP.
- Interface Identifier (Type = 3, Length = 8): An interface identifier of the link-local address for which the DBNG-UP will locally answer NS requests. The DBNG-CP can use this to derive a link-local address when generating IPv6 packets (e.g., RA) over CPRI.

6.9.4 BBF Outer Header Removal

The BBF Outer Header Removal IE MUST be encoded as Figure 87:




Figure 87: BBF Outer Header Removal

BBF Outer Header Removal Description

This field MUST be an 8-bit unsigned integer and indicates which headers need to be removed from an incoming packet. The values are as defined as follows:

Table 67: BBF Outer Header Removal Description

| Outer Header to be removed in the incoming packet | Value (Decimal) |
|---|-----------------|
| Ethernet | 1 |
| PPPoE / Ethernet | 2 |
| PPP / PPPoE / Ethernet | 3 |
| L2TP | 4 |
| PPP / L2TP | 5 |

- Ethernet: Removes the Ethernet header including S-Tags and C-Tags.
- PPPoE / Ethernet: Removes the PPP header and Ethernet header including S-Tags and C-Tags.
- PPP / PPPoE / Ethernet: Removes the PPP header, the PPPoE header, and the Ethernet header including S-Tags and C-Tags.
- L2TP: Removes only the L2TP header
- PPP/L2TP: Removes the PPP and the L2TP header together.

6.9.5 PPPoE Session ID

The PPPoE Session ID IE MUST be encoded as Figure 88.




Figure 88: PPPoE Session ID

PPPoE Session ID

Encode a PPPoE Session ID as specified in RFC 2516 [34].

6.9.6 PPP Protocol

The PPP Protocol IE MUST be encoded as Figure 89.

| Octets | Bits | | | | | | | |
|-------------|--|---|---|---|---------|------|----------|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32773 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 | Spare | | | | control | data | specific | |
| 8 to 9 | protocol | | | | | | | |
| 10 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 89: PPP Protocol

Exactly one flag **MUST** be set in octet 7. If more than one bit is set, this is an error and is handled according to 3GPP TS 29.244 [28] section 7.6

specific

Indicates a specific protocol value must be matched, further specified in the IE

data

Indicates any protocol value where the most significant bit equals zero, as per RFC 1661 [32].

control

Indicates any protocol value where the most significant bit equals one, as per RFC 1661 [32].

protocol

Octets “8 to 9” are only present if the specific bit is set and encode a valid PPP protocol value as assigned by Internet Assigned Numbers Authority.

6.9.7 Verification Timers

The Verification Timers IE **MUST** be encoded as Figure 90.

| Octets | Bits | | | | | | | |
|-------------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32774 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 to 8 | interval | | | | | | | |
| 9 | count | | | | | | | |
| 10 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 90: Verification Timers

interval

Specify an unsigned 16-bit interval in 10 milli-seconds that indicates how frequent the verification procedure is started.

count

Specifies an unsigned 8-bit integer on how many unanswered consecutive keepalive messages should be sent before connection is considered down.

6.9.8 PPP LCP Magic Number

The PPPoE LCP Magic Number IE **MUST** be encoded as Figure 91.

| Octets | Bits | | | | | | |
|-------------|--|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 |
| 1 to 2 | Type = 32775 | | | | | | |
| 3 to 4 | Length = n | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | |
| 7 to 10 | Tx Magic Number | | | | | | |
| 11 to 14 | Rx Magic Number | | | | | | |
| 15 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | |

Figure 91: PPP LCP Magic Number

Tx Magic Number

Encode a PPP LCP Magic Number as defined in RFC 1661 [32]. This is the magic number used when transmitting LCP keepalive messages.

Rx Magic Number

Only present if length >=10 and encode a PPP LCP Magic Number as defined in RFC 1661 [32].

When present, received LCP keepalive messages need to verify against this magic number. When not present (length < 10), magic numbers are not verified.

6.9.9 MTU

This IE MUST be encoded as indicated in Figure 92.

| Octets | Bits | | | | | | |
|------------|--|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 |
| 1 to 2 | Type = 32776 | | | | | | |
| 3 to 4 | Length = n | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | |
| 7 to 8 | MTU value | | | | | | |
| 9 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | |

Figure 92: MTU

MTU

The MTU MUST be encoded as an unsigned 16-bit integer value. Packets exceeding the MTU must either be fragmented (IPv4 without DF bit set RFC 791 [31]) or answered with an ICMP/ICMPv6 “Packet Too Big” error code (IPv4 with DF bit set RFC 791 [31], IPv6 RFC4443 [41]). MTU is applied on IP level, before any outer header or Ethernet encapsulation. A UP may apply a lower MTU if the associated forwarding construct requires so.

Note: The system MUST have a mechanism to control the rate of sending of ICMP Error Messages.

6.9.10 L2TP Tunnel Endpoint

This IE MUST be encoded as Figure 93.

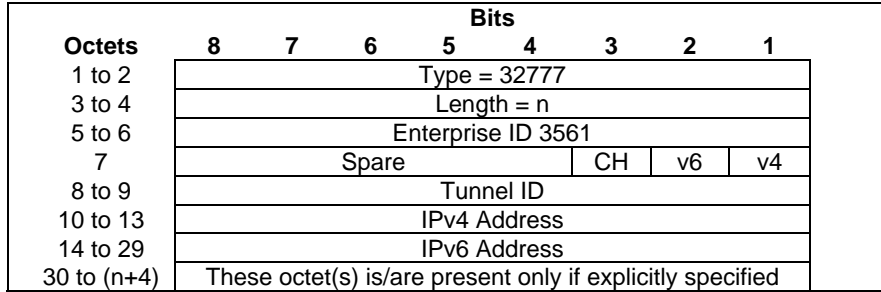


Figure 93: L2TP Tunnel Endpoint

Flags v4, v6, and CH are mutually exclusive.

v4

Indicates an IPv4 address is included

v6

Indicates an IPv6 address is included

CH

Indicates no IP is included, and the DBNG-CP function sets the CHOOSE (CH) bit to 1 if the DBNG-UP function supports the allocation of L2TP tunnel IP address and the DBNG-CP function requests the DBNG-UP function to assign the L2TP tunnel IP to the Traffic Endpoint.

Tunnel ID

Specifies the L2TP tunnel ID to match

IPv4 address

Specifies the L2TP IPv4 local terminating address

IPv6 address

Specifies the L2TP IPv6 local terminating address

6.9.11 L2TP Session ID

This IE MUST be encoded as specified in Figure 94.




Figure 94: L2TP Session ID

L2TP session ID

Specifies the L2TP session ID to match

6.9.12 L2TP type

This IE MUST be encoded as specified in Figure 95.

| Octets | Bits | | | | | | | |
|------------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32779 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 to 8 | Spare | | | | | | | T |
| 9 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 95: L2TP type

T

Identifies the l2tp type, 0 means l2tp data and 1 means l2tp control (RFC 2661 [35])

6.9.13 BBF Direction IE

| Octets | Bits | | | | | | | |
|------------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32803 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 | direction | | | | | | | |
| 8 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 96: BBF Direction IE

Octet 7 (direction) possible values:

0 means INPUT/INGRESS/UPSTREAM direction

1 means OUTPUT/EGRESS/DOWNSTREAM direction

6.9.14 BBF Family IE

| Octets | Bits | | | | | | | |
|------------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32804 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 | value | | | | | | | |
| 8 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 97: BBF Family IE

Octet 7 "value":

0 means "IPv4"

1 means "IPv6"

2 means "IPv46"

3 means "L2eth"

6.9.15 BBF SGRP Identifier

Subscriber Group (SGRP) Identifier must be encoded as in Figure 98 below.

| Octets | Bits | | | | | | | |
|-----------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32806 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 to 10 | SGRP-Identifier | | | | | | | |
| 11 to n+4 | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 98: BBF SGRP Identifier

SGRP-Identifier

SGRP identifier MUST encoded as an unsigned 32-bit integer value that uniquely identifies the Subscriber Group (SGRP).

6.9.16 BBF SGRP State

| Octets | Bits | | | | | | | |
|------------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32807 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 | Value | | | | | | | |
| 8 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 99: BBF SGRP State

The following enumerated values are within Octet 7:

- 0 - reserved
- 1 – Active: indicates SGRP State is “Active”
- 2 – Backup: indicates SGRP State is “Backup”
- 3 – Track-Logical-Port: SGRP State is set to “Track-Logical-Port”.

6.9.17 BBF SGRP Flags

The SGRP Flags IE contains the following flags:

- SGRP Route Advertisement State: refers the ability of the DBNG-UP acting as a backup for the Subscriber Group (SGRP) to optionally choose to advertise the routes towards the core. The SGRP Backup Route Advertisement State must be encoded as in Figure 100 below.
- SGRP Partial State Allowed: refers to the DBNG-UP resiliency being capable of creating partial state of subscriber or not. The SGRP Partial State Allowed must be encoded as in Figure 100 below.

| Octets | Bits | | | | | | | |
|------------|--|---|---|---|---|---|-----|-----|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32808 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 | Spare | | | | | | PSA | RAS |
| 8 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 100: BBF SGRP Flags

Octet 7 must be encoded as follows:

Bit 1 (RAS – Route Advertisement State): when set to “1” it indicates for the DBNG-UP to advertise routes when SGRP State is Backup.

Bit 2 (PSA – Partial State Allowed): when set to “1” it indicates partial state is allowed when SGRP is Backup for DBNG-UP.

6.9.18 BBF Operational Condition IE

| Octets | Bits | | | | | | | |
|------------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32809 (decimal) | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 | Value | | | | | | | |
| 8 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 101: BBF Operational Condition

Value: an enum indicating operational condition

- 0 - Up: indicates successful operation for SGRP on DBNG-UP (not applicable to track-logical port)
- 1 - Down: indicates SGRP is not operational due to a failure (applicable to all cases)
- 2 - Not-ready: indicates that the DBNG-UP is not-ready for state changes (applicable to all cases)
- 3 - Up-active: indicates successful operation for SGRP on DBNG-UP, SGRP is active (applicable only to track-logical-port)
- 4 - Up-backup: indicates successful operation for SGRP on DBNG-UP, SGRP is backup (applicable only to track-logical-port)

6.9.19 BBF IPv4 Prefix

BBF IPv4 Prefix must be encoded as in Figure 102 below.

| Octets | Bits | | | | | | | |
|-------------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32810 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 to 10 | IPv4 Prefix | | | | | | | |
| 11 | IPv4 Prefix Length | | | | | | | |
| 12 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 102: BBF IPv4 Prefix

IPv4 Prefix (4 octets)

Specifies the IPv4 subscriber prefix

IPv4 Prefix Length (1 octet)

Specifies the IPv4 subscriber prefix length. The IPv4 Prefix length shall be encoded as an 8 bits binary integer. Its value can be between 0 to 32. The prefix length value of "32" indicates an IPv4 host address.

An IPv4 default gateway is signaled this way.

6.9.20 BBF IPv6 Prefix

BBF IPv6 Prefix must be encoded as in Figure 103 below.

| Octets | Bits | | | | | | | |
|-------------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32811 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 to 22 | IPv6 Prefix | | | | | | | |
| 23 | IPv6 Prefix Length | | | | | | | |
| 24 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 103: BBF IPv6 Prefix

IPv6 Prefix (16 octets)

Specifies the IPv6 subscriber prefix

IPv6 Prefix Length (1 octet)

Specifies the IPv6 subscriber prefix length. The IPv6 Prefix Length shall be encoded as an 8 bits binary integer, e.g., if /72 prefix is used, the value shall be set to (decimal) 72, or if /56 prefix is used, the value shall be set to (decimal) 56. The prefix length value "128" indicates an individual /128 IPv6 address.

An IPv6 default gateway is signaled this way.

6.9.21 BBF Prefix Tag

BBF Prefix Tag must be encoded as in Figure 104 below.

| Octets | Bits | | | | | | | |
|-------------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32812 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 | BBF Prefix Tag Usage | | | | | | | |
| 8 to 11 | BBF Prefix Tag | | | | | | | |
| 12 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 104: BBF Prefix tag

BBF Prefix Tag Usage - In which conditions the tag is to be used, possible values:

- 0x0: Reserved
- 0x1 (Active SGRP): The tag is only to be used when the associated prefix is associated to an active SGRP.
- 0x2 (Backup SGRP): The tag is only to be used when the associated prefix is associated to a backup SGRP.

BBF Prefix tag

The tag, if present, MUST encoded as an unsigned 32-bit integer value that can be associated with a BBF IPv4 or IPv6 Prefix.

6.9.22 BBF Error Code

| Octets | Bits | | | | | | | |
|------------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32813 (decimal) | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 to 8 | Value | | | | | | | |
| 9 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 105: BBF Error Code

Value: an enum indicating error code, possible values:

- 0x0: Reserved
 - 0x1: SGRP programming error because of resource exhaustion
 - 0x2: SGRP programming error because of configuration mismatch or incomplete
 - 0x3: SGRP programming error related to vMAC
 - 0x4: SGRP Partial state is not supported on DBNG-UP
 - 0x5: SGRP Track logical port is not supported by DBNG-UP on the Logical ports specified for this SGRP
 - 0x6: ACL name not found
 - 0x7: Prefix programming error because of resource exhaustion
 - 0x8: Prefix programming error because of configuration mismatch or incomplete
 - 0x9: Prefix's Network Instance not found
 - 0xA: Prefix Active Tag matching not found
 - 0xB: Prefix Backup Tag matching not found
 - 0xC: Prefix's SGRP ID not found
 - 0xD: Prefix not found (on prefix modify or prefix delete)
 - 0xE: Prefix in use (when SGRP delete is performed before prefix delete)
 - 0xF: SGRP programming error because of ongoing SGRP deletion
- 0xffff: Other error

Note: if a prefix is not well formulated by the DBNG-CP, the DBNG-UP may include in the reply message a CAUSE IE specified in 3GPP TS 29.244 [28] section 8.2.1.

6.9.23 BBF Error Message

| Octets | Bits | | | | | | | |
|----------|-------------------------------------|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32814 (decimal) | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 to n+4 | Opaque Text String Describing Error | | | | | | | |

Figure 106: BBF Error Message

6.9.24 BBF Maximum ACL Chain Length

This IE indicates the length of ACL Chain possible on DBNG-UP. This IE MUST be present if the DBNG-UP supports ACL Chaining.

| Octets | Bits | | | | | | | |
|------------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32815 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 | ACL Chain Length | | | | | | | |
| 8 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 107: Maximum ACL Chain Length

ACL Chain Length (integer) indicates the maximum number of BBF ACL name entries that can be supported in an ACL Chain by DBNG-UP. For example, to indicate support for an ACL chain of maximum 3 entries {ACL1, ACL2, ACL3}, the DBNG-UP will send a value of 3.

6.9.25 BBF Forwarding Capability

The BBF Forwarding Capability IE is used to indicate the forwarding capability of a given logical-port.

| Octets | Bits | | | | | | | |
|------------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32816 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 | Forwarding Capability | | | | | | | |
| 8 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 108: BBF Forwarding Capability

Forwarding Capability indicates a percentage and may take binary coded integer values from and including 0 up to and including 100. Other values shall be considered as 0.

6.9.26 BBF Connectivity Status

The BBF Connectivity Status IE is used to indicate the connectivity status of a network instance.

| Octets | Bits | | | | | | | |
|------------|--|---|---|---|---|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32817 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 | Value | | | | | | | |
| 8 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 109: BBF Connectivity Status

Where Value is an enumeration as follows:

- 0 – reserved
- 1 – connected: The node reporting this status has connectivity to the rest of the network.
- 2 – isolated: The node reporting this status is isolated from the rest of the network.

6.9.27 Vendor-Specific Node Report Type

This Vendor Specific Node Report IE will use the BBF IANA Enterprise number “3561”. This IE indicates the type of BBF Node Report that the DBNG-UP sends to the DBNG-CP.

| Octets | Bits | | | | | | | |
|------------|--|-------|-------|-------|-------|-----|-----|-----|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32818 (decimal) | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 | Spare | Spare | Spare | Spare | Spare | NIR | SGR | LPR |
| 8 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 110: Vendor-Specific Node Report Type

Octet 7 shall be encoded as follows:

- Bit 1 – LPR (Logical Port Report): when set to "1", this indicates a Logical Port Report.
- Bit 2 – SGR (Subscriber Group Report): when set to "1", this indicates a SGRP Notification Report.
- Bit 3 – NIR (Network Instance Report): when set to "1", this indicates a Network Instance Report.

6.9.28 BBF C-Tag Range

This IE MUST be encoded as Figure 111.

| Octets | Bits | | | | | | | |
|-------------|--|---|---|---|-------------------|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32820 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 | start c-vid value | | | | | | | |
| 8 | end c-vid value | | | | start c-vid value | | | |
| 9 | end c-vid value | | | | | | | |
| 10 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 111: C-Tag Range

C-VID values are as specified in IEEE 802.1q.

- Octet 7/bit 8 shall be the most significant bit of the start c-vid value and octet 8/bit 1 shall be the least significant bit. See 3GPP 29.244 [28] clause 7.1 for more details on the encoding.
- Octet 8/bit 8 shall be the most significant bit of the end c-vid value and octet 9/bit 1 shall be the least significant bit. See 3GPP 29.244 [28] clause 7.1 for more details on the encoding.

The end c-vid value MUST be greater or equal to the start c-vid value.

It is expected that the actual Report corresponding to the marked bits will be with the same PFCP report message. Marking multiple bits is possible to allow the report to include multiple report types. It is also possible to include multiple report of the same type. i.e., multiple Logical Port Report.

6.9.29 BBF S-Tag Range

This IE MUST be encoded as Figure 112.

| Octets | Bits | | | | | | | |
|-------------|--|---|---|---|-------------------|---|---|---|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32821 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 | start s-vid value | | | | | | | |
| 8 | end s-vid value | | | | start s-vid value | | | |
| 9 | end s-vid value | | | | | | | |
| 10 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 112: S-Tag Range

S-VID values are as specified in IEEE 802.1q.

- Octet 7/bit 8 shall be the most significant bit of the start s-vid value and octet 8/bit 1 shall be the least significant bit. See 3GPP 29.244 [28] clause 7.1 for more details on the encoding.
- Octet 8/bit 8 shall be the most significant bit of the end s-vid value and octet 9/bit 1 shall be the least significant bit. See 3GPP 29.244 [28] clause 7.1 for more details on the encoding.

The end s-vid value MUST be greater or equal to the start s-vid value.

6.9.30 BBF Apply Action

The BBF Apply Action IE MUST be encoded as in Figure NN

| Octets | Bits | | | | | | | |
|------------|--|---|---|---|---|---|-----|-----|
| | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| 1 to 2 | Type = 32787 | | | | | | | |
| 3 to 4 | Length = n | | | | | | | |
| 5 to 6 | Enterprise ID 3561 | | | | | | | |
| 7 | | | | | | | RTR | NAT |
| 8 to (n+4) | These octet(s) is/are present only if explicitly specified | | | | | | | |

Figure 113: BBF Apply action Protocol

This IE acts as an extension to the 3GPP Apply Action IE to define new actions. Currently following additional bits are defined

NAT

Defined in TR 459.2

RTR (Reply-To-Request)

Indicates the DBNG-UP should reply to the matched packet locally if it was a request packet, otherwise it should be discarded

6.10 General DBNG-CP and DBNG-UP interoperability recommendations

There are instances where feature capabilities might not match between DBNG-CP and DBNG-UP. For example, software updates for new functional features, independent software upgrades on DBNG-CP and DBNG-UP, or necessary software rollback when software issues are discovered.

3GPP introduces basic protocol compatibility for PFCP in clause 7.6.1 of TS 29.244 [28] as follows:

The handling of unknown, unexpected or erroneous PFCP messages and IEs shall provide for the forward compatibility of PFCP. Therefore, the sending PFCP entity shall be able to safely include in a message a new conditional-optional or an optional IE. Such an IE may also have a new type value. Any legacy receiving PFCP entity shall, however, silently discard such an IE and continue processing the message.

The above statement ensures compatibility on a protocol level, but it does not provide compatibility on the higher application level. In addition, 3GPP TS 29.244 [28] section 7.6 describes Error handling including: Unknown IE, Missing IE, and unexpected IE. The following issues can occur when IEs are silently discarded:

1. The peer successfully validates the IEs in the message, applies the known IEs, and achieves the intended outcome. This is because the unknown IE are optional information. For example, most Node Reports IEs from DBNG-UP to DBNG-CP fall in this category.
2. The peer successfully validates the IEs in the message, applies the known IEs, and does not achieve the intended outcome. For example, a DBNG-CP sending LCP keepalive offload IEs to a DBNG-UP not supporting it may not be able to correctly handle LCP keepalive messages itself, leading to PPP session failures.
3. The peer cannot validate the IEs in the message, does not apply the IEs, and generates an error message. This error message will not provide any correlation to the ignored IEs but will likely be a generic error such as 'Rule creation/modification Failure' (cause 73).

It is not predictable whether the ignored IE will result in case 2 or 3 above. To ensure a deterministic behavior, 3GPP TS 29.244 [28] additionally introduces the concept of UP and CP Function Feature Flags. A new function feature can define:

1. A new set of IEs
2. Supporting existing IEs in new grouped IE contexts
3. Supporting new values for extensible IEs.

Based on this information a PFCP entity can take preventive actions. These preventive actions depend on the exact feature, and may include failing the procedure, falling back to another older method, or disabling the new feature altogether.

Therefore, to ensure interoperability between DBNG-CP and DBNG-UP, it is a best practice to group common features into a function feature flag. It is also best practice to specify the DBNG-CP default behavior if certain function feature flags are not signaled by the DBNG-UP and vice versa, the DBNG-UP default behavior if certain function feature flags are not signaled by the DBNG-CP. This is applicable to all TRs related to TR-459.

Annex A: Use Cases

A.1 Multi-access MS-BNG CUPS use case

Since TR-101 [4], the MS-BNG has evolved beyond basic broadband wireline access. The MS-BNG has enabled broadband services to be served over multiple access types which includes, fixed line users, hybrid access, and even public Wi-Fi. Figure 114 depicts a multiple server MS-BNG serving multiple access. Each “box” below represents a single physical network element. The redundancy mechanism within a single physical network element is out of the scope of this use case.

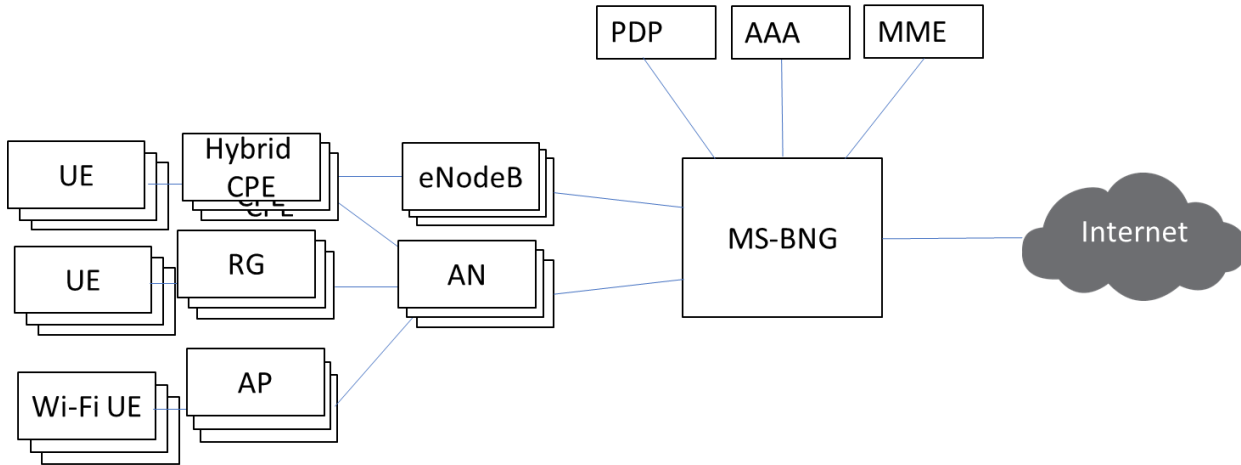


Figure 114: Multi-access MS-BNG

As subscribers scale and bandwidth scale increased, traditionally, this required introducing a new MS-BNG into the network. The new MS-BNG is a separate entity and requires separate commissioning.

BNG control and user plane separation addresses:

- Supporting increase of subscribers
- Supporting new types of subscribers connecting to the multi-access MS-BNG for broadband services enabled by new access technologies.

BNG CUPS simplifies in scaling up both subscribers and bandwidth. The CUPS architecture must allow the service provider to increase the scaling of subscribers on the DBNG-CP independent of the DBNG-UP. And vice versa, as more bandwidth is required, the DBNG-UP can be increased independent of the DBNG-CP. Please note that although DBNG-CP and DBNG-UP scaling can be increased independently, increase of either element requires a proportional increase of the other. The DBNG-CP provides a single point of management for all User Plane instances. The DBNG architecture shown in Figure 115 must continue to offer the same number of functions as today’s MS-BNGs deployed in production networks.

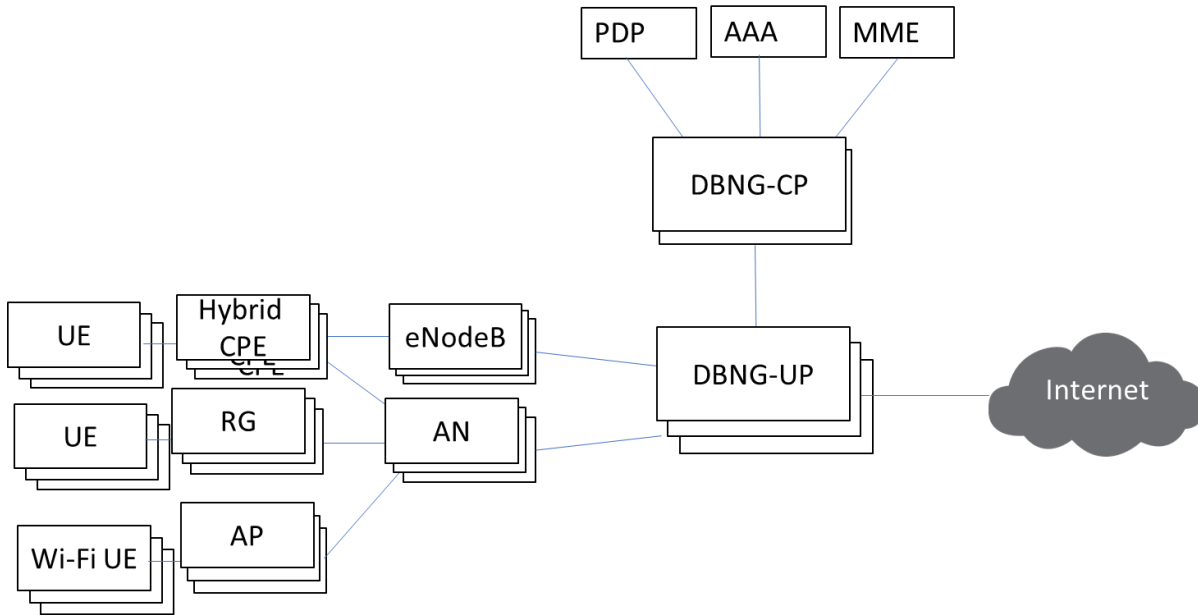


Figure 115: Multi-access DBNG

A.2 Wireline-Access disaggregated MS-BNG use case

From the perspective of wireline access services, it consists of residential subscriber access service and enterprise access service. Residential subscriber access services include HSI, IPTV multicast, VoIP, IMS (TR-69 [3]), and enterprise subscriber access service is usually a leased line service. The subscriber dials up through the access network, and MS-BNG User Plane redirects the access control packets to MS-BNG Control Plane which authenticates the subscriber and creates subscriber forwarding entries upon which MS-BNG User Plane forwards the subscriber data traffic. Usually after the completion of the subscriber access procedure, MS-BNG needs to do network address translation for the subscriber, both native IP and MPLS VPN could be the underlying technologies for MS-BNG network-side data traffic forwarding. The typical wireline access service scenarios include IPTV multicast, CGN, L2TP, and LI with access types of PPPoE, L2/L3 IPoE, dual-stack IPoE, L2-line, L3-line, L2TP, L2VPN (VLL/VPLS), and etc. Address management will require special consideration when the MS-BNG control plane and user plane functions are separated.

A.3 Data-trigger service use case

This type of service is common amongst enterprise customers who are provided static IP address. These static IP addresses are statically configured on CPE. Therefore, DHCP or PPPoE address request is not required. In this use case, the MS-BNG will verify the subscriber based on the data packet (e.g., Ethernet and IP header) and optionally other physical attributes (e.g., NAS port and NAS port ID). For this document, this use case will be referred to as data-trigger service where the DBNG-CP verification is triggered by the first subscriber data packet. In the DBNG case, the first subscriber data packet is the control packet and will be redirected by the DBNG-UP to the DBNG-CP for verification.

A.4 Wholesale Retail model with L2TP use case

As shown in Figure 116, the RG initiates PPPoE negotiation with the DBNG-CP, where the DBNG-CP cannot determine if the user should be terminated locally or tunneled through L2TP until authorization. The PPPoE packets are sent by the DBNG-UP to the DBNG-CP via the general CPR interface (tunnel). Once the DBNG-UP and DBNG-CP complete PPP LCP and PPP authentication, it is determined that this particular subscriber **session** is to be tunneled through L2TP. The DBNG-CP must initiate a L2TP tunnel/session with

the LNS, via the A10 interface. The DBNG-CP initiates the SCCRQ via the (subscriber-specific) packet redirect tunnel to the DBNG-UP, which would forward the packet out the A10 interface. Similarly, remaining SCCRP, SCCCN, ICRQ, ICRP, ICCN control packets are exchanged between the DBNG-CP through the DBNG-UP towards the LNS (via the A10). It should be noted that at any point in time, the session forwarding state can be updated for example due to bandwidth changes as specified in RFC 5515 [42]. Upon tunnel set-up completion, the DBNG-CP must be able to signal an update to the DBNG-UP to forward all packets from the RG directly to the A10 (via the DBNG-UP) without DBNG-CP involvement. The remaining, PPP NCP, LCP echo requests/responses, ND, RA, DHCPv6, etc. are all forwarded without DBNG-CP involvement. It is important that a single subscriber can have more than one PPPoE service, only some of which are offered by retailers. Therefore, not all PPPoE sessions are to be redirected to the LNS.

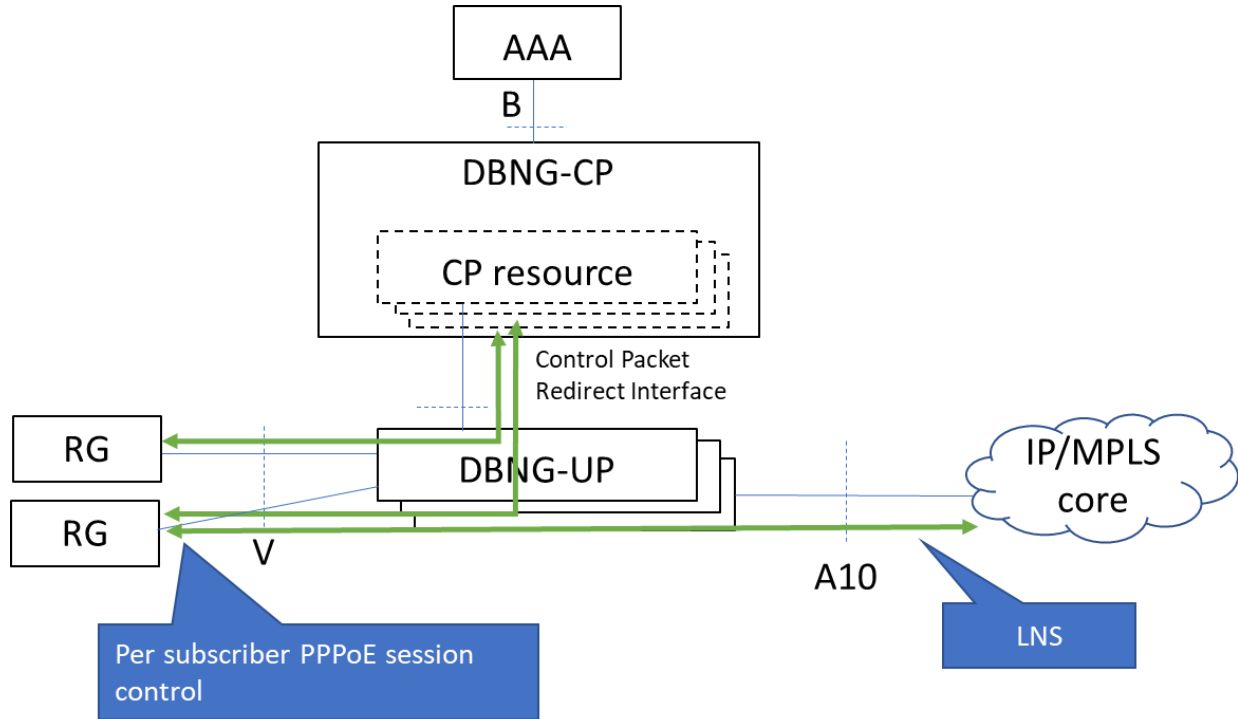


Figure 116: L2TP deployment model

A.5 IPoE DHCP and DHCPv6 Relay Service Model

DHCP relay and relay-proxy and DHCPv6 relay are used for deployments in which an external DHCP server is used to satisfy the address and configuration attributes requested in the DHCP DISCOVER or DHCPv6 SOLICIT. An external DHCP/DHCPv6 server may be used within the provider's network in conjunction with AAA or in support of L3 wholesale in which the DHCP/DHCPv6 server is part of the retailer's network. For the former case, the external DHCP/DHCPv6 server is connected to the DBNG-CP, whereas, for the latter, it is connected to the DBNG-UP via the A10 interface.

For all cases, upstream DHCP control packets from the RG are forwarded by the DBNG to an external DHCP server, which satisfies the addressing and configuration attributes requested in the DHCP DISCOVER or DHCPv6 SOLICIT. The DBNG will act as a stateful DHCP relay, meaning all DHCP control packets between the subscriber session and external DHCP server are snooped, the subscriber session authenticated with RADIUS, and subscriber QoS and forwarding policy optionally applied.

For L3 wholesale with multiple retailers, dedicated network instances are typically used for each retailer such that subscriber sessions are assigned to an instance by AAA during authorization and configured within the network instance associated with the corresponding retailer during Session Establishment. This means the A10 interface and subscriber sessions assigned to the retailer are co-located in the same network instance. One or more eligible servers are provisioned within the network instance, such that the upstream DHCP control packet is relayed to each of the eligible servers. The RG (client) will select the server from the DHCP OFFER or DHCPv6 ADVERTISE packets received by responding with a REQUEST containing the selected server's server-identifier.

When the external DHCP/DHCPv6 server is connected to the DBNG-UP via the A10 interface, DBNG support of the DHCP relay service model requires separate control packet exchanges, one between the DBNG DHCP relay and the RG and one between the DBNG DHCP relay and external DHCP server. The control packet exchange between the RG and DBNG DHCP relay is realized by the existing default CPRi and subscriber session CPRi that are used for existing access models described in this document. The control packet exchange between DBNG DHCP relay and external DHCP server, however, requires a new CPRi channel between DBNG-CP and DBNG-UP for the DBNG-CP to send DHCP control packets to the external server via the DBNG-UP and for the DBNG-UP to redirect external server originated DHCP control packets to the DBNG-CP. This new CPRi server channel, depicted by the orange arrow in Figure 117 below, is per network instance and used to exchange packets between configured external DHCP and DHCPv6 servers and the DBNG-CP via the DBNG-UP for that network instance. As the set of eligible DHCP and DHCPv6 servers are typically configured beforehand, this CPRi server channel is configured after a successful node association between the DBNG-CP and DBNG-UP.

The following diagram illustrates the CPRi channels for DHCP and DHCPv6 relay:

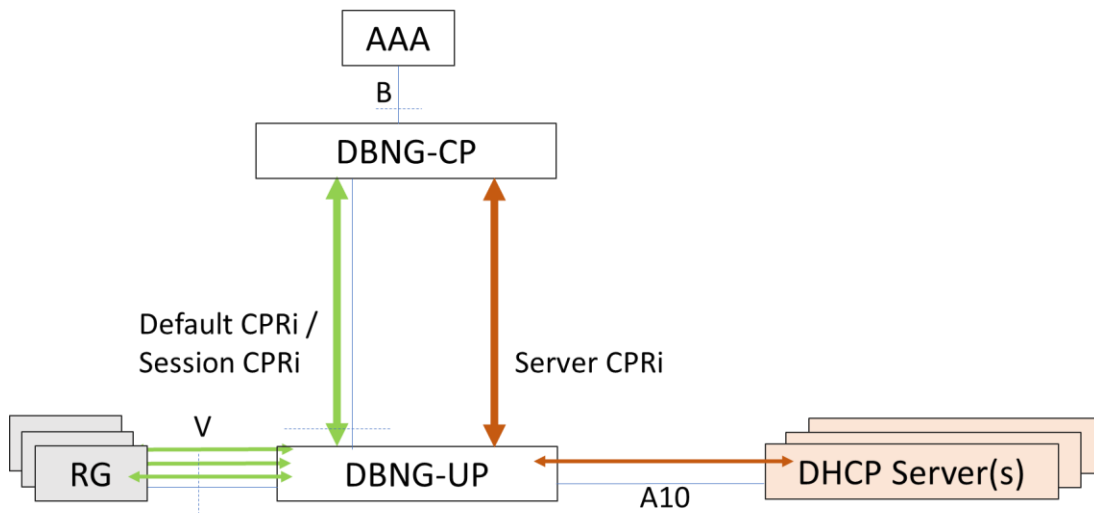


Figure 117: DBNG-UP Connected External DHCP Server Deployment Model

A.5.1 DHCPv4 Relay

For DHCPv4 relay, the DBNG-CP receives the upstream broadcast DHCPv4 DISCOVER or REQUEST packets from the default or per subscriber session CPR interface and converts the packet to unicast, including setting the relay agent IP address (GIADDR) to the DBNG-UP IPv4 address of the V (access) interface. A unicast packet is created and sent for each eligible external DHCP server.

For a DBNG-UP connected external DHCP server, the DBNG-CP forwards the unicast DHCP DISCOVER/REQUEST packet to the DBNG-UP via the server CPRi. The DBNG-UP will forward the unicast packet to the external DHCP server via the DHCP server-facing (A10) interface. The DBNG-UP will redirect external DHCP server originated packets (e.g., OFFER, ACK) to the DBNG-CP via the server CPRi.

The DBNG-CP will convert the downstream DHCP control packet (OFFER, ACK) to a broadcast or unicast packet, per the broadcast bit in the DISCOVER packet, and sends it to DBNG-UP via the default/common or subscriber session CPRi. The DBNG-UP will then forward the downstream packet to the RG via the V (access) interface.

Note that the unicast DHCP RENEW exchange between the RG and external DHCP server will be snooped by the DBNG-CP before sending it to the server (for a DBNG-UP connected external DHCP server, to the DBNG-UP via the server CPRi (REQUEST)) or the subscriber session CPRi (ACK).

A.5.2 DHCPv6 Relay

For DHCPv6 relay, the DBNG-CP receives the upstream multicast DHCPv6 SOLICIT or REQUEST packets from the default or per subscriber session CPR interface. The packets are received with a RELAY-FORW header inserted by the LDRA. The DBNG-CP will encapsulate with a RELAY-FORW header, setting the link address to the DBNG-UP local IPv6 address of the subscriber logical port (V interface) and convert the packet to unicast. A unicast packet is created and sent for each eligible external DHCPv6 server.

For a DBNG-UP connected external DHCPv6 server, the DBNG-CP forwards the unicast DHCPv6 SOLICIT/REQUEST packet to the DBNG-UP via the server CPRi. The DBNG-UP will forward the unicast packet to the external DHCPv6 server via the DHCPv6 server-facing (A10) interface. The DBNG-UP will redirect external DHCPv6 server originated packets (e.g., ADVERTISE, REPLY) to the DBNG-CP via the server CPRi.

The DBNG-CP will decapsulate the RELAY-REPLY header from this downstream DHCPv6 control packet before sending it to the DBNG-UP via the default/common or subscriber session CPRi. The DBNG-UP will then forward the downstream packet to the RG via the V (access) interface.

A.6 Subscriber Session ACL and ACL Chaining

Each of the subscriber session on DBNG-UP can have an ACL applied for filtering, routing, and other specific operations. Such ACL definitions exist on the DBNG-UP. DBNG-CP can apply these ACL(s) through SCi by specifying ACL name, family, and direction.

ACL Chaining allows attaching of multiple ACLs to the same subscriber session in ordered manner. This helps in splitting of large ACL definitions on DBNG-UP into common ACLs and session/subscriber specific or service specific ACLs. Below diagram shows how two ACLs (acl_a and acl_b) are chained for a subscriber session 1, whereas for subscriber session 2, the chaining is different (acl_a, acl_c). If no rules match data packet in the first ACL in the chain, the rules in next ACL are checked and so on.

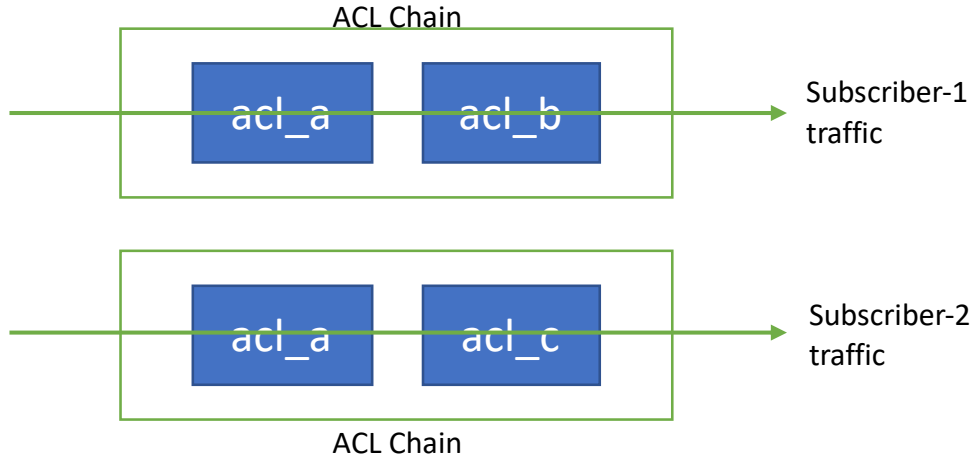


Figure 118: ACL chaining

Appendix I. Alternative Call Flows for PPPoE delayed session creation

This appendix shows possible alternative methods to implement delayed session creation for PPPoE based sessions (PPPoEv4, PPPoE dual-stack, LAC, LNS PPPoEv4, LNS Dual Stack). It is understood that any alternative where the Session Creation Request is sent after the PADO and before the IPCP conf-ack is a possible implementation of delayed session creation. Therefore, this appendix has a plain exemplificative scope.

I.1 PPPoE delayed session creation

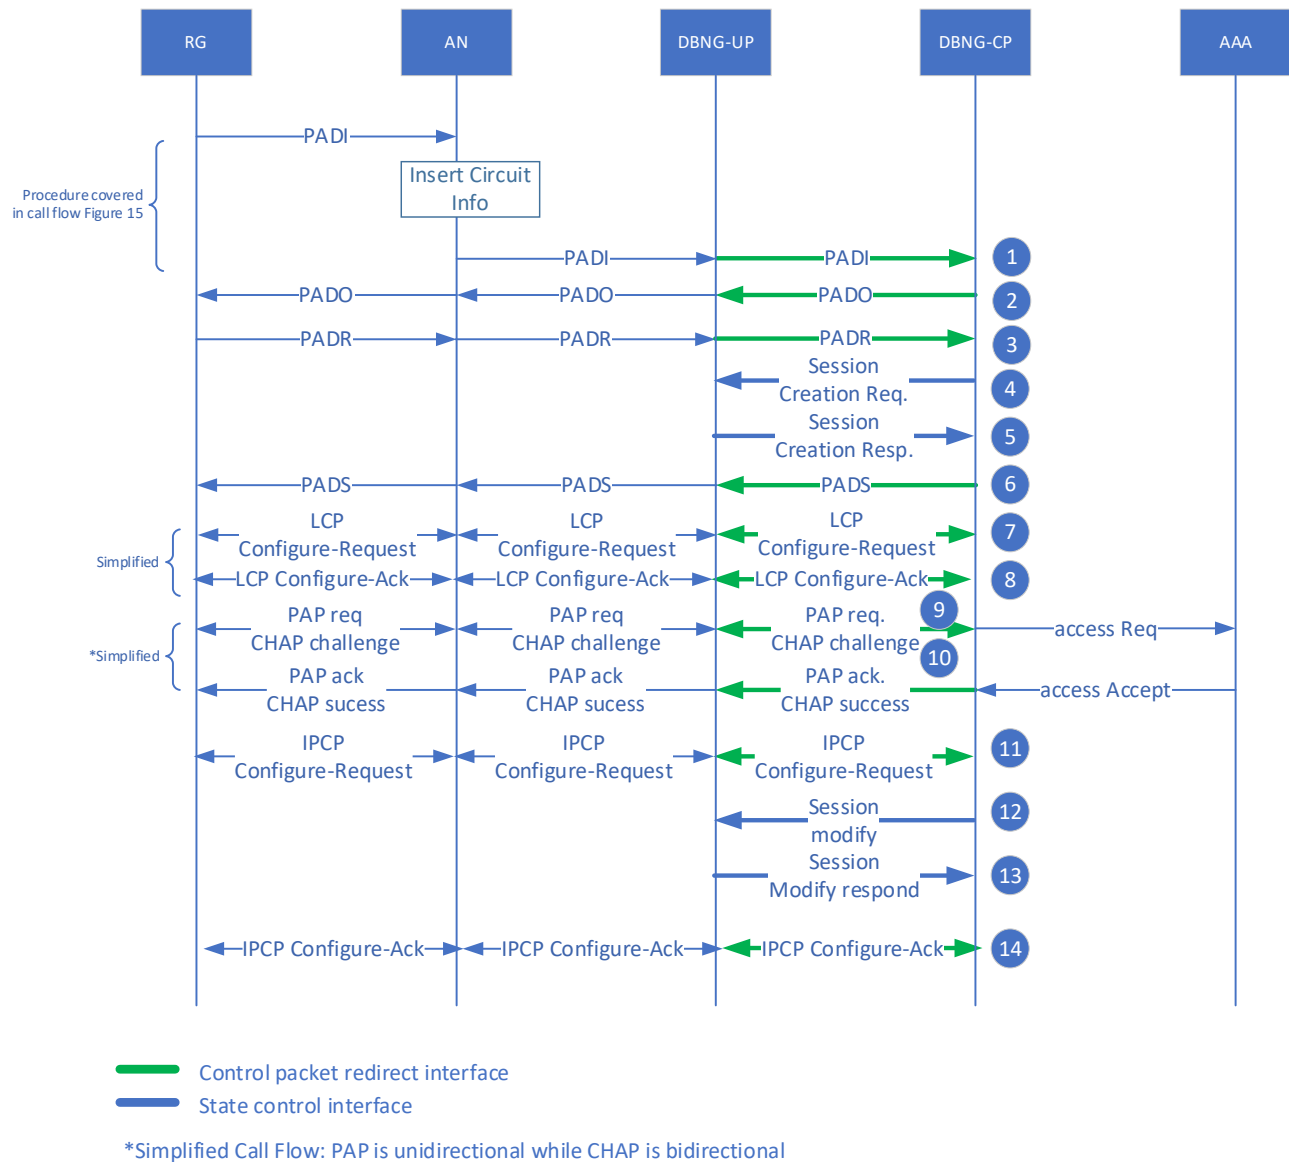


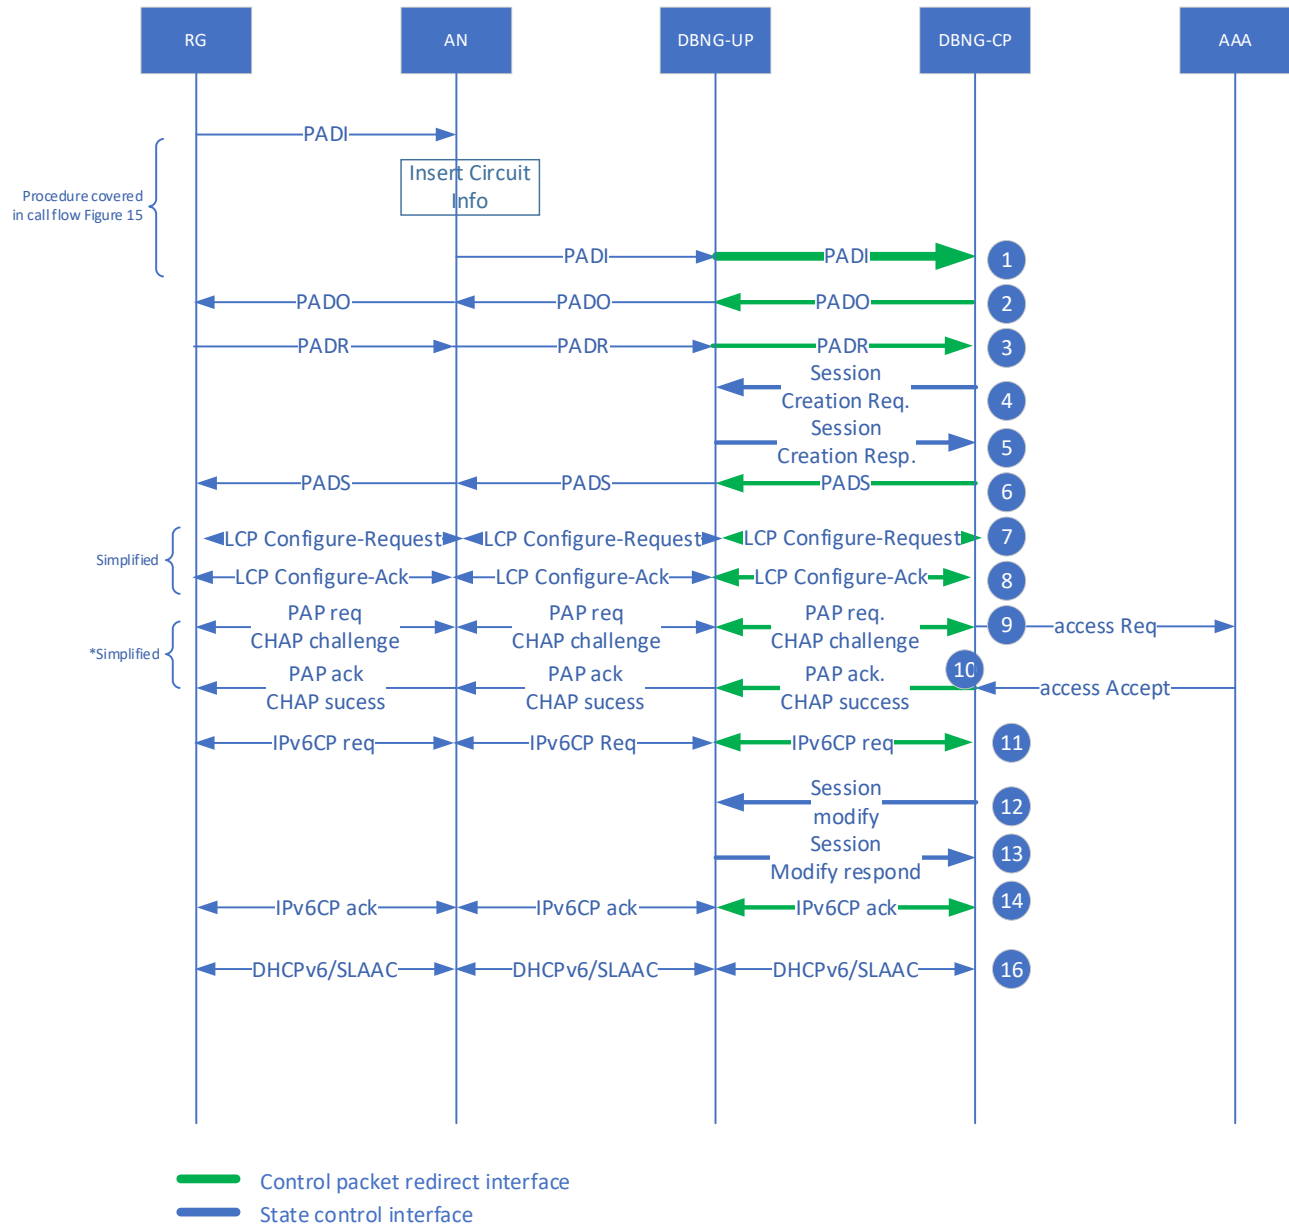
Figure 119: PPPoE call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions.

1. DBNG-CP will trigger and check for the desired service asked by the RG.
2. The DBNG-CP sends the PADO back to the RG through the DBNG-UP utilizing the downstream default redirect tunnel interface.
3. The PADR is sent from the RG through the DBNG-UP utilizing the upstream default redirect tunnel interface.
4. Upon receiving PADR, the DBNG-CP can at this point send a session creation request to create new packet forwarding states for the data packet. This updates the data plane state.

5. The DBNG-UP sends a response back to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready for forwarding.
6. The DBNG-CP sends the PADS back to the RG through the DBNG-UP utilizing the CPR interface.
7. The LCP Configure-Request is sent from the RG through the DBNG-UP utilizing the CPR interface.
8. The DBNG-CP sends the LCP Configure-Ack back to the RG through the DBNG-UP utilizing the CPR interface. The LCP Configure-Ack indicates either a PAP or CHAP authentication challenge.
9. Options:
 - Option 1: If the client chooses PAP, the RG sends a PAP request to the DBNG-CP through the DBNG-UP utilizing the CPR Interface. The credentials are sent in the Access-Request to the AAA server.
 - Option 2: If CHAP is required, the DBNG-CP initiates a challenge to the RG through the DBNG-UP utilizing the CPR Interface. The RG responds back to the challenge to the DBNG-CP. The challenge response is sent to the AAA server.
10. The AAA successfully authenticates the RG and replies to the RG with PAP/CHAP success.
11. Both DBNG-CP and RG send IPCP Configure-Request for parameter negotiation, utilizing a dedicated session control packet redirect tunnel. The RG is assigned an IPv4 address. Address could be assigned to the RG either through the AAA reply or through a local address server. As noted, if a session has not yet been established, the session must be established at this step.
12. Once the RG is assigned IP address, the DBNG-CP sends a session modification request to create new packet forwarding states for the data packet. The data plane is updated.
13. The DBNG-UP sends a response back to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready to forward the subscriber's IP data packets.
14. The IPCP Configure-Ack is sent from the DBNG-CP to the RG through the DBNG-UP utilizing a dedicated session control packet redirect tunnel.

I.2 PPPoEv6 delayed session creation



*Simplified Call Flow: PAP is unidirectional while CHAP is bidirectional

Figure 120: PPPoEv6 call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common/default CPR rule for upstream and downstream directions.

1. DBNG-CP will trigger and check for the desired service asked by the RG.
2. The DBNG-CP sends the PADO back to the RG through the DBNG-UP utilizing the downstream default redirect tunnel interface.
3. The PADR is sent from the RG through the DBNG-UP utilizing the upstream default redirect tunnel interface.

4. Upon receiving PADR, the DBNG-CP can at this point send a session creation request to create new packet forwarding states for the control packet. This updates the data plane state.
5. The DBNG-UP sends a response back to the DBNG-CP, informing that the states are installed, and the DBNG-UP is ready for forwarding.
6. The DBNG-CP sends the PADS back to the RG through the DBNG-UP utilizing the CPR interface.
7. The LCP Configure-Request is sent from the RG through the DBNG-UP utilizing the CPR interface.
8. The DBNG-CP sends the LCP Configure-Ack back to the RG through the DBNG-UP utilizing the CPR interface. The LCP Configure-Ack indicates either a PAP or CHAP authentication challenge.
9. Options:
 - Option 1: If the client chooses PAP, the RG sends a PAP request to the DBNG-CP through the DBNG-UP utilizing the upstream default redirect tunnel interface. The credentials are sent in the Access-Request to the AAA server.
 - Option 2: If CHAP is required, the DBNG-CP initiates a challenge to the RG through the DBNG-UP utilizing upstream and downstream default redirection tunnel interfaces. The RG responds back to the challenge to the DBNG-CP. The challenge response is sent to the AAA server.
10. The AAA successfully authenticates the RG and replies to the RG with PAP/CHAP success.
11. The IPv6CP Configure-Request is sent from the RG through the DBNG-UP utilizing the upstream default redirect tunnel interface.
12. At this point, the DBNG-CP had obtained the IPv6 addresses and prefixes for the RG either from the local address server or from AAA returned VSAs. The DBNG-CP sends a session modification request to specify packet forwarding states for data packet. The data plane state is updated.
 - Traffic management rule for data packets: match data packet and perform forwarding action.
13. The DBNG-UP sends a response back to the DBNG-CP, informing that the states are installed and the DBNG-UP is ready to forward the subscriber's IP data packets.
14. The DBNG-CP sends the IPv6CP Configure-Ack to the RG through the DBNG-UP utilizing the CPR interface.
15. In the case of SLAAC prefix assignment, the DBNG-CP sends an RA to the RG informing the Link Local Address which is described in detail in section 4.5.10. In the case of DHCPv6 assignment, please refer to section 4.5.7 delayed session creation mode.

I.3 PPPoE Dual Stack delayed session creation

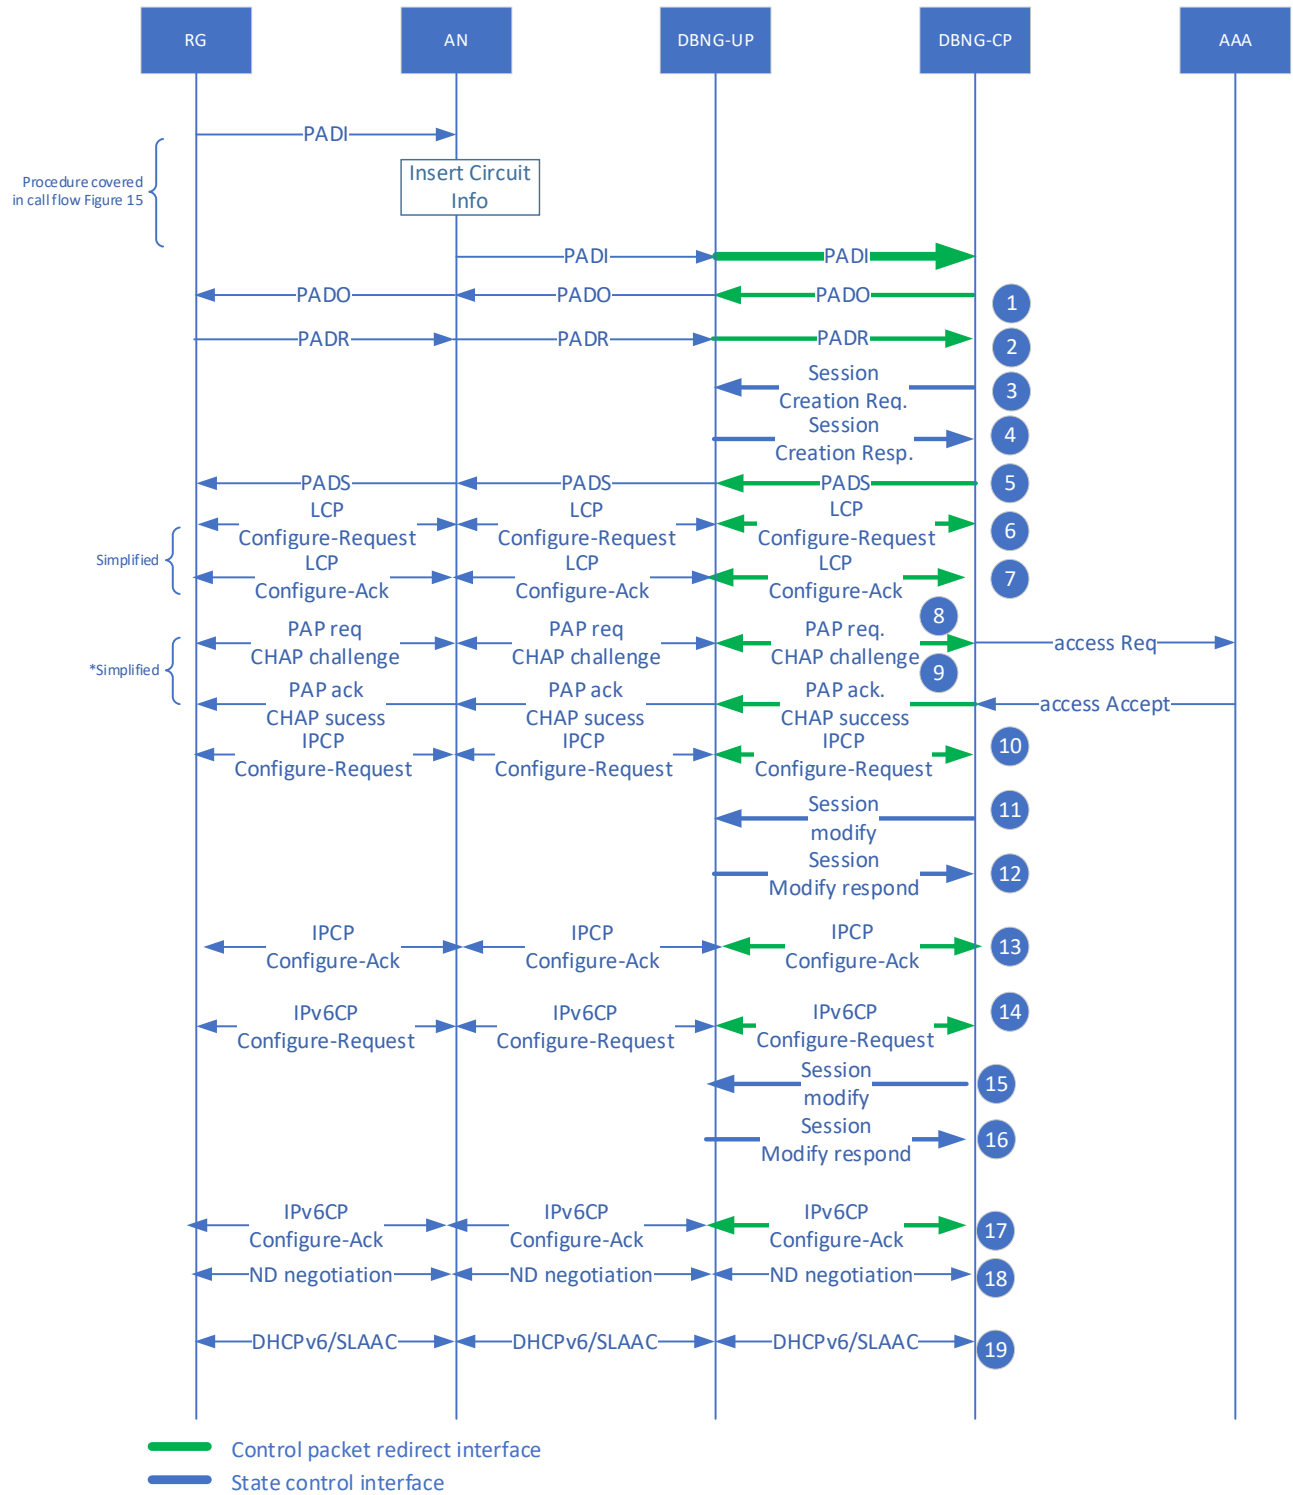


Figure 121: Delayed PPPoE Dual Stack call flow

Prior to step 1, call flow in section 4.5.2 covers the generic common/default CPR rule for upstream and downstream directions.

1-13. Follows the same procedure as section I.1delayed session creation model step 1-13

14-19. Follows the same procedure as section I.2delayed session creation model step 11-16

Note: The dual-stack access processes could be concurrent, one before the other, or one after the other.

I.4 LAC delayed session creation

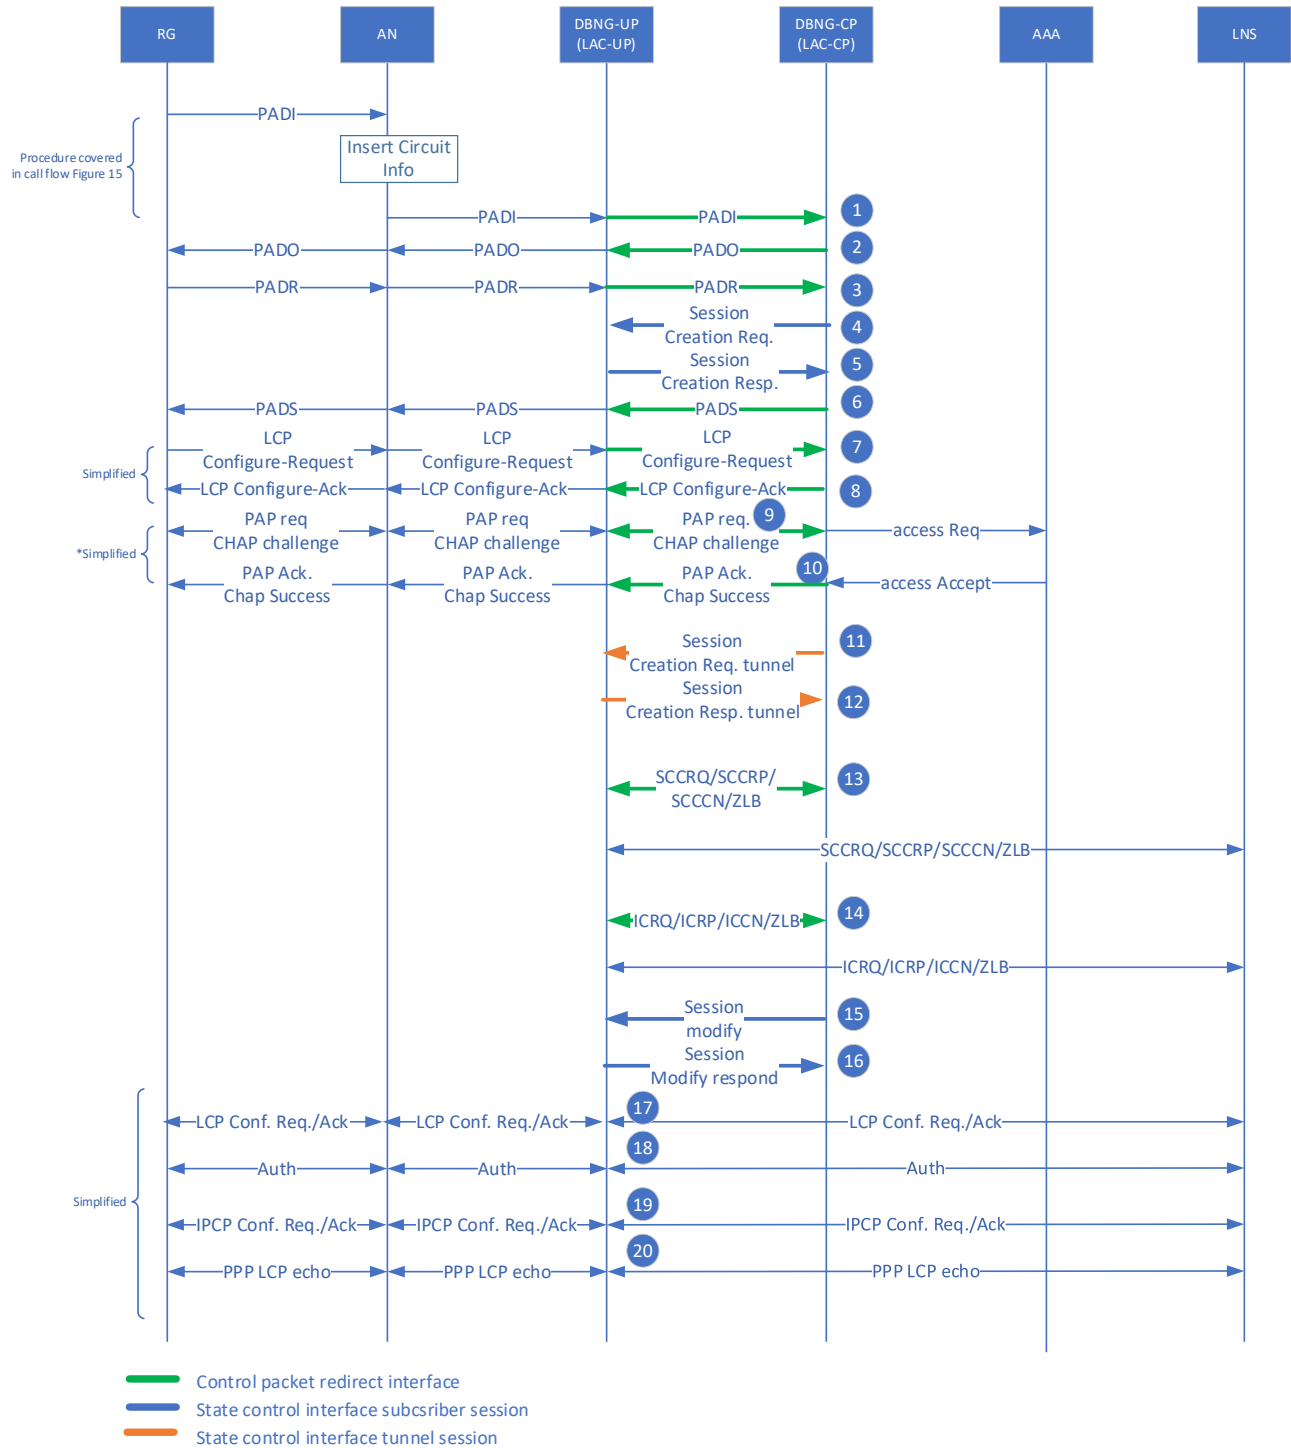


Figure 122: LAC call flow – Delayed Session Creation

This call flow describes about delayed session creation for the use case described in section 4.4.12 of TR459 issue 1.

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions.

Step 1&2 of section 4.4.12 are delayed until the completion of step 4 of section 4.4.12 of TR459.
In step 2, BNG CP sends PADO packet to RG via BNG UP using default downlink redirection rule.

I.5 LNS – PPPoEv4 delayed session creation

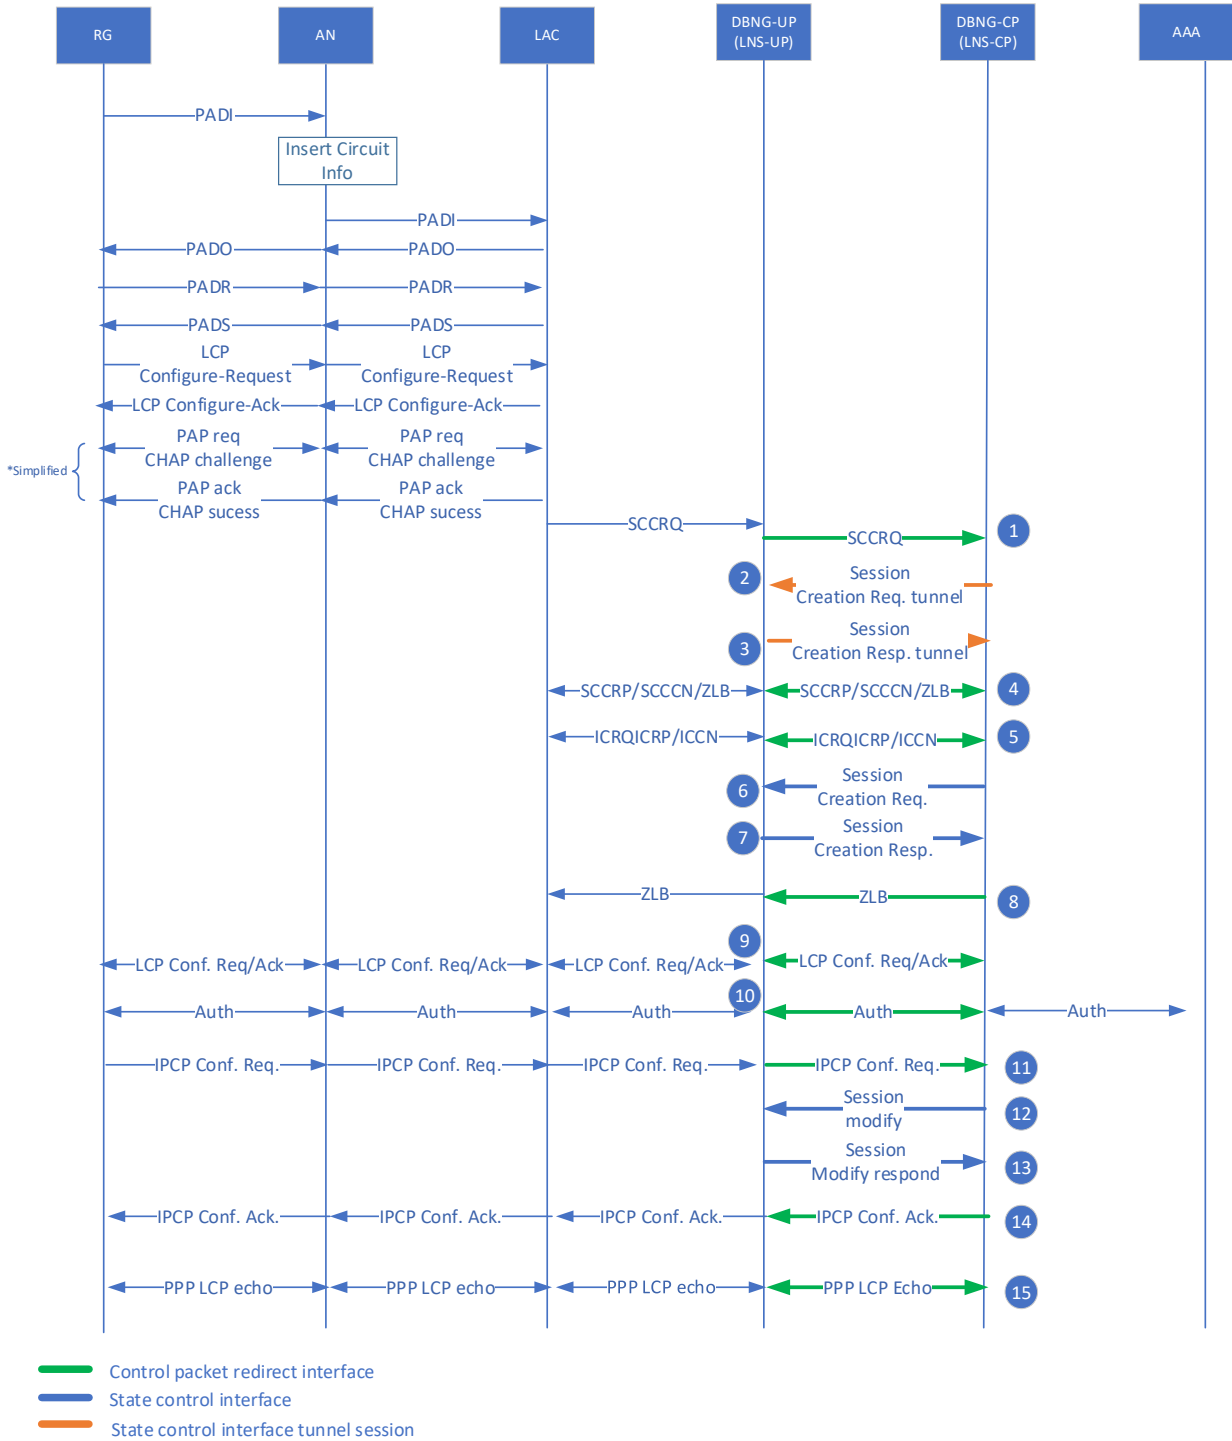


Figure 123: LNS IPv4 call flow Delayed Session Creation

Prior to step 1, call flow in section 4.5.2 covers the generic common CPR rule for upstream and downstream directions.

This section describes an option where the I2tp session creation is delayed till the I2tp ICRP message is received by BNG CP LNS i.e., step 6&7 of section 4.4.13 of TR459 are delayed until step 8 is complete without ZLB, after the I2tp session is created, ZLB packet will be sent by BNG CP.

I.6 LNS - Dual Stack delayed session creation

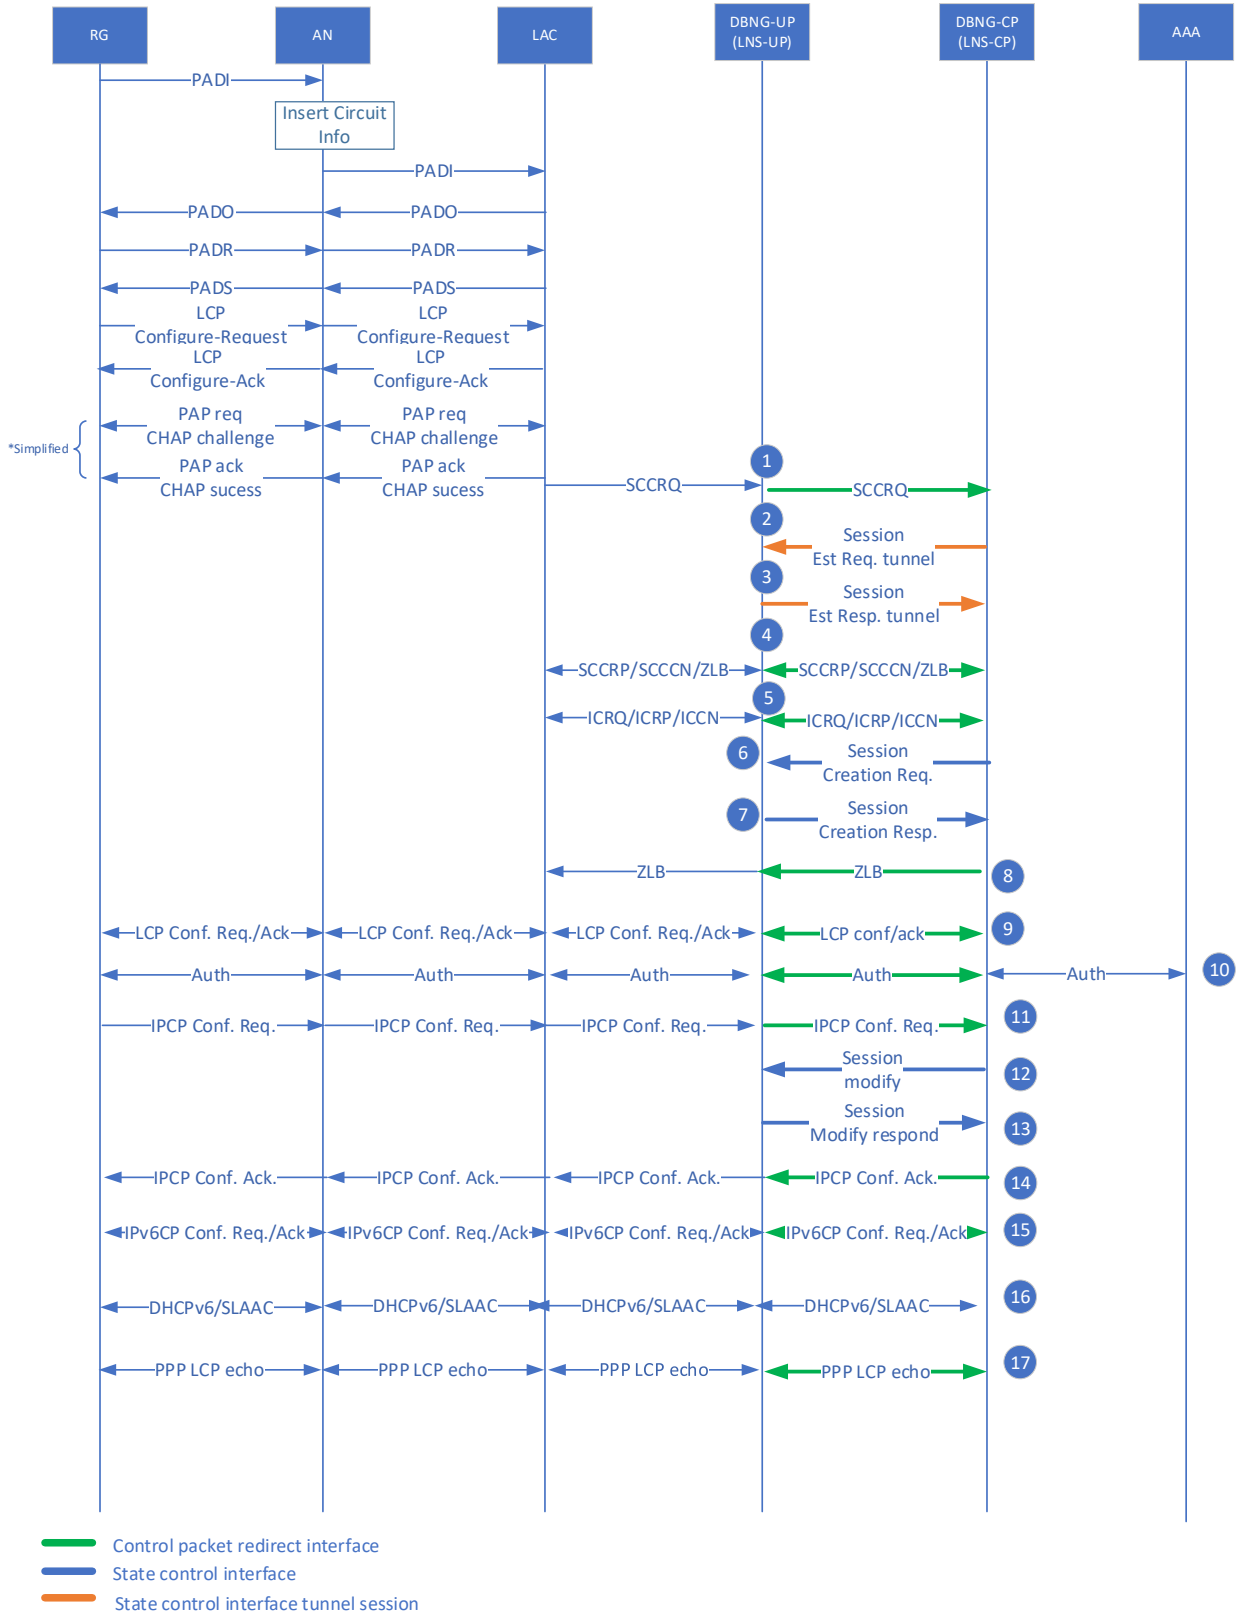


Figure 124: LNS Dual Stack call flow Delayed Session Creation

This case is same as the call flow in section I.5 of this document, but for dual stack case.

End of Broadband Forum Technical Report TR-459