

TR-459.2
**Multi-Service Disaggregated BNG with CUPS:
Integrated Carrier Grade NAT function. Reference
Architecture, Deployment Models, Interface, and
Protocol Specifications**

Issue: 2
Issue Date: January 2025

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	Issue Date	Issue Editor	Changes
1	29 October 2021	Kenneth Wan, Nokia	Original
2	29 January 2025	Nagaraj S Turaiyur, Juniper Networks	Subscriber group (SGRP) extended for CGNAT private/public prefixes

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

Editor: Nagaraj S Turaiyur, Juniper Networks

Work Area Director(s): Jonathan Newton, Vodafone

Project Stream Leader(s): Hans-Joerg Kolbe, Radisys
Jonathan Newton, Vodafone

Table of Contents

Executive Summary	8
1 Purpose and Scope	9
1.1 Purpose	9
1.2 Scope	9
2 References and Terminology	10
2.1 Conventions	10
2.2 References	10
2.3 Definitions	11
2.4 Abbreviations	11
3 DBNG with integrated Carrier Grade NAT (CG-NAT) Overview	14
3.1 NAT Concepts	14
3.1.1 CG-NAT Introduction	14
3.1.2 CG-NAT Logging	15
3.2 DBNG Functional Architecture	16
3.2.1 Control Plane Function Overall:	16
3.2.2 User Plane Function Overall:	17
3.3 DBNG CGNAT models	17
3.3.1 DBNG-CP centric model	17
3.3.2 DBNG-UP centric model	18
3.3.3 Common Dynamic NAT block allocation for both models	19
4 Call Flows	20
4.1 DBNG-CP centric model call flows	20
4.1.1 Initial NAT block allocation for DHCP subscribers	20
4.1.2 Initial NAT block allocation for PPPoE subscribers	22
4.1.3 Session termination	23
4.2 DBNG-UP centric model	24
4.2.1 Initial NAT block allocation for DHCPv4 Subscriber	24
4.2.2 Initial NAT block allocation for PPPoE Subscriber	25
4.2.3 Session termination	26
4.3 Common Dynamic NAT block allocation for all models	27
4.4 Establishment of resilient session (CP-centric model and UP-centric model)	27
5 Technical Requirements	29
5.1 General NAT requirement on DBNG	29
5.2 DBNG-CP requirements	29
5.2.1 DBNG-CP centric model	29
5.2.2 DBNG-UP centric model	30
5.2.3 Common dynamic NAT block requirements for both models	30
5.2.4 Subscriber group (SGRP) requirements for both models	30
5.3 DBNG-UP requirements	30
5.3.1 DBNG-CP centric model	30
5.3.2 DBNG-UP centric model	31
5.3.3 Common dynamic NAT block requirements for both models	31
5.4 Management Interface requirements	31
5.5 NAT Logging Timestamp requirements	31
6 PFCP CUPS protocol	32

6.1	DBNG-CP centric model use cases and IEs exchanges	32
6.1.1	Use case: Initial NAT block allocation.....	32
6.1.2	Use case: Dynamic NAT block allocation	32
6.2	DBNG-UP centric model use cases and IE exchanges	33
6.2.1	Use case: Initial and dynamic NAT block allocation	33
6.3	Common IEs NAT block allocation/deallocation reporting	33
6.3.1	Create PFCP IE Session Report Rule (SRR) Extensions	33
6.3.2	PFCP session reporting rule	34
6.4	PFCP IEs requirement for NAT SGRP.....	34
6.5	PFCP IE summary.....	34
6.6	PFCP Grouped IE extensions	38
6.6.1	BBF UP Function Features IE.....	38
6.6.2	Create SRR IE	40
6.6.3	BBF Report Trigger	41
6.6.4	Session Report IE within PFCP Session Report Request.....	41
6.6.5	BBF Apply Action IE.....	41
6.6.6	BBF NAT External Port Range	42
6.6.7	BBF NAT Port Forward	42
6.6.8	BBF Dynamic NAT Block Port Range.....	43
6.6.9	BBF Event Time Stamp	43
Appendix I.	44
DBNG dynamic NAT block use case:	44
Use cases for prefix routing announcement and resiliency		45

Table of Figures

Figure 1: CG-NAT NAT 44 scenario15

Figure 2: CG-NAT NAT 444 scenario15

Figure 3: DBNG with integrated CG-NAT function architecture16

Figure 4: DBNG-CP centric model with initial NAT block allocation for DHCP subscribers.....20

Figure 5: DBNG-CP centric model with initial NAT block allocation for PPPoE subscribers22

Figure 6: DBNG-CP centric model subscriber session termination23

Figure 7: DBNG-UP centric model with initial NAT block allocation for DHCP subscribers.....24

Figure 8: DBNG-UP centric model with initial NAT block allocation for PPPoE subscribers25

Figure 9: DBNG-UP centric model subscriber session termination26

Figure 10: Dynamic NAT block allocation for all models27

Figure 11: Establishment of resilient session (CP-centric model and UP-centric model)28

Figure 12: BBF UP Function Features38

Figure 13: BBF Apply Action41

Figure 14: BBF NAT External Port Range.....42

Figure 15: BBF NAT port forward42

Figure 16: BBF Dynamic NAT Block Port Range43

Figure 17: BBF Event Time Stamp43

Table of Tables

Table 1: BBF UP Function Features Flag and applicable BBF PFCP IEs34

Table 2: PFCP Session Establishment Request and BBF Grouped IEs structure34

Table 3: PFCP Session Modify Request and BBF Grouped IEs structure36

Table 4: PFCP Session Report Request and BBF Grouped IEs structure36

Table 5: BBF Extended Information Element and Applicability Table for association and default redirection
tunnels for NAT use case.....37

Table 6: BBF Extended Information Element and Applicability Table for DBNG-CP centric NAT model37

Table 7: BBF Extended Information Element and Applicability Table for DBNG-UP centric NAT model37

Table 8: BBF UP Function Features.....38

Table 9: BBF extended Create Traffic Endpoint IE(s) within PFCP Session Establishment Request.....39

Table 10: BBF extended Forwarding Parameters IE in FAR.....40

Table 11: BBF extended Update Forwarding Parameters IE(s) in FAR.....40

Table 12: Create SRR IE within PFCP Session Establishment Request.....40

Table 13: BBF extended Session Report IE.....41

Table 14: Log events44

Executive Summary

This Technical Report extends the TR-459: Multi-Service Disaggregated BNG with CUPS - Reference Architecture, Deployment Models, Interface and Protocol Specification to include CGNAT support.

1 Purpose and Scope

1.1 Purpose

TR-459: Multi-Service Disaggregated BNG with CUPS: Reference Architecture, Deployment Models, Interface, and Protocol Specifications [1] serves as a baseline document for the DBNG architecture, requirements, and PFCP protocol extension specification. In this document, the DBNG solution is further extended to incorporate the fixed access Carrier Grade Network Address Translation (CG-NAT). This document will extend DBNG architecture, technical requirements, and PFCP protocol to cover CG-NAT.

1.2 Scope

The following is in scope of this document:

- NAT standards following the best practice defined in RFC 7857 [10]
- Use cases that involves the following:
 - o CG-NAT
 - o Dynamic NAT block allocation
- CG-NAT Functional Definitions
- CG-NAT Functional Requirements
- DBNG Architecture with integrated CG-NAT function
- CG-NAT Call flows
- PFCP IE extensions to support CG-NAT

In this issue, the subscriber group (SGRP) concept in TR-459[1] is extended for CGNAT subscribers to support resiliency by defining requirements and a call flow. SCi extensions have been made to SGRP to carry and advertise CGNAT private/ public prefixes.

Out of scope for this document:

- NAT for IPv6 whether for NAT 66, other variations, or IPv6 transition strategies.

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 [3].

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-459	Multi-Service Disaggregated BNG with CUPS. Reference Architecture, Deployment Models, Interface, and Protocol Specification	BBF	2025
[2] 3GPP 29.244	TS Interface between the Control Plane and the User Plane nodes v16.5.0	3GPP	Dec 2019
[3] RFC 2119	Key words for use in RFCs to Indicate Requirements Levels	IETF	1997
[4] RFC 3022	Traditional IP Network Address Translator	IETF	2001
[5] RFC 5382	NAT Behavioral Requirements for TCP	IETF	2008
[6] RFC 4787	Network Address Translation (NAT) Behavioral Requirements for Unicast UDP	IETF	2007

[7]	RFC 5508	NAT Behavioral Requirements for ICMP	IETF	2009
[8]	RFC 2663	IP Network Address Translator (NAT) Terminology and Considerations	IETF	1999
[9]	RFC 7422	Deterministic Address Mapping to Reduce Logging in Carrier-Grade NAT Deployments	IETF	2014
[10]	RFC 7857	Updates to Network Address Translation (NAT) Behavioral Requirements	IETF	2016
[11]	RFC 5905	Network Time Protocol Version 4: Protocol and Algorithms Specification	IETF	2010
[12]	RFC 5424	The Syslog Protocol	IETF	2009

2.3 Definitions

The following terminology is used throughout this Technical Report.

NAT block	A range of ports from a given public IP address
NAT pool	A set of outside IP address(es) that are defined and used for source address and port translation.
Private IP address	An IP address from the Private/Local network as defined in RFC 2663 [8]
Public IP address	An IP address from the Public/Global/External network as defined in RFC 2663 [8]
Outside IP address	An IP address to which the inside IP address has been mapped by NAT..
Inside IP address	An IP address by which a subscriber is identified on the private network.
Outside IP subnet	The IP subnet of outside IP addresses.
Inside IP subnet	The IP subnet of inside IP addresses.
Port Forwarding	When a private IP address and port are statically bound to a public IP address and port

2.4 Abbreviations

This Technical Report uses the following abbreviations:

3GPP	3rd Generation Partnership Project
AAA	Authentication, Authorization & Accounting
ALG	Application Level Gateway
BBF	The Broadband Forum
BNG	Broadband Network Gateway
CG-NAT	Carrier Grade Network Address Translator
CP	Control Plane
CPE	Customer Premises Equipment

CPR	Control Packet Redirect
CUPS	Control and User Plane Separation
DBNG	Disaggregated BNG
DHCP	Dynamic Host Configuration Protocol
EMS	Element Management System
ETH	Ethernet
FTP	File Transfer Protocol
FAR	Forward Action Rule
GW	Gateway
IANA	Internet Assigned Numbers Authority
ICMP	Internet Control Message Protocol
IE	Information Element
IETF	Internet Engineering Task Force
IP	Internet Protocol
IPCP	IP Control Protocol
IPoE	IP over Ethernet
IPoEv4	IPoE version 4
ISP	Internet Service Provider
LCP	Link Control Protocol
MAC	Media Access Control
Mi	Management Interface
MME	Mobile Management Entity
MPLS	Multi-Protocol Label Switching
MS-BNG	Multi-Service BNG
MTU	Maximum Transfer Unit
NAT	Network Address Translation
NAPT	Network Address Port Translation
NCP	Network Control Protocol
NETCONF	Network Configuration Protocol
PFCP	Packet Forwarding Control Protocol
PDI	Packet Detection Information
PDP	Policy Decision Point
PDR	Packet Detection Rule
PEP	Policy Enforcement Point
PFCP	Packet Forwarding Control Protocol
PGW	PDN GW
PPP	Point to Point Protocol
PPPoE	PPPoE over Ethernet
PPPoEv4	PPPoE for IP version 4
PTA	PPP Termination and Aggregation
QoS	Quality of Service
RADIUS	Remote Authentication Dial In User Service
RG	Residential Gateway

RIR	regional Internet registry (RIR)
SRR	Session Reporting Rule
SCCCN	Start Control Connection Connected
SCCRQ	Start Control Connection Request
SCi	State Control interface
SGW	Serving GW
SNMP	Simple Network Management Protocol
SGRP	Subscriber Group
TEID	Traffic Endpoint Identifier
TR	BBF Technical Report
TWAG	Trusted WLAN Access Gateway
UP	User Plane
VLAN	Virtual Local Area Network
VLL	Virtual Leased Line
Wi-Fi	Wireless Fidelity
WLAN	Wireless Local Area Network
WLAN-GW	WLAN Gateway
WT	BBF Working Text

3 DBNG with integrated Carrier Grade NAT (CG-NAT) Overview

3.1 NAT Concepts

Network Address Translator (NAT) RFC2663 [8] RFC3022 [4] is a mechanism that allows the translation of a source IP address to another source IP address. NAT supports two modes: NAT and Network Address Port Translation (NAPT). While NAT supports only IP address translation, NAPT supports both IP address translation and port mapping. NAPT enables multiple hosts to share the same public IP address. IETF RFC 5382 [5], RFC 4787 [6], and RFC 5508 [7] define NAT requirements for TCP, UDP, and ICMP respectively.

For some applications, such as FTP, the payload of the packet contains the source IP address and port. The IP packet header might become inconsistent with the packet payload after NAT translation. To address this issue, the Application level Gateway (ALG) is used. When NAPT is performed on an IP packet, the ALG for specified protocol also changes both the IP address and port number inside the packet. Typical protocols that require ALG include FTP, ICMP, and SIP.

In general, each subscriber initiates public network access from a private network such as a web browsing. Simultaneously, a NAT session is initiated to allow the subscriber to access the internet from the private network and vice versa allowing the requested internet content to return back to the subscriber. In this case, a NAT session is only established when a subscriber from the private network starts to transmit data to allow end to end network communication. The NAT session creates a mapping table that allows return traffic from the public network back to the subscriber in the private network. However, in some cases, traffic can be initiated from the public network to the private network without a prior NAT session. In this case, a NAT mapping must be pre-provisioned to allow external public networks to access the private network. The pre-provisioned NAT mapping is called “port forwarding”, where a private IP address and port are statically bound to a public IP address and port. External users can access services provided by private hosts by accessing the pre-provisioned public IP address and port.

In most broadband high-speed internet service today, the residential gateway (RG) performs both a routed function and a NAT function. NAT is used as a solution to conserve IPv4 public address. Service providers are required to carefully manage their public IPv4 addresses because the regional Internet registry (RIR) assigns a limited number of public IPv4 addresses per service provider. Further details of NAT can be found in IETF RFC 2663 [8] “IP Network Address Translator (NAT) Terminology and Considerations”.

3.1.1 CG-NAT Introduction

Carrier-Grade NAT (CG-NAT), also called Large-Scale NAT, is a NAT function performed in the service provider network to reduce service provider IPv4 address consumption. There are typically two use cases of CG-NAT:

1. Translation of a subscriber private IPv4 to public IPv4 address is referred to as NAT44.
2. It is common for home RG to have already performed NAT function and in this case the subscriber NAT traffic goes through NAT again at a CG-NAT within the service provider network. This is often referred to as NAT 444 (IPv4 to IPv4 to IPv4).

CG-NAT enables multiple subscribers to share one public IPv4 address. With CG-NAT, the NAT is placed in an ISP's network and translates the traffic of many subscribers. Subscribers have limited or no control over the CG-NAT. The following figure shows a typical CG-NAT deployment scenario.




Figure 1: CG-NAT NAT 44 scenario

NAT444 is the most widely deployed CG-NAT solution. The following figure shows the how NAT444 works.




Figure 2: CG-NAT NAT 444 scenario

To conserve public addresses, NAPT is often used. NAPT allows multiple subscribers to share a single public address. A CG-NAT device typically allocates a range of consecutive ports from a single public IP address to a private IP address representing a subscriber. Since the IP TCP/UDP port number ranges from 0 to 64K, if each NAT block provides 1000 ports per subscriber, a public IPv4 address can be divided into 64 NAT blocks to support 64 subscribers. Sometimes a single port range might not be adequate for a subscriber. If a subscriber exhausts the first allocated port range, the CG-NAT can allocate a new port range to the subscriber. These port ranges are called extended port ranges.

Further, using port ranges can also improve the scalability and efficiency of NAT logging.

3.1.2 CG-NAT Logging

NAT logging is a function that records the mapping between a subscriber's private address and a combination of the public address and port used on the public Internet. NAT logging is an important function when CG-NAT is integrated with BNG. RFC7422 [9] introduces an efficient and scalable way to implement NAT logging. Which can be simply called port-range based NAT logging, where NAT logging records the mapping between the private IP address, the corresponding public IP address and port range, and a time stamp. This information can be reported to the AAA or Syslog server for logging, more details can be found in RFC 5424 [12].

AAA accounting records are normally sent when subscribers go online and offline, recording public IP address allocation and release. After a subscriber goes online, if NAT information is updated, for example, an extended port range is allocated to the subscriber, an Accounting Request (Update) carrying the updated NAT information needs to be sent to the AAA server to update the corresponding NAT logs.

3.2 DBNG Functional Architecture

CG-NAT is a solution that is widely deployed in operators' broadband networks today. The following section defines how CG-NAT functions are incorporated to the DBNG architecture. The architecture allows for two different CGNAT models: DBNG-CP centric vs. DBNG-UP centric. The model key differences are discussed in section 3.3.

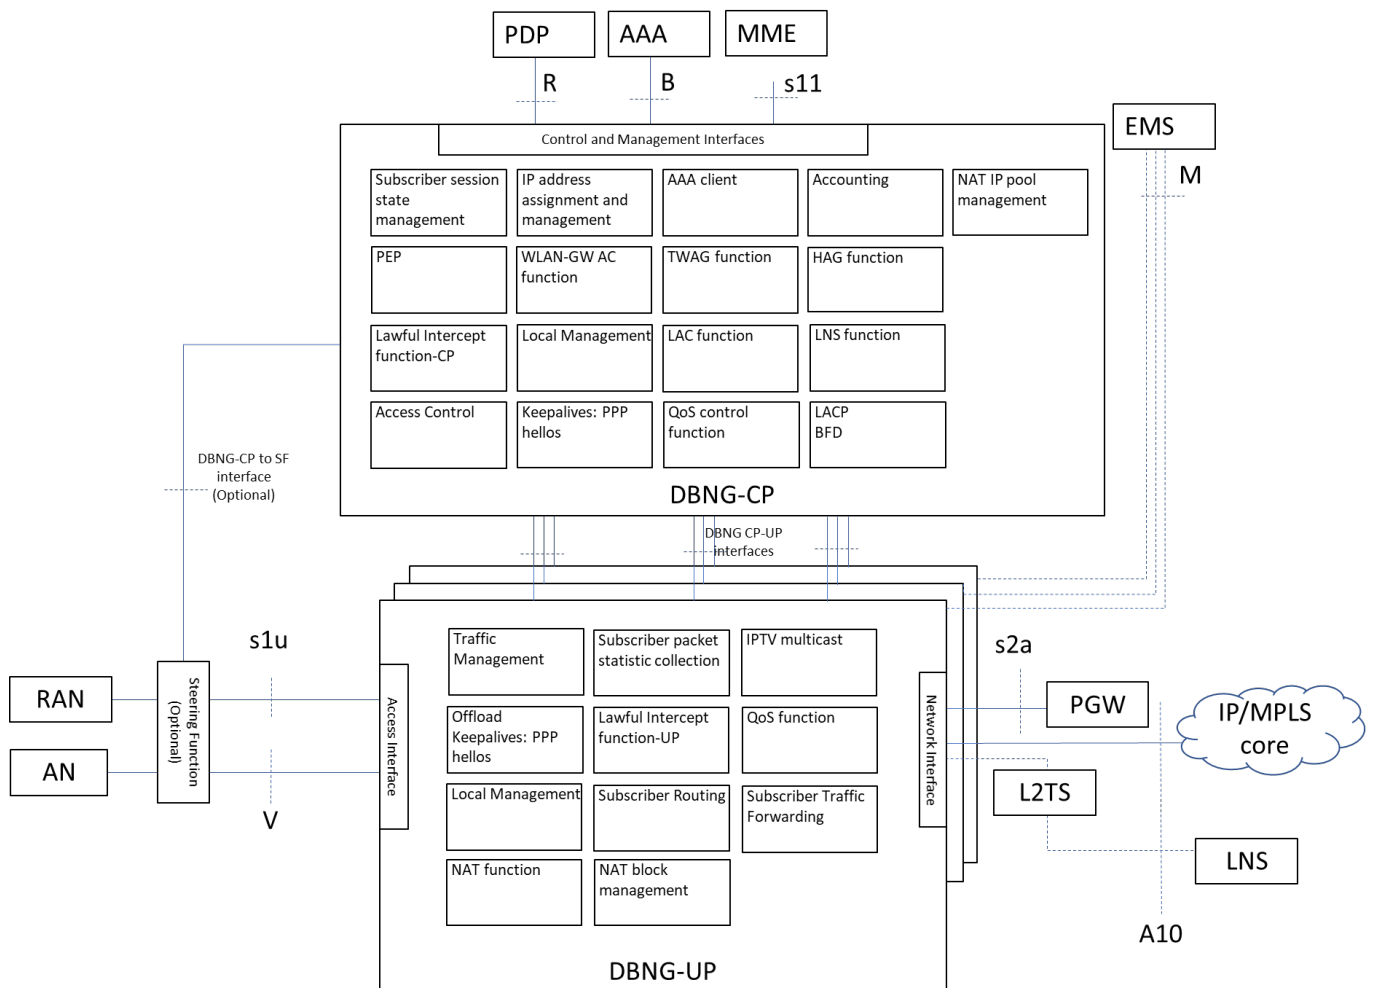


Figure 3: DBNG with integrated CG-NAT function architecture

3.2.1 Control Plane Function Overall:

The DBNG-CP is responsible for:

- NAT pool management where NAT blocks are assigned to individual subscribers residing on the DBNG-UP.
- The DBNG-CP must ensure minimum public address fragmentation as NAT pools are allocated to DBNG-UPs for routing advertisement.

3.2.2 User Plane Function Overall:

The DBNG-UP NAT function is responsible for the following:

- The DBNG-UP receives traffic forwarding instruction from the DBNG-CP to perform NAT on the subscriber data plane traffic.

3.3 DBNG CGNAT models

CG-NAT is a solution that is widely deployed in operators' broadband networks today. This Technical Report specifies two different NAT models, a DBNG-CP centric one and another a DBNG-UP centric one. Both models share the same DBNG architecture. Key facts of each model is specified below along with cost/ benefit analysis.

3.3.1 DBNG-CP centric model

In this model, the DBNG-CP manages NAT pools. For each subscriber that requires CG-NAT service, the DBNG-CP would assign the initial NAT block at subscriber session creation time. The DBNG-CP also instructs the DBNG-UP to reserve a group of NAT blocks for subscribers that require additional NAT blocks assignment. When subscriber requires additional port blocks, the DBNG-UP would assign the NAT blocks to the subscriber from the reserved port range.

Benefits

- Supports both single NAT port block and dynamic NAT port blocks assignment.
- No additional PFCP messages required for initial NAT port block assignment.
- Guarantees NAT port block is available at start of PFCP session.
- Guarantees dynamic NAT blocks are reserved for additional NAT block assignment.
- Requires only SCi interface.

Costs

- A separate PFCP message is used to inform new NAT pool to the DBNG-UP and this may temporarily stall subscriber NAT service.
- Dynamic NAT block assignment utilize a different PFCP reporting mechanism than initial NAT block assignment.
- Blocks of NAT port blocks are reserved to guarantee dynamic port block assignment might remain unused for a period of time.

Ideal use case of the DBNG-CP centric model

This model better guarantees both initial NAT block assignment and additional NAT blocks assignment. This model is the most efficient when operators assign a single NAT block per subscriber. A group of NAT blocks per public address can be reserved to better guarantee subscriber additional NAT blocks request. However, the reserved NAT blocks might be unused for a period of time.

3.3.1.1 Control Plane Centric NAT model Details:

In this model, the DBNG-CP is responsible for these additional following:

- On-Demand NAT pool allocation: DBNG-CP distribute NAT pools based on subscriber scale. There is no need to pre-allocate NAT pools to DBNG-UP.
- NAT logging accuracy: The DBNG-CP has the full subscriber information: the subscriber IP addresses, the NAT public IP address and port range, subscriber MAC address or other circuit information, and time of day information.

- Use of resource: E.g., If NAT pools are limited, DBNG-CP can halt subscriber state creation and eliminate unnecessary PFCP session establishments to DBNG-UP.
- Policy server communication: Policy server may apply unique NAT policy per subscriber. E.g., a port forwarding rule for a specific subscriber. The DBNG-CP connected directly to policy server can apply the NAT policy and forwarding rules to the DBNG-UP.
- For NAT pool reallocation: The DBNG-CP manages subscriber session state. Once a subscriber session terminates, the DBNG-CP can remove the NAT resources along with the subscriber PFCP session. There is no need to rely on NAT timers to free up NAT resources.
- For NAT port forwarding: This can be performed statically or dynamically through protocol such as PCP. The DBNG-CP is the centralized location to obtain port forwarding information from configuration, AAA policy server, or dynamic PCP servers. The DBNG-CP programs the forwarding rules on the DBNG-UP for the subscriber accordingly.
- For PFCP procedure: DBNG-CP programs traffic detection and forwarding rules to the DBNG-UP. By placing the NAT management function on the DBNG-CP, it follows the typical procedure of programming forwarding rules from the CP to the UP, in this case NAT.

3.3.2 DBNG-UP centric model

In this model, the DBNG-CP manages the NAT pools and the DBNG-UP manages the NAT blocks. The DBNG-CP allocates NAT pool(s) to DBNG-UP using subscriber NAT policies described below. NAT policy is assigned to each subscriber during subscriber login time. Based on the subscriber NAT policy, the DBNG-UP would assign each subscriber one or more NAT blocks.

Each NAT policy specifies the NAT block allocation rules and may contain:

- List of public addresses
- NAT block size
- Maximum number of NAT blocks allowed per subscriber
- NAT block aging time.

Benefits

- Any unused NAT block can be allocated to new or existing subscriber. No NAT blocks are reserved
- Consistent PFCP message used to report NAT block allocation.

Cost

- Subscriber may be denied additional NAT block request if NAT policy is not crafted properly
- NAT pool replenishment on the DBNG-UP may temporarily result in subscriber session creation failure
- Requires utilization of both SCi and Mi interface
- Two additional PFCP messages are required for initial port block allocation notification.

Ideal use case of the DBNG-UP centric model

This model better guarantees efficient use of all NAT blocks. Subscribers are allocated NAT blocks on a first-come-first-serve basis until all NAT blocks are exhausted. Since NAT blocks are never reserved, subscribers have the flexibility to request any number of NAT blocks. However, other subscribers may consequently be denied additional NAT blocks when NAT blocks are exhausted.

3.3.2.1 User Plane Centric NAT model Details:

In this model, the DBNG-CP is responsible for these additional following:

- The DBNG-CP manages NAT pools and policies and distributes these to the DBNG-UPs.
- The DBNG-CP configures routes for the public addresses from the respective DBNG-UP to reach the Internet.

- The DBNG-CP specifies the NAT policy to use for each subscriber during session creation on the DBNG-UP and reports the NAT mappings chosen by DBNG-UP for the subscriber to the Accounting or Log servers.

In this mode, the DBNG-UP is responsible for these additional following:

- The DBNG-UP performs individual public address and NAT block selection and assignment for each subscriber based on NAT policy given by DBNG-CP during session creation. The DBNG-UP locally decides additional NAT blocks for allocation and release of unused blocks, and reports block usage status to the DBNG-CP for each subscriber using SCi messages. NAT function is performed on the upstream and downstream traffic of the subscriber.

3.3.3 Common Dynamic NAT block allocation for both models

Overall requirements:

- NAT pool management
 - o To avoid IP route fragmentation in routing advertisement, it is best practice for the DBNG-CP to assign a NAT pool as a contiguous IP subnet to the DBNG-UP for NAT purpose. Each IP within the NAT pool must serve two purposes: one for initial NAT block assignment and the second for dynamic NAT block allocation.
 - Subscribers requiring NAT would start the subscriber session with an initial NAT block.
 - For subscriber that requires additional NAT blocks, each IP address will have some NAT blocks reserved for dynamic NAT block allocation. This dynamic NAT block is not pre-assigned to any particular subscriber but is shared amongst subscribers.
 - o Every public IP address must reserve NAT blocks for dynamic NAT block allocation. It is essential for subscriber to utilize the same public IP address even when additional NAT blocks are requested for most internet application.
- NAT logging
 - o To allocate additional NAT blocks to be used by subscribers, the SCi interface is used to report the current set of NAT blocks consumed by the subscriber.
 - o It is best practice the keep logging to minimal to not overload the SCi interface with NAT block allocation reporting messages.

4 Call Flows

4.1 DBNG-CP centric model call flows

The call flows are split into DBNG-CP and DBNG-UP centric model. Please note that call flows in general are informational only.

4.1.1 Initial NAT block allocation for DHCP subscribers

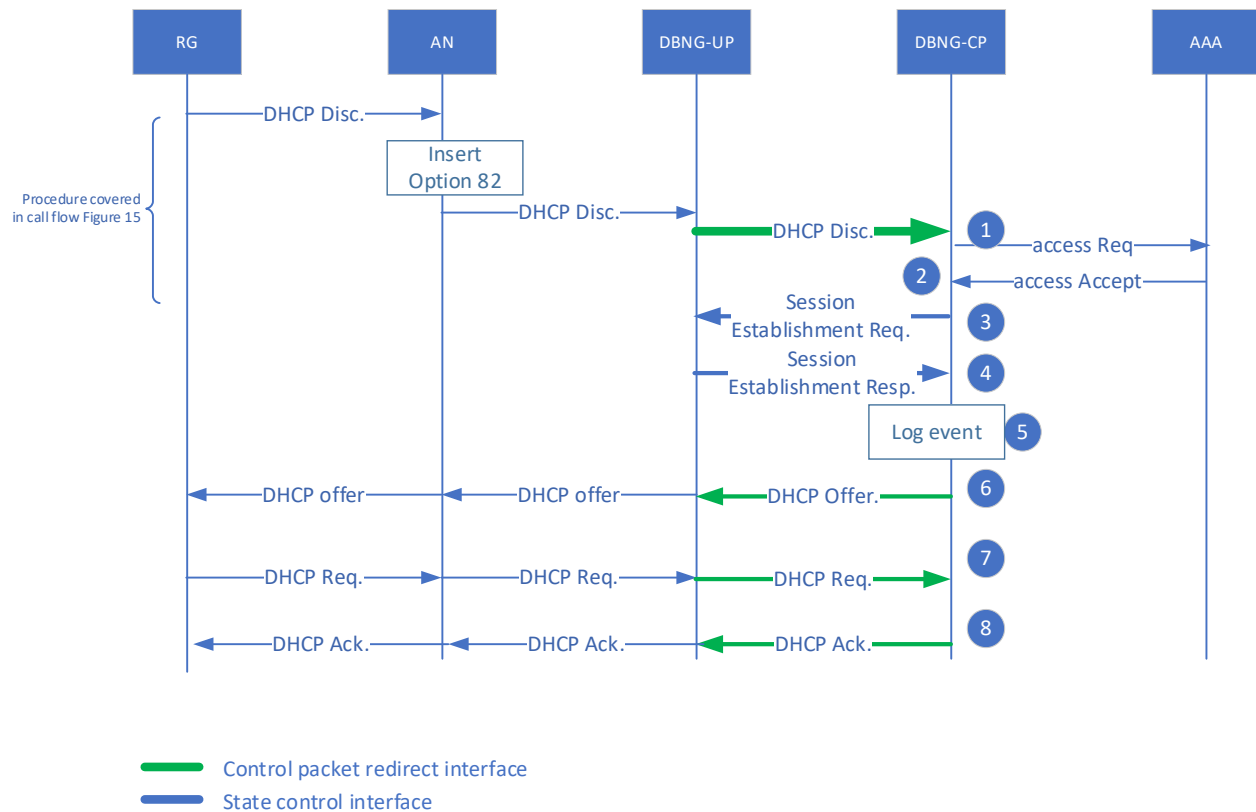


Figure 4: DBNG-CP centric model with initial NAT block allocation for DHCP subscribers

Steps:

1. Follows TR-459 [1] DHCP call flow step 1.
2. Follows TR-459 [1] DHCP call flow step 2. In addition, the AAA service might indicate a NAT policy for the subscriber.
3. Follows TR-459 [1] DHCP call flow step 3. *The DBNG-CP knows the IP address to be assigned to the RG, either obtained through the local address server or from a AAA returned attribute. At the same time, from the public IP subnet assigned to the DNBG-UP, the DBNG-CP has the IP address and port range to be assigned for the RG. At this point the DBNG-CP can send a session creation request to create new packet forwarding states for the data packet including NAT. This updates the data plane state.
4. Follows TR-459 [1] DHCP call flow step 4. 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 NAT'ed.
5. The DBNG-CP will trigger a log event which could be in a form for RADIUS accounting to report the public address and port ranges used by the subscriber with the current time.
- 6 to 8. Follows TR-459 [1] DHCP call flow step 5 to 7.

4.1.2 Initial NAT block allocation for PPPoE subscribers

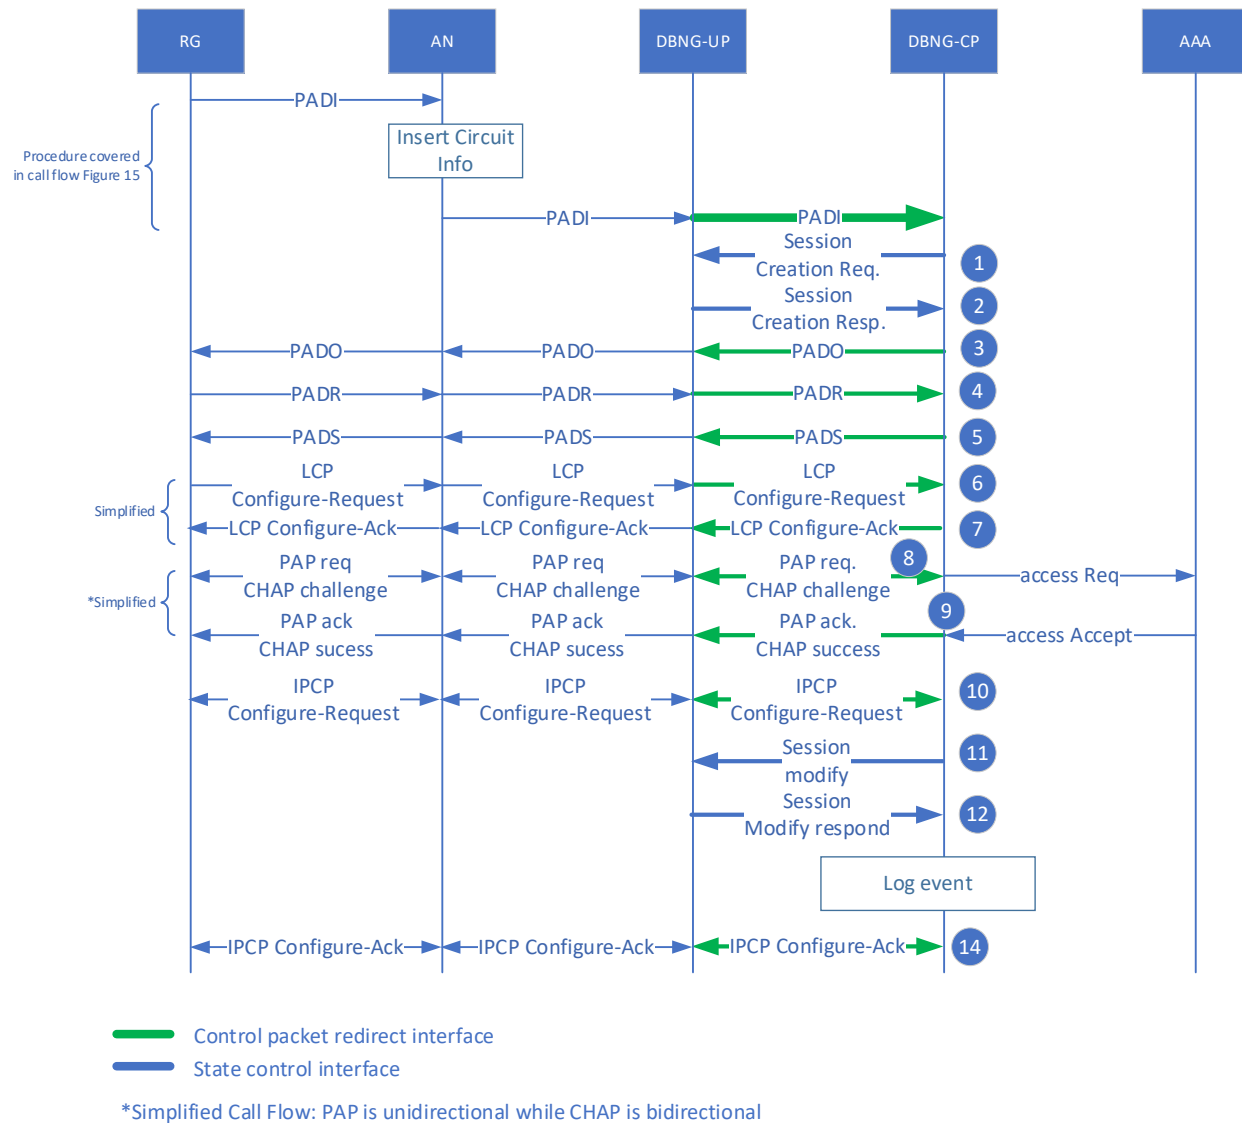


Figure 5: DBNG-CP centric model with initial NAT block allocation for PPPoE subscribers

Steps:

- 1 to 8. Follows TR-459 [1] PPPoE call flow step 1-8.
- 9 Follows TR-459 [1] PPPoE call flow step 9. In addition, AAA can return a NAT policy.
- 10 Follows TR-459 [1] PPPoE call flow step 10.
- 11 Follows TR-459 [1] PPPoE call flow step 11. *The DBNG-CP knows the IP address to be assigned to the RG, either obtained through the local address server or from a AAA returned attribute. At the same time, the DBNG-CP also has the public IP address and port range to assigned to the RG. At this point the DBNG-CP can send a session creation request to create new packet forwarding states for the data packet including NAT. This updates the data plane state.
- 12 Follows TR-459 [1] PPPoE call flow step 12. 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 NAT'ed.

- 13 The DBNG-CP will trigger a log event which could be in a form for RADIUS accounting to report the public address and port ranges used by the subscriber with the current time.
- 14 Follows TR-459 [1] PPPoE call flow step 13.

4.1.3 Session termination

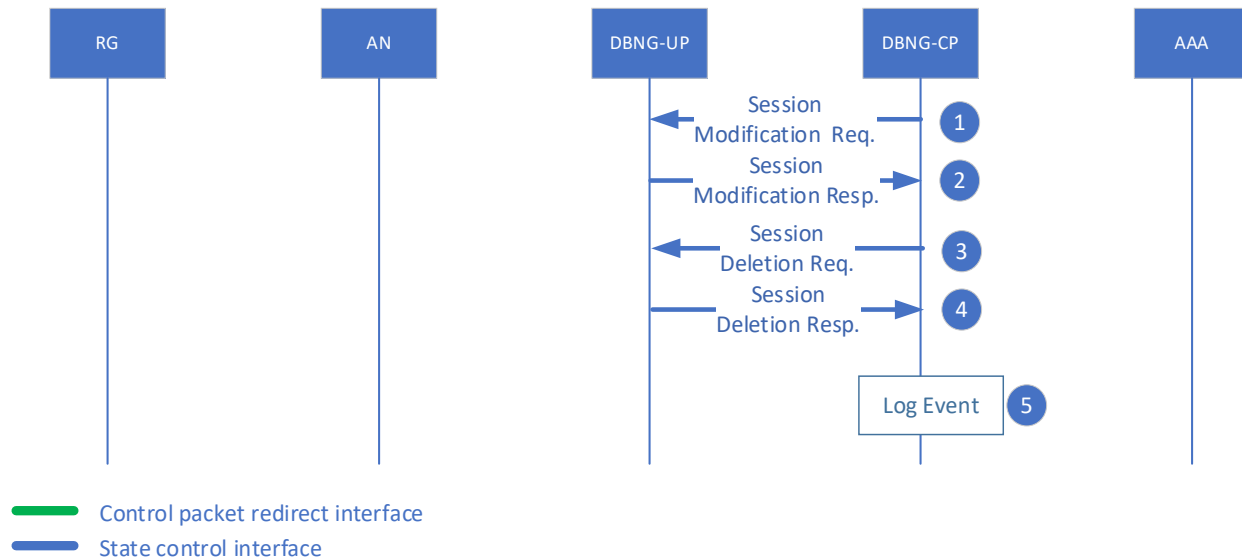


Figure 6: DBNG-CP centric model subscriber session termination

Steps:

1. [conditional] If the subscriber terminates the IPoEv4 session only, the DBNG-CP will trigger a session modification request to remove the traffic rules for the IPoEv4 session including the NAT public address and NAT block.
2. [conditional] If the subscriber terminates the IPoEv4 session only, the DBNG-UP acknowledge the request and any NAT block allocated for the subscriber is released. Now skip to step 5.
3. [conditional] If the DBNG-CP triggers a session deletion for the entire IPoE session or PPPoE session. The termination of the session could be triggered by the subscriber, by the operator, or the network.
4. [conditional] If the DBNG-CP triggers a session deletion for the entire IPoE session or PPPoE session. The DBNG-UP notifies the DBNG-CP that the subscriber session is removed on the DBNG-UP. The NAT block allocated to the subscriber is released for future allocation.
5. The DBNG-CP sends a log message capturing the termination event which includes the list of NAT block(s) used by the subscriber.

4.2 DBNG-UP centric model

4.2.1 Initial NAT block allocation for DHCPv4 Subscriber

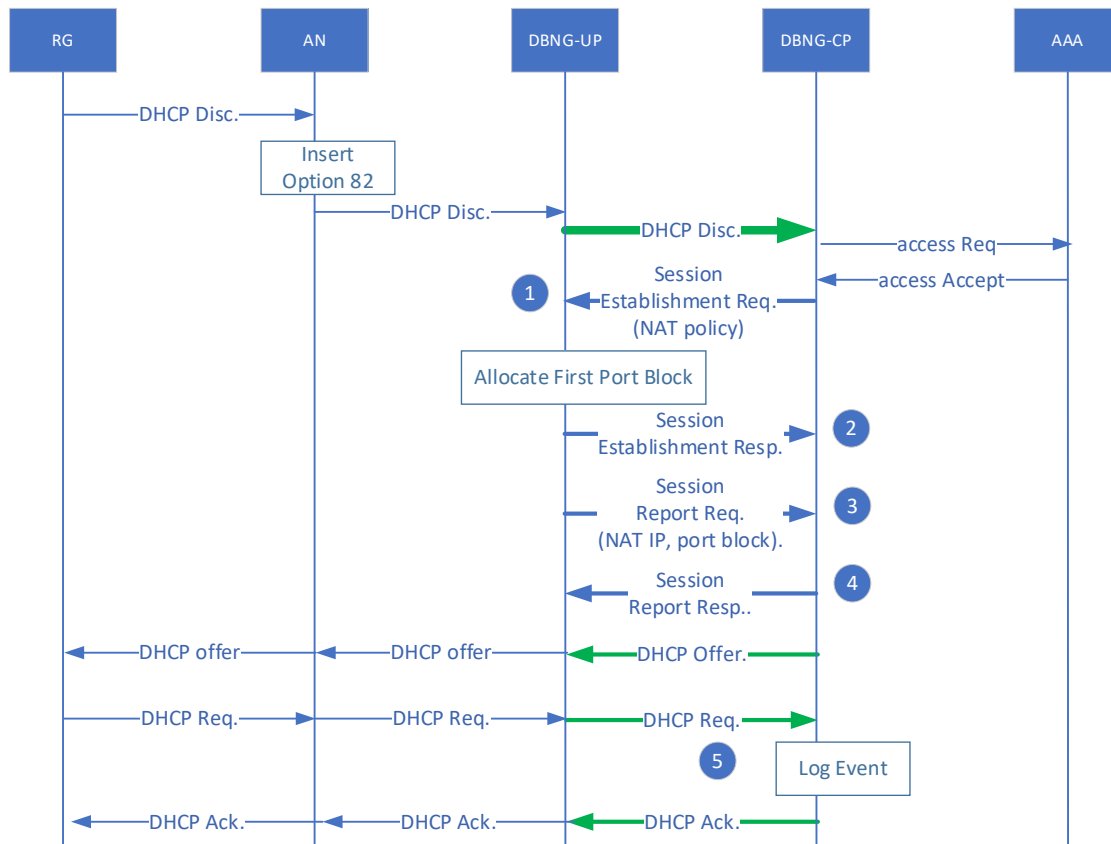


Figure 7: DBNG-UP centric model with initial NAT block allocation for DHCP subscribers

IPoE DHCPv4 subscriber connects and on successful authentication is assigned a private IP address on the DBNG-CP.

1. The DBNG-CP selects a NAT policy for the subscriber session*. The policy is in the Forwarding Policy IE and is sent in Session Establishment Request message.
2. DBNG-UP sends the Session Establishment Response.
3. DBNG-UP based on the NAT policy of subscriber allocates NAT public address and first NAT block and sends the NAT block information in Session Report Request message**.
4. DBNG-CP sends the Session Report Response
5. The DBNG-CP sends the RADIUS Acct Start message to AAA Server for the subscriber including the private IP address, NAT public address and NAT block information. Subscriber originates data traffic which makes use of the first NAT block.

***Note:** The NAT policy name can come from AAA server.

****Note:** Alternatively, the first block can be allocated by the DBNG-UP on first data packet and reported in a PFCP Session Report Request later.

4.2.2 Initial NAT block allocation for PPPoE Subscriber

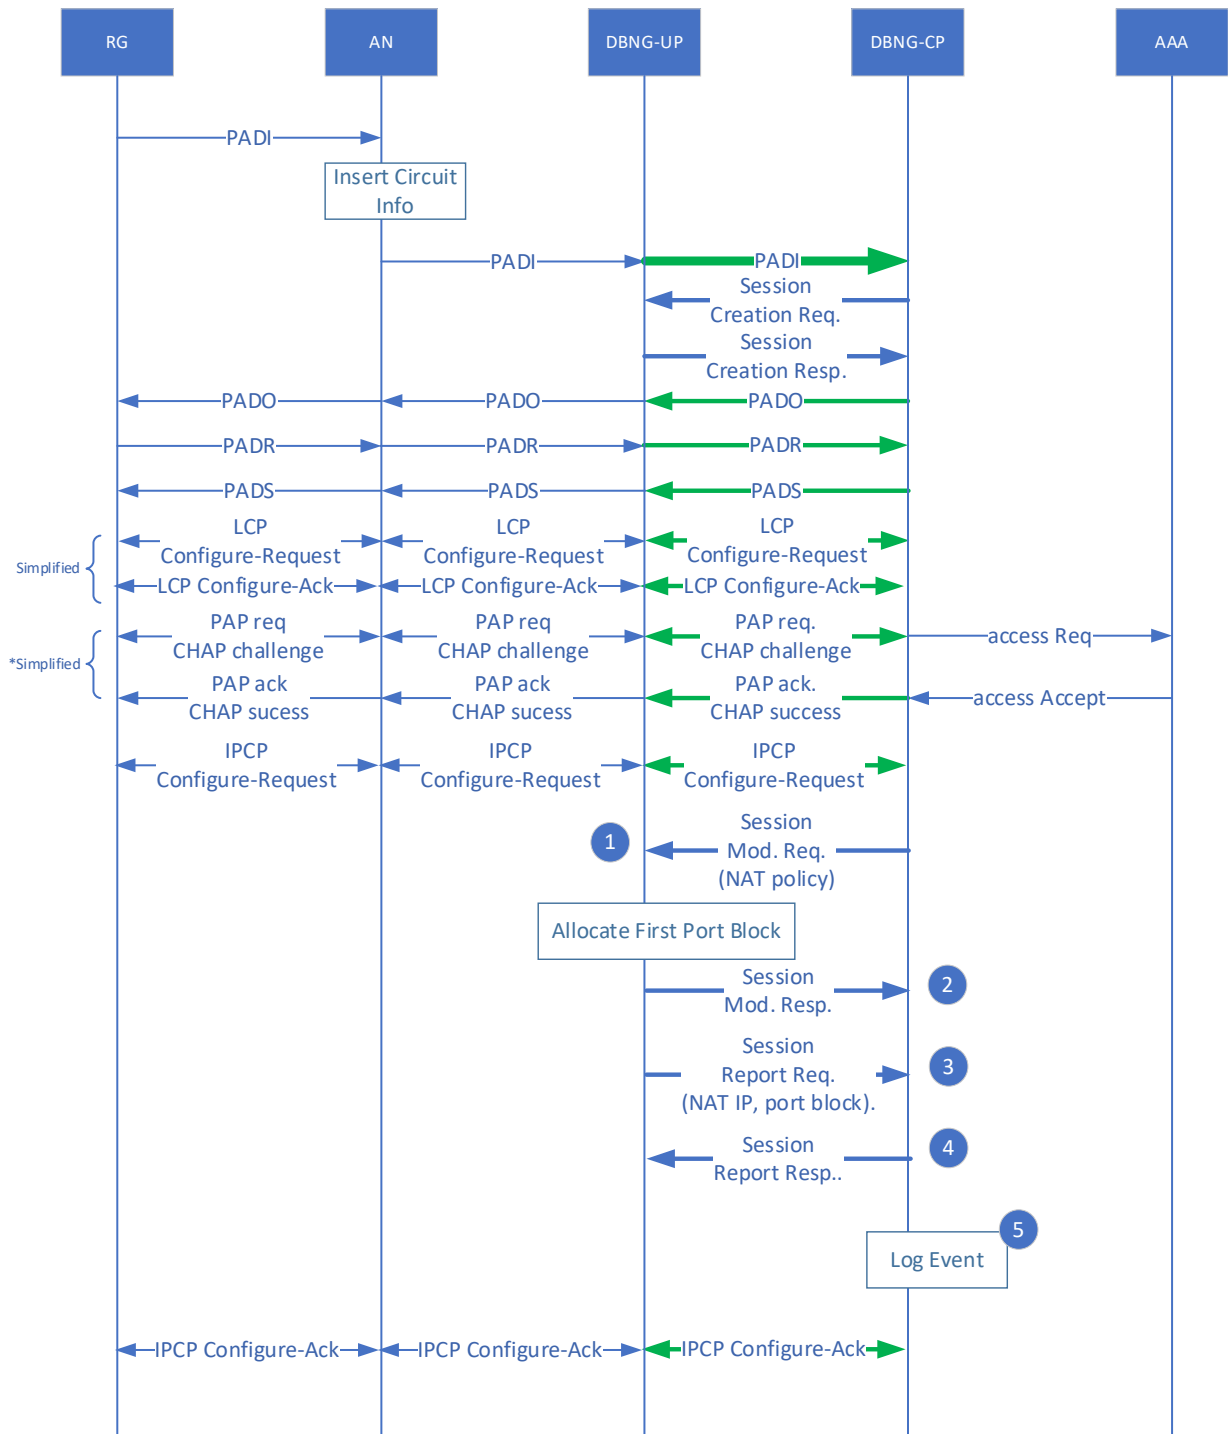


Figure 8: DBNG-UP centric model with initial NAT block allocation for PPPoE subscribers

PPPoE subscriber connects and on successful authentication is assigned a private IP address on the DBNG-CP.

1. On receipt of PPP IPCP Conf Request, DBNG-CP based on private IP address populates NAT policy name* in the Forwarding Policy IE and sends in Session Modification Request message.
2. DBNG-UP sends the Session Modification Response.
3. DBNG-UP based on the NAT policy of subscriber allocates NAT public address and first NAT block and sends the NAT block information in Session Report Request message**.
4. DBNG-CP sends the Session Report Response.
5. The DBNG-CP sends the RADIUS Acct Start message to AAA Server for the subscriber including the private IP address, NAT public address and NAT block information. Subscriber originates data traffic which makes use of the first NAT block.

***Note:** The NAT policy name can come from AAA server.

****Note:** Alternatively, the first block can be allocated by the DBNG-UP on first data packet and reported in a PFCP Session Report Request later.

4.2.3 Session termination

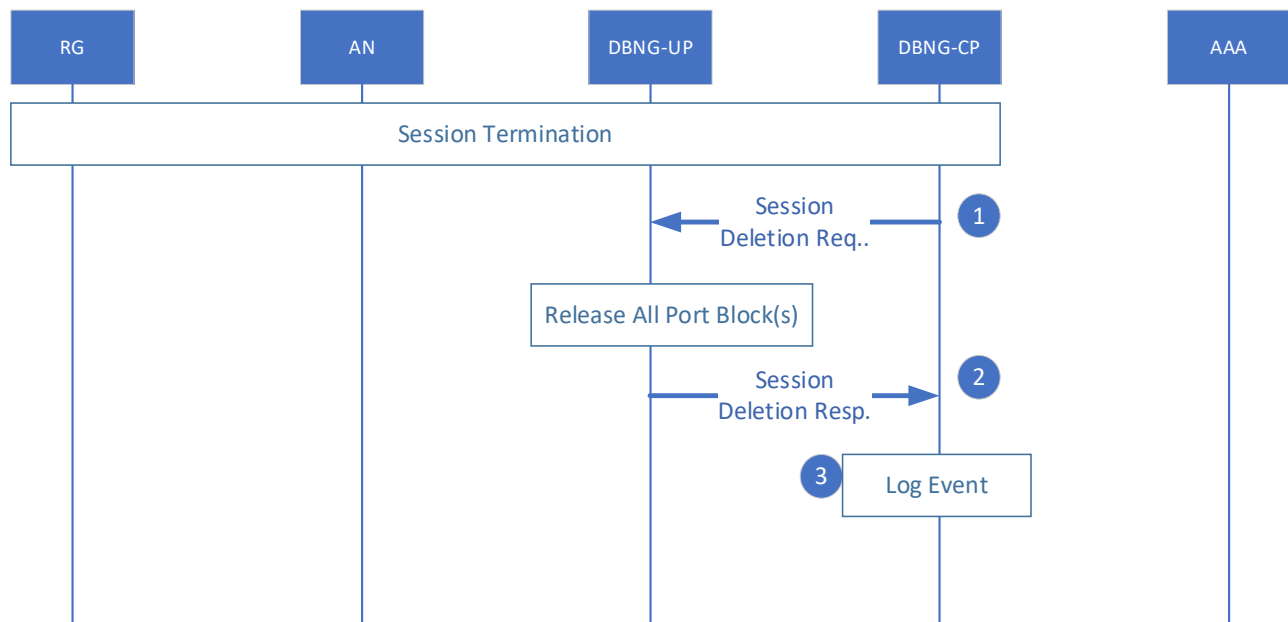


Figure 9: DBNG-UP centric model subscriber session termination

Subscriber disconnects from the DBNG-CP by sending a DHCP Release or PPPoE PADT.

1. DBNG-CP sends PFCP Session Deletion Request message to DBNG-UP.
2. DBNG-UP releases NAT blocks and bindings and sends Session Deletion Response message.
3. The DBNG-CP sends a RADIUS Acct Stop message to the AAA server for the subscriber with NAT blocks reported from the last PFCP Session Report.

4.3 Common Dynamic NAT block allocation for all models

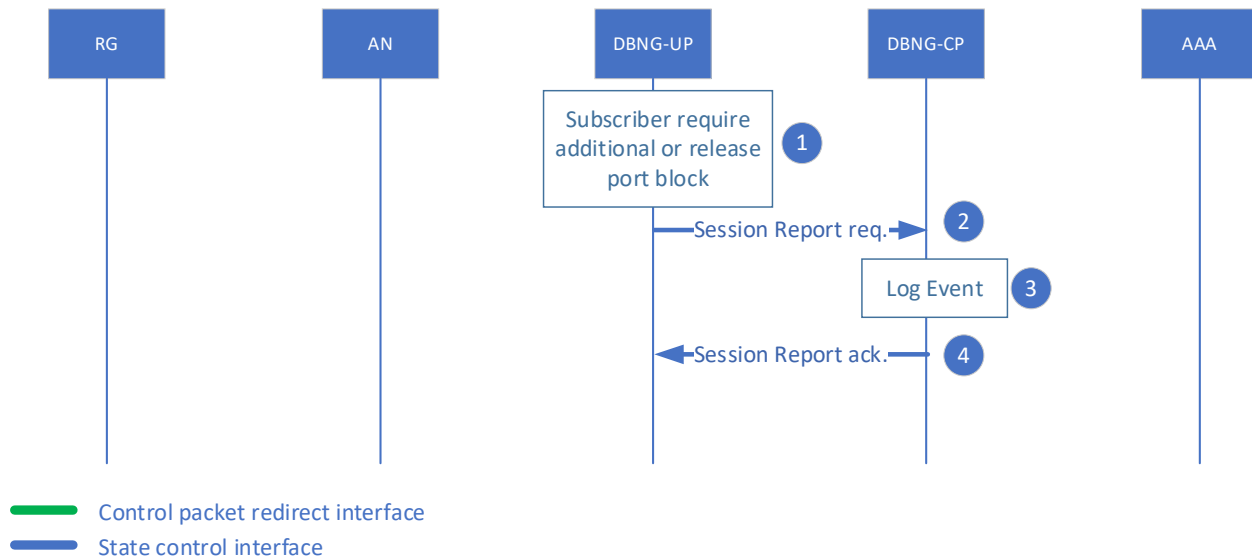


Figure 10: Dynamic NAT block allocation for all models

Steps:

1. After DHCP or PPPoE process completes, it is possible for some subscribers to have policies that allows them to consume additional NAT blocks.
2. When the UP allocates new NAT block to the subscriber, the UP will use Session Report message to inform the CP.

It is also possible for the subscriber to release previously allocated NAT block and the DBNG-UP can reclaim previously allocated NAT block. This will also trigger a log event with the use of a Session Report request message.

3. The Session Report request message will contain NAT event which includes the current set of all NAT blocks consumed the time event of the allocation.
4. The DBNG-CP sends the entire list of NAT block range(s) to the log server.
5. The DBNG-CP acknowledges the Session Report message.

4.4 Establishment of resilient session (CP-centric model and UP-centric model)

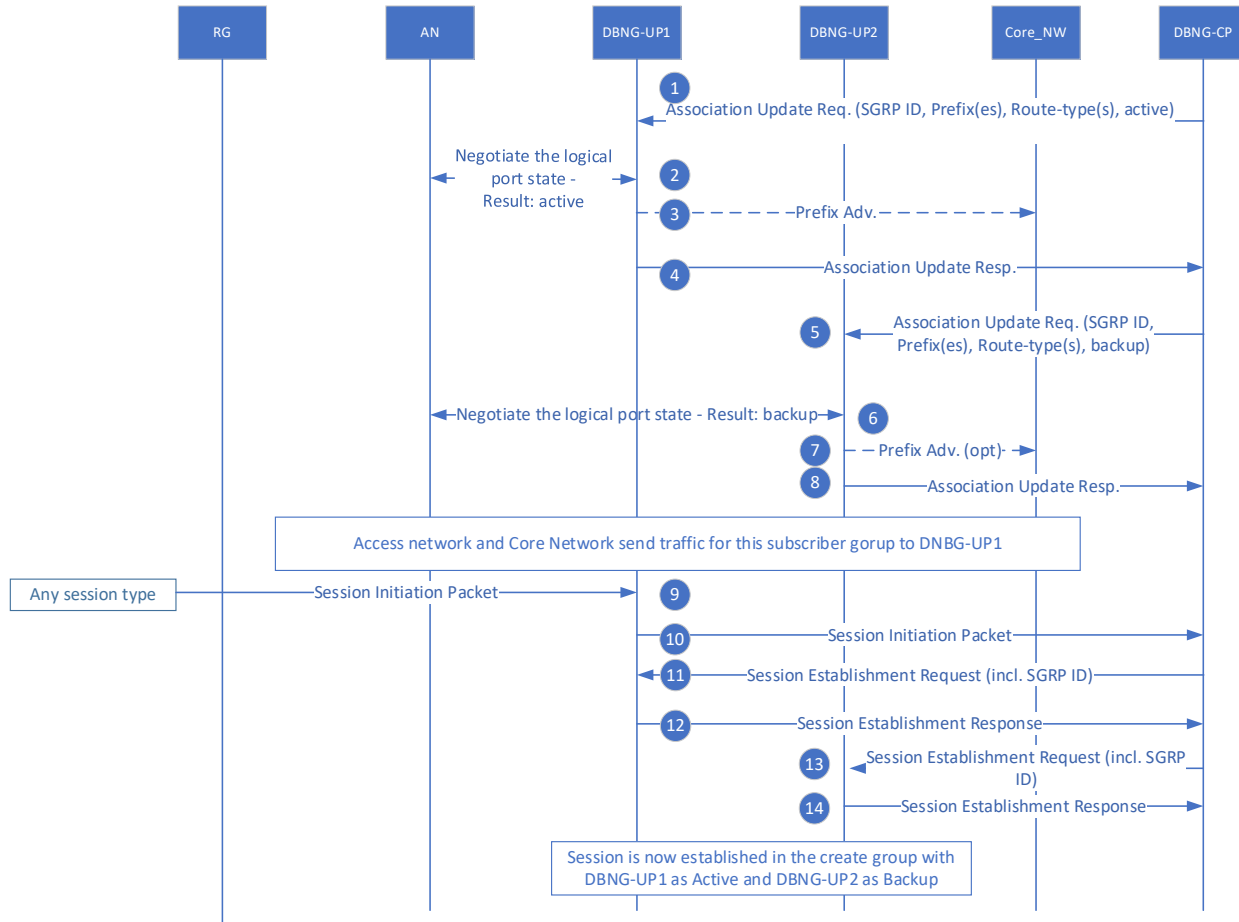


Figure 11: Establishment of resilient session (CP-centric model and UP-centric model)

Steps 1-14 follow the same procedure as TR-459 issue section 4.5.47.1. Below are the key differences.

Step 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. Each of the prefixes would identify their route type (private or public). The DBNG-UP1 advertises the prefix to the core network based on local route policies.

Step 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. Each of the prefixes would identify their route type (private or public). The DBNG-UP2 advertises the prefix to the core network based on local route policies.

5 Technical Requirements

Below is a list of key requirements to support the carrier grade NAT function on DBNG.

5.1 General NAT requirement on DBNG

- [R-1] NAT pool MUST be dynamically allocated from DBNG-CP to the DBNG-UP.
- [R-2] NAT pool allocation to a DBNG-UP MUST be able to scale up or down based on the number of subscribers.
- [R-3] NAT pools which are no longer required MUST be made available for other DBNG-UP to use.
- [R-4] The DBNG-CP MUST be able to reserve a range of ports per outside IP address for static port forwarding purpose.
- [R-5] A NAT policy MUST inform the DBNG-CP that the subscriber requires static port forwarding.
- [R-6] If a subscriber requires a certain port forwarding mapping, then the DBNG-CP MUST ensure that the assigned outside IP will have the port available for assignment.
- [R-7] The DBNG-CP MUST support both DBNG-CP and DBNG-UP centric models for CG-NAT.
- [R-8] The DBNG-UP MUST support either DBNG-CP and DBNG-UP centric models for CG-NAT.
- [R-9] The DBNG-CP MUST support the CG-NAT models identified by the DBNG-UP using the BBF UP function feature flag.
- [R-10] In the case where a DBNG-UP can support both DBNG-CP and DBNG-UP centric models and sends both BBF UP function features flags, the DBNG-CP MUST select either DBNG-CP centric model or DBNG-UP centric model for that particular DBNG-UP instance and MUST only exchange PFCP IEs associated to the selected model.

NAT logging can be in the form of system logs or in the form of RADIUS accounting messages. The log will contain information of the subscriber private address mapped a port range of a public address. The logging must contain enough information to identify the subscriber. For example, in the case where duplicate private IP are used, additional identifiers are required such as MAC address and circuit information. The NAT logging can be sent to log servers or within RADIUS accounting message.

- [R-11] NAT logging MUST be supported and include subscriber private IP, public IP, public port range, and time of day.

5.2 DBNG-CP requirements

5.2.1 DBNG-CP centric model

- [R-12] The DBNG-CP MUST manage NAT pool to be allocated per DBNG-UP.
- [R-13] The DBNG-CP MUST assigning the initial public address and port range per subscriber.
- [R-14] The DBNG-CP MUST assign NAT policy per subscriber using the SCi Session Establishment Request message.
- [R-15] The DBNG-CP MUST be able to assign multiple NAT blocks from different NAT pool to a subscriber.
- [R-16] The DBNG-CP SHOULD assign public IP prefix per DBNG-UP to optimize routing advertisement.
- [R-17] The DBNG-CP MUST provide NAT logging capability.

5.2.2 DBNG-UP centric model

- [R-18] The DBNG-CP MUST manage and assign NAT pools for each DBNG-UP supporting NAT and configure corresponding routes to each DBNG-UP.
- [R-19] The DBNG-CP MUST manage NAT policies using the Mi.
- [R-20] The DBNG-CP MUST be able to assign NAT policy per subscriber using the SCi Session Establishment Request message.
- [R-21] The DBNG-CP MUST support reporting NAT block allocations/ releases for each subscriber to the Accounting/ Logging server based on the SCi Session Report messages from DBNG-UP.

5.2.3 Common dynamic NAT block requirements for both models

- [R-22] The DBNG-CP MUST monitor NAT pool usage for each DBNG-UP by allocating additional addresses and reclaiming unused addresses from the DBNG-UP.
- [R-23] The DBNG-CP MUST be able to assign outside IP subnet to the DBNG-UP for NAT. Where each outside IP address would reserve some NAT blocks for dynamic NAT block allocation.
- [R-24] The DBNG-CP MUST provide NAT logging capability for dynamic NAT block allocation.

5.2.4 Subscriber group (SGRP) requirements for both models

- [R-25] The DBNG-CP MUST be able to group subscribers together that share both private and public IP prefixes.
- [R-26] In addition to [R-25], the DBNG-CP MUST be able to control the resiliency state of both the private and public IP prefixes.

Public prefix requirements

- [R-27] The DBNG-CP MUST be able to assign public IP prefixes to the DBNG-UP for routing advertisement purposes.
- [R-28] The DBNG-CP MUST be able to control routing advertisement based on the resiliency status of the IP prefix.

Private prefix properties

- [R-29] The DBNG-CP MUST be able assign a virtual MAC address for the private IP prefix for resiliency purposes.
- [R-30] The DBNG-CP MAY use the SCi protocol to advertise private IP prefix(es).

5.3 DBNG-UP requirements

5.3.1 DBNG-CP centric model

- [R-31] The DBNG-UP MUST support at least one NAT block.

5.3.2 DBNG-UP centric model

- [R-32] The DBNG-UP MUST dynamically select an initial NAT blocks based on the NAT policy rules for the subscriber provided by DBNG-CP in the SCi Session Establishment Request message.
- [R-33] The DBNG-UP MUST report every new NAT block allocation for the subscriber using the SCi Session Report message.
- [R-34] The DBNG-UP SHOULD support releasing of NAT blocks before session termination, such that NAT block releases for the subscriber are reported using SCi Session Report messages.
- [R-35] The DBNG-UP MUST release all currently assigned NAT blocks for the subscriber on receipt of SCi Session Deletion Request.

5.3.3 Common dynamic NAT block requirements for both models

- [R-36] The DBNG-UP MUST be able to assign new NAT blocks for subscribers that have fully utilized their current NAT blocks.
- [R-37] The DBNG-UP MUST be able to release NAT blocks that are no longer used by subscribers.
- [R-38] If there is a new NAT block allocation or a NAT block deallocation, the DBNG-UP MUST report to the DBNG-CP the most up to date NAT blocks used.

5.4 Management Interface requirements

- [R-39] The DBNG-CP MAY assign one or more NAT pool to each DBNG-UP for routing advertisements.
- [R-40] In the case of DBNG-UP performing initial NAT block allocation, the DBNG-CP MUST configure one or more NAT policies on each DBNG-UP. Each NAT policy specifies the NAT block allocation rules – public address(es) to use, NAT block size, maximum number of blocks for a private address, NAT block aging time (if any), etc.

5.5 NAT Logging Timestamp requirements

- [R-41] The DBNG MUST be able to log NAT messages with granularity as specified in Section 6.6.9.
- [R-42] The DBNG-CP MUST support NTP be able to include the timestamp for logging NAT messages.
- [R-43] The DBNG-UP MUST support NTP be able to include the timestamp for logging NAT messages.

6 PFCP CUPS protocol

Note: Advertisement of the inside subscriber prefix and outside subscriber prefix is performed by static provisioning on the DBNG-UP by utilizing interfaces such as Mi. The use of PFCP for subscriber prefix advertisement is for further study.

6.1 DBNG-CP centric model use cases and IEs exchanges

6.1.1 Use case: Initial NAT block allocation

NAT can be used in combination with IPoE and/or PPPoE subscriber session. The data forwarding rules are split between control packet forwarding and data traffic forwarding. Control packets must be redirected through the CPR Interface to DBNG-CP for address assignment. And subscriber IPoE data traffic is forwarded through the network interface.

6.1.1.1 PFCP Control Packet Redirection Rule

Control packet redirection PFCP rules follow the IPoE TR-459 section 6.3.2.1 [1] for IPoE and section 6.3.2.2 [1] for PPPoE based subscriber session.

6.1.1.2 PFCP Data Packet Forwarding Rule

Data packet forwarding PFCP rules follow the general description highlighted in TR-459 section 6.2.2 [1]. Below are further extensions required to support IPoE NAT data forwarding:

- **BBF Outer Header Removal:** BBF IE in TR-459 [1] to remove the Ethernet header from data packets before forwarding to the network interface. Details of the extension are in TR-459 [1]
- **BBF Outer Header Creation:** BBF IE in TR-459 [1] to construct the Ethernet header for packet forwarding to the subscriber. Details of the extension are in TR-459 [1]
- **BBF Apply Action:** BBF IE extension to perform the NAT function
- **Forwarding Parameters IE extensions required:** In the downstream PDR from the DBNG-UP to the RG, the Forwarding parameter would include the **BBF NAT Port Forward IE**. Detailed information of the extension is in TR-459 section 6.5 [1]. The existing conditional "Forwarding Policy" (IE Type = 41) will carry the NAT policy name 3GPP 29.244 Section 7.5.2.3 [2].
- **Traffic-Endpoint IE extensions required:** In the upstream PDR from the RG to the DBNG-UP, the traffic endpoint of the access would include the Ethernet header information such as C-Tag, S-Tag, IP address, and the logical port of the subscriber traffic. And the linked traffic endpoint of the network would include the public IP address (UE IP address IE) and the **BBF NAT External Port Range IE**. In the downstream PDR, the traffic endpoint would be public IP address (UE IP address IE) and the **BBF NAT External Port Range IE**. Detailed information of the extension is in TR-459 section 6.5 [1].

6.1.2 Use case: Dynamic NAT block allocation

The DBNG-CP at PFCP session establishment request will include a create session reporting report (SRR) to request the DBNG-UP to report the current NAT block. When a subscriber allocates or deallocates a NAT block, the DBNG-UP must utilize PFCP session reporting request mechanism. Afterwards, the DBNG-CP will reply to the DBNG-UP with a session report response to acknowledge the request.

6.1.2.1 PFCP Data Packet Forwarding rule

The rule will be same as 6.1.1.2 and will include the following additional IE to support Dynamic port block allocation. Please note Dynamic port block allocation will also require the use of Session Reporting.

- **Traffic-Endpoint IE extensions required:** In the upstream PDR from the RG to the DBNG-UP, the traffic endpoint of the access would include the Ethernet header information such as C-Tag, S-Tag, IP address, and the logical port of the subscriber traffic. And the linked traffic endpoint of the network would include the public IP address (UE IP address), the **BBF NAT External Port Range IE**, and the **BBF Dynamic NAT Block Port Range**. In the downstream PDR, the traffic endpoint would be public IP address (UE IP address IE) and the **BBF NAT external Port Range IE** and the **BBF Dynamic NAT Block Port Range**. Detailed information of the extension is in TR-459 section 6.5 [1].

6.2 DBNG-UP centric model use cases and IE exchanges

6.2.1 Use case: Initial and dynamic NAT block allocation

This section describes the PFCP extensions required to support NAT Outside IP address and NAT block allocation from UP.

The PFCP Session Establishment Request requires additional (optional) information for CG-NAT purpose.

6.2.1.1 PFCP Data Packet Forwarding Rule

Data packet forwarding PFCP rules follow the general description highlighted in section 6.3 of TR-459 [1] for DHCP and PPPoE. Below are further extensions required to support NAT data forwarding.

The FAR shall be extended to activate NAT:

- **BBF Apply Action:** BBF IE extension to perform the NAT function
- **Forwarding Parameters:** Refer 3GPP 29.244 Section 7.5.2.3 [2]. The existing conditional "Forwarding Policy" (IE Type = 41) will carry the NAT policy name.

6.3 Common IEs NAT block allocation/deallocation reporting

6.3.1 Create PFCP IE Session Report Rule (SRR) Extensions

The Session Report Rule has two use case:

1. In the DBNG-CP initial NAT block allocation use case, it is used to report additional NAT block allocation by the subscriber. It is also used to report NAT block released that are no longer used by the subscriber.
2. Use by the DBNG-UP initial NAT block allocation to report both the initial and additional NAT block

Create SRR IE is extended to instruct the DBNG-UP to report NAT events depending on the use case specified above. The IE Required to support BBF CG-NAT use case is shown below.

- **SRRID** – The ID assigned to the Session Report Rule.
- **Traffic Endpoint** - The Traffic Endpoint (public address) to report the NAT event.
- **BBF Report Trigger:** To report NAT block status changes.

6.3.2 PFCP session reporting rule

The PFCP Session Report IE shall be originated by the DBNG-UP referencing the SRRID received in Session Establishment Request.

Session Report IE shall be used to report NAT outside IP/ NAT block changes:

- **SRRID** – The ID assigned to the Session Reporting Rule.
- **BBF NAT Outside Address**: This is only applicable to DBNG-UP based initial NAT block allocation and it is used to report public/ outside IP assigned for subscriber.
- **BBF NAT External Port Range**: BBF IE NAT block(s).
- **Event Time Stamp**: IE as defined in 3GPP 29.244 [2].

6.4 PFCP IEs requirement for NAT SGRP

[R-44] PFCP BBF SGRP IE MUST be supported for NAT DBNG-UP resiliency.

The assignment of the prefix of type public or private would use the PFCP IE BBF UP subscriber prefix defined in [1]. Below values mentioned in PFCP IE BBF Prefix Type defined in [1] are applicable:

0x1: CGNAT Private Prefix

0x2: CGNAT Public Prefix

Note: the IPv6 prefix will not be applicable for NAT.

The same SGRP ID MUST be used to advertise the public and private prefix for a subscriber.

6.5 PFCP IE summary

The tables below are consistent and reproduced from TR-459 [1] with additions for CGNAT.

Table 1: BBF UP Function Features Flag and applicable BBF PFCP IEs

BBF UP Function Features flag	IE Type value (Decimal)	Information Elements	Section
NAT-CP	32787	BBF Apply Action	6.6.5
	32788	BBF NAT External Port Range	6.6.6
	32789	BBF NAT Port Forward	6.6.7
	32790	BBF Report Trigger	6.6.3
	32791	BBF Dynamic NAT Block Port Range	6.6.8
	32792	BBF Event Time Stamp	6.6.9
NAT-UP	32787	BBF Apply Action	6.6.5
	32788	BBF NAT External Port Range	6.6.6
	32790	BBF Report Trigger	6.6.3
	32792	BBF Event Time Stamp	6.6.9

If the DBNG-UP supports CGNAT functionality, the following IEs are to be supported on both DBNG-CP and DBNG-UP.

Table 2: PFCP Session Establishment Request and BBF Grouped IEs structure

PFCP message type	PFCP Grouped IEs and IEs		P	Source Reference
	Create FAR (Grouped)	BBF Apply Action	C	6.6.5

PFCP Session Establishment Request		Forwarding Parameters (Grouped)	BBF NAT port forward	C	6.6.7
			Forwarding Policy	C	3GPP TS 29.244 8.2.23 [2]
	Create Traffic Endpoint (Grouped)	BBF NAT External Port Range		C	6.6.6
		BBF Dynamic NAT Block Port Range		C	6.6.8
	Create SRR (Grouped)	SRR ID		C	3GPP TS 29.244 7.5.2.9 [2]
		Traffic Endpoint ID		C	3GPP TS 29.244 8.2.92 [2]
		BBF Report Trigger		C	6.6.3

The following IEs can be sent from the DBNG-CP to the DBNG-UP during a subscriber mid-session. The use cases are:

- Modify existing NAT port forwarding rules for the subscriber
- Allowing changes to NAT blocks

Table 3: PFCP Session Modify Request and BBF Grouped IEs structure

PFCP message type	PFCP Grouped IEs and IEs			P	Source Reference
PFCP Session Modify Request	Create FAR (grouped)	Forwarding Parameters (grouped)	BBF NAT port forward	C	6.6.7
			Forwarding Policy	C	3GPP TS 29.244 8.2.23 [2]
	Create Traffic Endpoint (Grouped)	BBF NAT External Port Range		C	6.6.6
		BBF Dynamic NAT Block Port Range		C	6.6.8
	Update FAR (grouped)	Update Forwarding Parameters (grouped)	BBF NAT Port Forward	C	6.6.7
			Forwarding Policy	C	3GPP TS 29.244 8.2.23 [2]
	Update Traffic Endpoint (grouped)	BBF NAT External Port Range		C	6.6.6
		BBF Dynamic NAT Block Port Range		C	6.6.8

Table 4: PFCP Session Report Request and BBF Grouped IEs structure

PFCP message type	PFCP Grouped IEs and IEs			P	Source Reference
Session Report Request	Report Type			M	3GPP TS 29.244 8.2.21 [2]
	Session Report (Grouped)	SRR ID		M	3GPP TS 29.244 8.2.151 [2]
		BBF NAT Outside Address		C	3GPP TS 29.244 8.2.62 [2]
		BBF NAT External Port Range		C	6.6.6
		BBF Event Time Stamp		C	6.6.9

C is conditional to the use case

If the DBNG-UP supports CGNAT functionality, the following IEs are to be supported depending on the use case outlined below.

CP UL: control packets uplink from access to DBNG-CP

CP DL control packets downlink from DBNG-CP to access

DP UL: data packets uplink from access to network

DP DL: data packets downlink from network to access

Create SRR: Create SRR are sent from DBNG-CP to DBNG-UP

UP Session Report: Session Reports are sent from DBNG-UP to DBNG-CP

Table 5: BBF Extended Information Element and Applicability Table for association and default redirection tunnels for NAT use case

IE Type value (Decimal)	Information Elements	Section	Use Case	
			CP and UP Association	Default CPR PDR CP UL
32768	BBF UP Function Features IE	6.6.1	Yes	No
32787	BBF Apply Action	6.6.5	No	No
32788	BBF NAT External Port Range	6.6.6	No	No
32789	BBF NAT Port Forward	6.6.7	No	No
32790	BBF Report Trigger	6.6.3	No	No
32791	BBF Dynamic NAT Block Port Range	6.6.8	No	No
32792	BBF Event Time Stamp	6.6.9	No	No

Table 6: BBF Extended Information Element and Applicability Table for DBNG-CP centric NAT model

IE Type value (Decimal)	Information Elements	Section	Use Case: DHCPv4, PPPoEv4, LNS subscriber session CGNAT					
			CP UL PDR	CP DL PDR	DP UL PDR	DP DL PDR	Create SRR	UP Session Report
32768	BBF UP Function Features IE	6.6.1	No	No	No	No	No	No
32787	BBF Apply Action	6.6.5	No	No	Yes	Yes	No	No
32788	BBF NAT external Port Range	6.6.6	No	No	Yes	No	No	Yes
32789	BBF NAT Port Forward	6.6.7	No	No	No	Yes	No	No
32790	*BBF Report Trigger	6.6.3	No	No	No	No	Yes	No
32791	*BBF Dynamic NAT block port range	6.6.8	No	No	Yes	No	No	No
32792	*BBF Event Time Stamp	6.6.9	No	No	No	No	No	Yes

*Only applicable when dynamic NAT block assignment is required

Table 7: BBF Extended Information Element and Applicability Table for DBNG-UP centric NAT model

IE Type value (Decimal)	Information Elements	Section	Use Case: DHCPv4, PPPoEv4, LNS subscriber session CGNAT					
			CP UL PDR	CP DL PDR	UP UL PDR	UP DL PDR	SR Create SRR	Session Report
32768	BBF UP Function Features IE	6.6.1	No	No	No	No	No	No
32787	BBF Apply Action	6.6.5	No	No	Yes	Yes	No	No
32788	BBF NAT external Port Range	6.6.6	No	No	No	No	No	Yes
32789	BBF NAT Port Forward	6.6.7	No	No	No	No	No	No
32790	BBF Report Trigger	6.6.3	No	No	No	No	Yes	No
32791	BBF Dynamic NAT block port range	6.6.8	No	No	No	No	No	No
32792	BBF Event Time Stamp	6.6.9	No	No	No	No	No	Yes

6.6 PFCP Grouped IE extensions

6.6.1 BBF UP Function Features IE

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 12: 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. Spare bits must be set to zero by senders and shall be ignored by the receiver.

The following table specifies the features defined on the UP.

Table 8: BBF UP Function Features

Feature Octet / Bit	Feature	Description
7/7	NAT-CP	Informs the CP that the UP supports NAT and port forwarding function. Subscriber initial public IP assignment performed by CP.
7/8	NAT-UP	Informs the CP that the UP supports NAT and port forwarding function. Subscriber initial public IP assignment performed by UP.

Both NAT-CP and NAT-UP supports dynamic NAT block allocation.

6.6.1.1 Create Traffic Endpoint

The table below is the 3GPP defined Create Traffic Endpoint IE. Details of this grouped IE can be found in 3GPP TS 29.244 [2]. The grouped IE traffic endpoint already contains a list of IEs to specify properties of the endpoint including the GTP tunnel TEID or the subscriber IP address. To cover the wireline case, further extensions are required to describe the endpoint.

Table 9: 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
Reused 3GPP IEs below			
MAC address	C	If present, this IE MUST be used to identify the MAC address of the traffic endpoint. (see 3GPP TS 29.244 [2] 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 [2] 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 [2] for IE details)	S-Tag
BBF Extended IEs below			
Logical Port	C	If present, this IE MUST be used to provide an opaque value obtained from the NSH header to indicate the logical port for the subscriber. (see TR-459 section 6.6.3.1 [1] for IE details)	Logical port Details in TR-459 section 6.6.2 [1]
PPPoE Session ID	C	If present, this IE MUST be used to identify the PPPoE session ID of the subscriber. (see TR-459 section 6.6.5 [1] for IE details)	PPPoE Session ID Details in TR-459 section 6.6.5 [1]
L2TP tunnel	C	If present, this IE MUST be present if a L2TP tunnel is required. (see TR-459 section 6.5.8 [1] for IE details)	L2TP Tunnel Details in TR-459 section 6.5.8 [1]
BBF NAT External Port Range	C	This is only present with NAT-CP flag enabled from BBF UP Function Feature. If present, this IE Must be present if the subscriber requires NAPT	BBF NAT External Port Range
BBF Dynamic NAT Block Port Range	C	This is only present with NAT-CP flag enabled from BBF UP Function Feature. If present, this IE MUST be present if the subscriber requires dynamic NAT block allocation. This specifies the starting port and the ending port for the dynamic NAT block allocation.	BBF Dynamic NAT Block Port Range

6.6.1.2 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 10: 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 TR-459 section 6.6.3 [1].
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 TR-459 section 6.6.9 [1]
BBF NAT port forward	C	This is only present with NAT-CP flag enabled from BBF UP Function Feature. If present, this IE Must be present if the subscriber requires static NAT port forwarding	BBF NAT port forward

6.6.1.3 Update Forwarding Parameters

The Update forwarding Parameters IE in FAR can include the BBF Outer Header Creation IE and the MTU IE. The BBF Outer Header Creation IE is used to encapsulate the subscriber data packet in various wireline encapsulation. The MTU IE is primarily used in the case of PPPoE.

Table 11: 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
BBF NAT Port Forward	C	This is only present with NAT-CP flag enabled from BBF UP Function Feature. If present, this IE Must be present if the subscriber requires static NAT port forwarding	BBF NAT Port Forward

6.6.2 Create SRR IE

Table 12: Create SRR IE within PFCP Session Establishment Request

Octet 1 and 2	Create SRR IE Type = 212 (decimal)		
Octets 3 and 4	Length = n		
Information elements	P	Condition / Comment	IE Type
SRR ID	M	This IE shall uniquely identify the SRR among all the SRRs configured for this PFCP session.	SRR ID
Traffic Endpoint ID	C	This IE shall be present if the BBF Report Trigger indicates NAT. In this case the referred traffic endpoint ID contains the public NAT context for which reporting is enabled. In the case of NAT, this should be the same traffic endpoint the Outside IP address	Traffic Endpoint ID
BBF Report Trigger	C	This IE identify events to be reported	BBF Report trigger

6.6.3 BBF Report Trigger

Octets	Bits							
	8	7	6	5	4	3	2	1
1 to 2	Type = 32790							
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 11: BBF Report Trigger

The following enum are within Value (Octet 7):

- 0 - reserved
- 1 – Indicates DBNG-UP allocated NAT blocks need to be reported

6.6.4 Session Report IE within PFCP Session Report Request

Table 13: BBF extended Session Report IE

Octet 1 and 2	Create SRR IE Type = 214 (decimal)		
Octets 3 and 4	Length = n		
Information elements	P	Condition / Comment	IE Type
SRR ID	M	This IE shall uniquely identify the SRR among all the SRRs configured for this PFCP session.	SRR ID
BBF NAT Outside Address	C	This IE shall be present if the UPF support initial NAT block allocation and assigned a NAT outside address to the subscriber. This IE is only used for the initial NAT block assignment.	UE IP Address
BBF NAT External Port Range	C	This IE shall be present if the UPF is required to report NAT block allocation or deallocation	BBF NAT External Port Range
BBF Event Time Stamp	C	This IE shall be present to express the time of the NAT event	BBF Event Time Stamp

6.6.5 BBF Apply Action IE

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	Spare	Spare	Spare	Spare	Spare	Spare	Spare	NAT
8 to (n+4)	These octet(s) is/are present only if explicitly specified							

Figure 13: BBF Apply Action

This IE acts as an extension to the 3GPP Apply Action IE to define new actions. Currently following additional bits are defined:

- 7/1 – NAT – Apply NAT functionality, this requires the FORW action to be set. This is only valid if the UP indicated support for the NAT feature, and must only be used in combination with the two following FAR Destination Interfaces:
 - o Access: there is no requirement for the FAR to be linked to a traffic endpoint (unless outer header creation indicates this). The UP performs NAT functionality by looking up existing NAT mappings or port forwards. If such a lookup fails, the packet should be dropped.
 - o Network: the FAR must be linked to a traffic endpoint that includes a UE IP Address and NAT Port range. Packets will go through a NAT function that check if a mapping exists and forward accordingly. If no mapping exists a new mapping may be created based on the provided IP and port ranges.

6.6.6 BBF NAT External Port Range

	Bits							
Octets	8	7	6	5	4	3	2	1
1 to 2	Type = 32788							
3 to 4	Length = n							
5 to 6	Enterprise ID 3561							
7 to 8	Start Port							
9 to 10	End Port							
11 to (n+4)	These octet(s) is/are present only if explicitly specified							

Figure 14: BBF NAT External Port Range

This IE contains the starting and ending port assigned for the subscriber. There can be multiple starting and ending ports pair if multiple NAT blocks are assigned to the subscriber.

Start port – the starting IP port for the NAT block

End port – the ending IP port for the NAT block

6.6.7 BBF NAT Port Forward

	Bits							
Octets	8	7	6	5	4	3	2	1
1 to 2	Type = 32789							
3 to 4	Length = n							
5 to 6	Enterprise ID 3561							
7 to 10	Inside IP Address							
11 to 12	Inside Port							
13 to 14	Outside Port							
15	Protocol							
16 to (n+4)	Multiple of the quadruplets defined above							

Figure 15: BBF NAT port forward

If this attribute is not present, port forwarding is not required.

If this attribute has a length of 0, then delete all port forwarding entries for this.

If this attribute has a length other than 0, then replace existing port forwarding entries with this IE(s).

6.6.8 BBF Dynamic NAT Block Port Range

Octets	Bits							
	8	7	6	5	4	3	2	1
1 to 2	Type = 32791							
3 to 4	Length = n							
5 to 6	Enterprise ID 3561							
7 to 8	Start Port							
9 to 10	End Port							
11 to (n+4)	These octet(s) is/are present only if explicitly specified							

Figure 16: BBF Dynamic NAT Block Port Range

This IE indicates the starting and ending NAT port for dynamic NAT block assignment.
Start port – a reference to the traffic endpoint that is associated to the NAT outside endpoint.
End port – a reference to the traffic endpoint that is associated to the NAT outside endpoint.

6.6.9 BBF Event Time Stamp

The BBF Event Time Stamp IE indicates the time stamp when the event occurs in a given usage report. It shall be encoded as shown in Figure 8.2.114-1.

Octets	Bits							
	8	7	6	5	4	3	2	1
1 to 2	Type = 32792							
3 to 4	Length = n							
5 to 6	Enterprise ID 3561							
6 to 13	BBF Event Time Stamp							
14 to (n+4)	These octet(s) is/are present only if explicitly specified							

Figure 17: BBF Event Time Stamp

The BBF Event Time Stamp field shall contain a UTC time. Octets 6 to 13 shall be encoded in the same format as the eight octets of the 64-bit timestamp format as defined in clause 6 of IETF RFC 5905 [11].

Note: The encoding is defined as the time in seconds and fraction of a second relative to 00:00:00 on 1 January 1900.

Appendix I.

DBNG dynamic NAT block use case:

Network Address Translation (NAT) conserves IPv4 addresses by allows a single public IPv4 address to be divided amongst multiple subscribers into NAT blocks for Internet access. Below are some key factors in determining NAT blocks allocation.

- Port 0 to 1,023 are well known IANA registered ports which leaves ports only 1,024 to 65,535 available for NAT. As an example, if each subscriber is given 1000 ports with approximately 64000 ports available. A single IPv4 public address can then support 64 subscribers in total.
- The number of ports a subscriber require depends directly on the internet applications being used. Each application requires a variable amount of ports, some requiring only one while other requiring. The time of day also affects the number of ports consumed, where a typical household would utilize internet more in the evening. Nonetheless, to deliver the best end user experience, an adequate number of ports must be provided.
 - Allocation too small of a NAT block will help conserve public IPv4 address but provide an unpleasant end user experience as NAT ports run out. Allocation of a large NAT block alleviate the unpleasant subscriber experience but is wasteful on IPv4 public address consumption.
 - Typically, operators utilize a pre-trial phase to monitor typical ports consumption and determine the best NAT block size per household. This set block size would satisfy majority of the subscriber in the network.
 - In the atypical case, some subscribers such as enterprise customer may require larger NAT blocks. There are at least two possible solutions to address this:
 - Allocate a multiple NAT blocks at initialization.
 - Dynamically allocate additional NAT blocks when demand rises and deallocate NAT blocks when no longer in need.
- Local jurisdiction often requires operators to identify subscriber based on the IPv4 public address. Therefore, service providing must keep a persistent record of the allocated assigned IPv4 public address and the specific port range assigned to a particular subscriber at specify time of date. Logging can be quite intensive for both the DBNG-CP and the log servers:
 - The system requires processing cycles to gather information and generate a log entry which will contain the subscriber identification (ie. private IP, Mac, circuit ID), the NAT block ranges and its associate NAT IP address(es), and the current time stamp. This must be designed with disaster recovery scenario in mind where all subscribers will log into the system at once.
 - This log entry must be stored as non-volatile data and must account for write time such as write time on disk.
- To ease logging requirements a few standard practices is used:
 - Avoid over subscription of ports that requires constant NAT block allocation and deallocation. Allocate a block of ports that is adequate for the entire subscriber session. Reducing the entire subscriber into two logs: one for the NAT block consumed and one for releasing the NAT block no longer in use.
 - In the case where the subscriber MUST use dynamic NAT block range, the number of allocation/deallocation should be kept at a minimum.

Table 14: Log events

Subscriber type	NAT block dimensioning	Logging requirement
Typical broadband Subscribers	Single NAT block	2 log events
Premium broadband Subscribers	Single NAT block (medium size)	2 log events
Business Enterprise subscribers	Single NAT block (large size)	2 log events
Remaining subscribers	Variable NAT blocks	2*N log events (N = number of blocks allocated)

Use cases for prefix routing announcement and resiliency

Use Case 1: (Advertisement of Public IP prefix for subscriber routing)

To allow bi-directional routing of subscriber NAT'ed traffic to the A10 interface. The public prefix assigned by the DBNG-CP per DBNG-UP will require prefix announcements. In addition, in a resilient setup, the public prefix would require tagging to control the route advertisement.

Use Case 2: (Advertisement of Private IP prefix for specific services)

In some instances, it is required for the private IP prefix to be advertised to the rest of the network for IP services. These services include multicast and walled-garden services.

End of Broadband Forum Technical Report TR-459.2