ETSI TR 123 975 V19.0.0 (2025-10)

TECHNICAL REPORT

**Digital cellular telecommunications system (Phase 2+) (GSM);
Universal Mobile Telecommunications System (UMTS);
LTE;
IPv6 migration guidelines
(3GPP TR 23.975 version 19.0.0 Release 19)**

Reference
RTR/TSGS-0223975vj00

Keywords
GSM,LTE,UMTS

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - APE 7112B
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° w061004871

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The declarations pertaining to these essential IPRs, if any, are publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI IPR online database.

Pursuant to the ETSI Directives including the ETSI IPR Policy, no investigation regarding the essentiality of IPRs, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

DECT™, **PLUGTESTS™**, **UMTS™** and the ETSI logo are trademarks of ETSI registered for the benefit of its Members. **3GPP™**, **LTE™** and **5G™** logo are trademarks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners. **oneM2M™** logo is a trademark of ETSI registered for the benefit of its Members and of the oneM2M Partners. **GSM**® and the GSM logo are trademarks registered and owned by the GSM Association.

Legal Notice

This Technical Report (TR) has been produced by ETSI 3rd Generation Partnership Project (3GPP).

The present document may refer to technical specifications or reports using their 3GPP identities. These shall be interpreted as being references to the corresponding ETSI deliverables.

The cross reference between 3GPP and ETSI identities can be found at 3GPP to ETSI numbering cross-referencing.

Modal verbs terminology

In the present document "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the ETSI Drafting Rules (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

Contents

Foreword

This Technical Report has been produced by the 3rd Generation Partnership Project (3GPP).

The contents of the present document are subject to continuing work within the TSG and may change following formal TSG approval. Should the TSG modify the contents of the present document, it will be re-released by the TSG with an identifying change of Release date and an increase in version number as follows:

Version x.y.z

where:

x the first digit:

1 presented to TSG for information;

2 presented to TSG for approval;

3 or greater indicates TSG approved document under change control.

y the second digit is incremented for all changes of substance, i.e. technical enhancements, corrections, updates, etc.

z the third digit is incremented when editorial only changes have been incorporated in the document.

Introduction

With the depletion of IPv4 addresses and the development of data service, demands for deploying IPv6 are higher than before. This document analyzes different IPv6 migration scenarios and applicable mechanisms as well as identifies impacts on 3GPP network elements.

1 Scope

The technical report identifies various scenarios of transition to IPv6 and co-existence of IPv4 and IPv6, deployment options and impacts on 3GPP network elements. In particular:

- Identify the transition and co-existence scenarios of interest for operators and the respective assumptions and requirements.
- Analyze existing IP address allocation mechanism for IPv6 migration if necessary.
- Investigate IPv6 transition mechanisms for the scenarios identified during the study and investigate their applicability for 3GPP network, and identify the compatibility among applicable transition mechanisms.
- Identify any impact on 3GPP network elements.
- Provide recommendations on IPv6 transition and co-existence of IPv4 and IPv6 and identify if any normative work is needed.

2 References

The following documents contain provisions which, through reference in this text, constitute provisions of the present document.

- References are either specific (identified by date of publication, edition number, version number, etc.) or non-specific.
- For a specific reference, subsequent revisions do not apply.
- For a non-specific reference, the latest version applies. In the case of a reference to a 3GPP document (including a GSM document), a non-specific reference implicitly refers to the latest version of that document *in the same Release as the present document.*

[1] 3GPP TR 21.905: "Vocabulary for 3GPP Specifications".

[2] IETF RFC 6333: "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion".

[3] IETF RFC 6346: "The Address plus Port (A+P) Approach to the IPv4 Address Shortage".

[4] http://tools.ietf.org/html/draft-boucadair-port-range-02.

[5] http://tools.ietf.org/html/draft-bajko-pripaddrassign-01.

[6] IETF RFC 6144: "Framework for IPv4/IPv6 Translation".

[7] IETF RFC 6535: "Dual-Stack Hosts Using "Bump-in-the-Host" (BIH)".

[8] IETF RFC 6145: "IP/ICMP Translation Algorithm".

[9] 3GPP TS 23.401: "General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access".

[10] 3GPP TS 23.402: "Architecture enhancements for non-3GPP accesses".

[11] 3GPP TS 23.060: "General Packet Radio Service (GPRS) Service description; Stage 2".

[12] Durand, A., Droms, R., Haberman, B., Woodyatt, J., Lee, Y., and R. Bush: "Dual-stack Lite broadband deployments post IPv4 exhaustion", IETF draft, draft-ietf-softwire-dual-stack-lite-07 (work in progress).

[13] Brockners, F., Gundavelli, S., Speicher, S., Ward, D: "Gateway Initiated Dual-Stack Lite Deployment", draft-ietf-softwire-gateway-init-ds-lite-03 (work in progress).

[14] 3GPP TR 23.981: "Interworking aspects and migration scenarios for IPv4 based IMS implementations".

[15] Void.

[16] Void.

[17] IETF RFC 6146: "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers".

[18] IETF RFC 6147: "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers".

[19] IETF RFC 6619: "Scalable Operation of Address Translators with Per-Interface Bindings".

[20] IETF RFC 4364: "BGP/MPLS IP Virtual Private Networks (VPNs)".

[21] 3GPP TS 23.203: "Policy and charging control architecture".

[22] Void.

[23] IETF Internet-Draft, draft-boucadair-behave-bittorrent-portrange-02: "Behaviour of Bit Torrent service in an IP Shared Address Environment" work in progress.

[24] Haverinen, H., Siren, J., and P. Eronen: "Energy Consumption of Always-On Applications in WCDMA Networks", VTC'07-Spring, Dublin Ireland, 20-25 April 2007.

[25] 3GPP TS 24.008: "Mobile radio interface Layer 3 specification; Core network protocols; Stage 3".

[26] IETF RFC 1918: "Address Allocation for Private Internets".

[27] IETF RFC 2784: "Generic Routing Encapsulation (GRE)".

[28] IETF RFC 2890: "Key and Sequence Number Extensions to GRE".

[29] IETF RFC 3338: "Dual Stack Hosts Using "Bump-in-the-API" (BIA)".

[30] IETF RFC 2767: "Dual Stack Hosts using the "Bump-In-the-Stack" Technique (BIS)".

3 Definitions, symbols and abbreviations

3.1 Definitions

For the purposes of the present document, the terms and definitions given in TR 21.905 [1] and the following apply. A term defined in the present document takes precedence over the definition of the same term, if any, in TR 21.905 [1].

NAT: A function which provides NAT44, NAPT44, NAT64, NAPT64 or combinations of these.

Attachment circuit: as used by RFC 4364 [20], term to refer generally to means of attaching to a router, such as: PPP connections, ATM Virtual Circuits (VCs), Frame Relay VCs, Ethernet interfaces, Virtual Local Area Networks (VLANs) on Ethernet interfaces, GRE tunnels, Layer 2 Tunnelling Protocol (L2TP) tunnels, IPSec tunnels, etc. An attachment circuit identifies uniquely the MPLS VPN used by all traffic using that circuit.

CE: as used by RFC 4364 [20], stands for Customer Edge router or Customer Edge device. It represents an IP device using a BGP/MPLS IP Virtual Private Networks (VPN) to communicate with other CE devices using the same VPN, without the need to be routing peers of each other and without visibility of MPLS or the MPLS backbone providing connectivity between CEs in different sites. CEs are connected to PEs using attachment circuits. If CEs use dynamic routing protocols (CE routers) to route traffic in the VPN, then they are routing peers of the directly attached PEs.

PE: as used by RFC 4364 [20], stands for Provider Edge router. PEs use MPLS to tunnel traffic among each other enabling IP traffic between the CEs attached to them.

3.2 Abbreviations

For the purposes of the present document, the abbreviations given in TR 21.905 [1] and the following apply. An abbreviation defined in the present document takes precedence over the definition of the same abbreviation, if any, in TR 21.905 [1].

NAPT44	Network Address and Port Translation IPv4 to IPv4
NAT44	Network Address Translation IPv4 to IPv4
NAPT64	Network Address and Port Translation IPv6 to IPv4
NAT64	Network Address Translation IPv6 to IPv4
PCC	Policy and Charging Control

4 Baseline Architecture for IPv4 and IPv6 Co-existence

This clause describes how dual-stack connectivity has been specified for the EPS and GPRS networks.

The Release 8 3GPP EPS architecture supports and optimises the co-existence of IPv4 and IPv6 with dual-stack operation. Dual-stack operation means that native IPv4 and native IPv6 packets are transported in parallel by tunnelling them from the UE to the PDN GW within a single EPS bearer/PDP context. This dual-stack EPS bearer/PDP context is associated with both an IPv4 address and an IPv6 prefix.

In comparison, dual stack connectivity to a given PDN in the pre-Release 9 GPRS network (with Gn/Gp SGSN and/or GGSN elements) requires the activation of two parallel PDP contexts, one for IPv4 traffic and one for IPv6 traffic. It should be noted that these parallel PDP contexts enable the same dual-stack connectivity for an application as the dual-stack EPS Bearers/PDP Contexts in the Release 8 EPS.

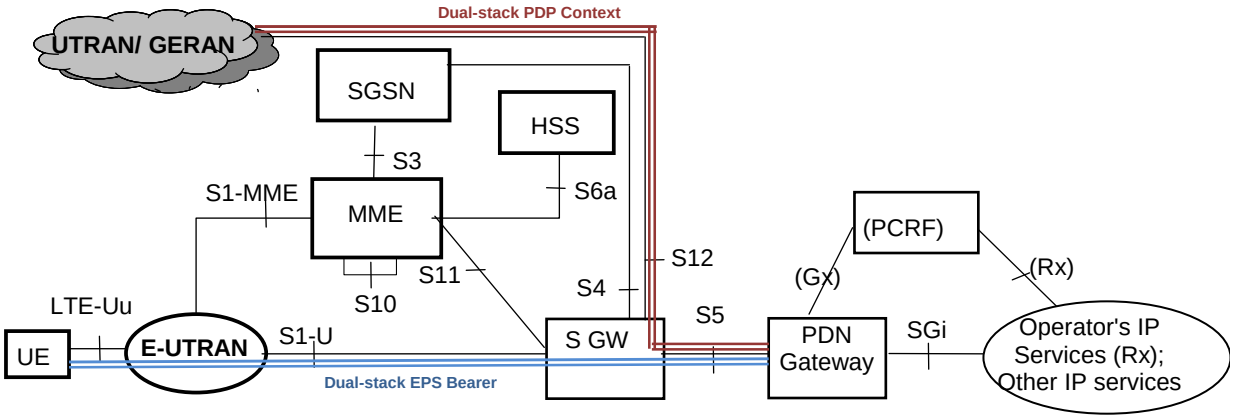

Figure 4.1: EPS Non-roaming architecture for 3GPP accesses in Release 8

Figure 4.1 depicts the Release 8 3GPP reference architecture for EPS according to TS 23.401 [9]. Upon request from the UE, the MME and S4-SGSN can activate a dual-stack EPS bearer/PDP context, which is identified in signalling by the PDN/PDP type 'v4v6'. A dual-stack EPS bearer tunnels IPv4 and IPv6 traffic in parallel from the UE to PDN GW.

In order to support dual-stack connectivity where possible, it has been specified in the Release 8 EPS specifications TS 23.401 [9], TS 23.060 [11], that if a Release 8 UE/MS supports both IPv4 and IPv6, the UE/MS shall always start off by requesting for a dual-stack (PDN/PDP type v4v6) bearer. It is also assumed that the UE/MS has no knowledge of the IPv4 and/or IPv6 capabilities of a given PDN. Neither does the UE/MS have any awareness of whether dual-stack bearers/ contexts are supported by the network to which it is attaching.

In Release 8, the EPS control plane elements (MME, S4-SGSN) and user plane elements (SGW, PGW) are all able to identify and handle requests to activate a dual-stack bearer/ context, and to enforce the type of bearers/ contexts that are allocated to the UE/MS. The network may downgrade the request for the PDN/PDP type v4v6 if a given PDN supports/allows only one of the address types (i.e. IPv4 or IPv6) as configured in the HSS. This limitation may stem from an operator policy. Another reason for downgrading may be that there are Gn/Gp SGSNs in the operator's network that have not been upgraded to support the PDP type 'v4v6'. The outcome of a PDN/PDP Type request depends on HSS provided subscription data, PGW configuration and home (and possibly visited) network core configuration. Any of these factors may downgrade the request to a single address type.

The EPS interworking architecture for Gn/Gp SGSNs is shown in Figure 4.2.

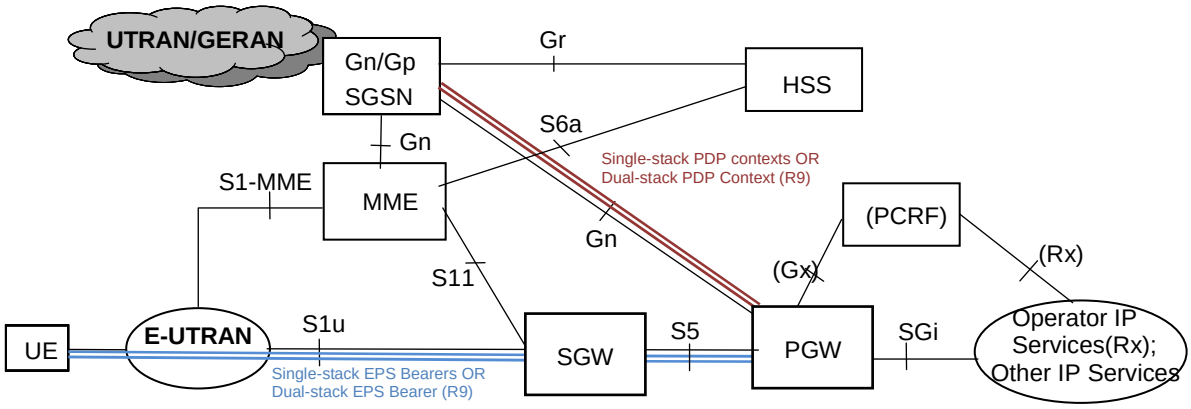

Figure 4.2: EPS Non-roaming Architecture for interoperation with Gn/Gp SGSNs in Rel-8

Dual stack PDP context support for the GPRS core network (GGSN, Gn/Gp SGSN) is specified in Release 9. To use this functionality a MS need to support Release 8 or higher (same as S4-SGSN access in EPS) in order to successfully request a PDP type 'v4v6' connection in UTRAN/GERAN. A pre-Release 9 Gn/Gp SGSN handles PDP type 'v4v6' as an 'unknown' PDP type, meaning that it handles a request for the PDP type 'v4v6' as if it were a request for the PDP type 'v4'.

NOTE 1: The 3GPP specification TS 24.008 [25] is not entirely unambiguous on the treatment of unknown PDP types. Even if the information element coding for "PDP type" specifies that a request for an "unknown PDP type" shall be treated as if it were a request for PDP type v4, the error signalling elsewhere in the specification include the possibility to signal an error code "unknown PDP address or PDP type".

In order to support inter-RAT mobility to/from a pre-Release 9 Gn/Gp SGSN, parallel v4 and v6 bearers/PDP contexts to a given PDN must to be used instead of dual-stack contexts.

The request to activate two parallel single stack bearers/PDP contexts is always initiated by the UE/MS. If the Release 8 network assigns a single-stack bearer to the UE/MS in response to a request for a dual-stack bearer, the network also signals to the UE/MS an indication on whether parallel single stack bearers are allowed to the same PDN or not. If the UE/MS fails to activate a dual-stack bearer/context, and it receives a single-stack IPv4 or IPv6 bearer/context, it shall attempt to activate a parallel single-stack bearer/context for the other IP address type to the same PDN, unless the UE has received an explicit indication from the network that parallel single stack bearers/contexts are not allowed.

In GPRS core networks, dual-stack connectivity is also possible with a pre-Release 9 GGSN and SGSN. These network elements do not support dual-stack PDP contexts, but dual-stack usage may be possible by activating a parallel IPv4 PDP context and IPv6 PDP context to the same PDN. In order to establish dual-stack connectivity in this case, a dual-stack UE shall attempt to open parallel single-stack v4/v6 PDP contexts to the same PDN even without receiving an explicit indication on support for parallel single stack bearers to the same PDN.

For end-hosts, the activation and mobility of dual-stack bearers/ contexts is simpler in comparison to handling of parallel IPv4 and IPv6 bearers/contexts. The usage of dual stack bearers/ contexts also simplifies the handling of parallel IPv4 and IPv6 traffic within the network after early EPS deployment phase when SGSNs are upgraded to support the PDP type 'v4v6'.

TR 23.981 [14] describes a scenario where old SGSNs do not support PDP type IPv6. Considering the fact that PDP type IPv6 has been specified since R'97 and SGSNs shipped during the last couple of years have support for PDP type IPv6, the assumption is that all SGSNs support PDP type IPv6.

5 IPv6 migration scenarios

5.1 Scenario 1: Dual-stack connectivity with Limited Public IPv4 Address Pools

In this IPv6 transition scenario, the operator runs the user plane in dual stack mode, i.e., the UEs are assigned both an IPv6 prefix and an IPv4 address in order to allow UEs to utilise both IPv4 and IPv6 capable applications. This scenario relies on the availability of dual-stack UEs, which are able to support parallel IPv4 and IPv6 connectivity to a single PDN. It is further assumed that the proportion of IPv6 capable applications will start to increase as soon as UEs and networks starts to become dual-stack capable. As popular services start to support IPv6, a part of IPv4 traffic will gradually be offloaded into the IPv6 domain. Services that are operator owned and deployed (for example LTE voice and other IMS based services) could be IPv6 enabled (in addition to IPv4) and hence accessible by the dual-stack capable UEs.

During transition phase, the depletion of public IPv4 addresses may become an issue in some operators' networks. The lack of public IPv4 address availability in the near future will inhibit the growth of data services and mobile broadband networks. The shortage of public addresses will be aggravated by always-on packet data connectivity, which is expected to prevail in newer network deployments.

To alleviate the shortage of public IPv4 addresses, the usage of private IPv4 addresses can be considered (e.g. the RFC 1918 [26] addresses). The utilisation of private IPv4 addresses should not require new procedures to be specified for the UE in order to ensure maximum applicability in a network with an early dual-stack UE population.

5.2 Scenario 2: Dual Stack connectivity with Limited Private IPv4 Address Pools

This migration scenario is based on the Dual stack model: The operator assigns both an IPv6 prefix and an IPv4 address to UEs in order to ensure that both IPv4 and IPv6 capable applications can be utilised.

The IPv4 addresses assigned to UEs are taken from one of the private address ranges as specified in RFC 1918 [26]. To enable global connectivity, network address translation (NAT) is performed on the (S)Gi interface for IPv4 packets originated from or destined to the UEs.

NOTE: Depletion of public IPv4 addresses while transitioning to IPv6 might be one reason for operators to assign private as opposed to public IPv4 addresses to UEs.

The challenge of this scenario lies in the limited number of private IP addresses. In case more than 16 million UEs are active (i.e. have an active PDP context/EPS bearer) in the same network at the same time, the network will run out of private IPv4 addresses. In order to avoid this, the operator may have to consider assigning the same IPv4 address to multiple UEs. Nevertheless, the operator expects that legacy IPv4 applications continue to work in this situation.

When defining solutions for this scenario, it additionally needs to be taken into account that in existing deployments some operators currently use the private IPv4 address assigned to a given UE to identify the respective customer (note that for this reason private IPv4 addresses are currently unique within these networks). Therefore, a solution for this scenario needs to ensure that IPv4 flows on the Gi interface can be uniquely traced back to a given UE/customer.

5.3 Scenario 3: UEs with IPv6-only connection and applications using IPv6

The operator decides to only assign IPv6 prefixes to the UEs due to e.g. shortage of IPv4 addresses or to address use cases, in which it appears beneficial - from an operational perspective - to only assign IPv6 addresses (e.g. m2m scenarios). UEs with IPv6-only connectivity running applications using IPv6 should however still be able to access both IPv4- or IPv6-enabled services.

Based on this scenario description, two use cases need to be considered:

1) The UE, configured only with an IPv6 prefix, has to be able to access IPv4 services.

2) The UE, configured only with an IPv6 prefix, has to be able to access IPv6 services.

5.4 Scenario 4: IPv4 applications running on a Dual-stack host with an assigned IPv6 prefix and a shared IPv4 address and having to access IPv4 services

In this scenario an IPv4 application running on a dual-stack UE requires to access IPv4 services without the operator having to allocate a unique non-shared (private or public) IPv4 address to the UE. The dual-stack UE running these applications uses an IPv4 address that is shared amongst many other UEs, and uses an IPv6 prefix.

6 High level requirements

The high-level requirements are to cover all the scenarios described in clause 5 in roaming and non-roaming cases. The IPv6 migration architecture should take into consideration any possible impacts to the policy architecture.

7 Solutions and functional flows description

7.1 Solution 1 - Dual-stack deployment combined with NAPT44

7.1.1 Overview

Enabling MS/UE dual stack communication and moving traffic from IPv4 to IPv6 can achieve a significant reduction of the number of dedicated public IPv4 addresses assigned to a NAPT44. Public IPv4 are assigned to the NAPT44 for the purpose of sharing this resource among several users at a specific ratio determined by the number of ports dynamically available to a single user as calculated by the operator. Real port usage is determined by the user applications and their need for connections which could range from one to several thousands per user. Offloading these IPv4 resources by moving the traffic to IPv6 will end up freeing a significant amount of public IPv4 addresses that can be used elsewhere in the operator network.

The use of traditional NAT has a size limitation due to the maximum 16 million available RFC 1918 [26] private IPv4 addresses. The description below depicts possible solutions to how the operational impact of this limitation can be overcome.

7.1.2 Description

7.1.2.1 Dual-stack deployment

MS/UE attaches to network APN(s) using applicable procedures described in TS 23.401 [9], TS 23.402 [10] and TS 23.060 [11] in order to get dual stack connectivity to Internet (IPv4 and IPv6). The operator assigns private IPv4 addresses to the UEs and uses NAPT44 to provide access to the Internet. The operator may multiplex multiple UEs onto a single public IPv4 address using traditional NAPTs. The operator assigns IPv6 prefixes to the UEs allowing native IPv6 access to the Internet.

The MS/UE will now use IPv6 to communicate with dual stack reachable services/peers and thus offloading the NAPT44 assigned public IP address/ports resources that would have been made available for the UE if it not had been able to use IPv6. When communicating with Services/peers only served by IPv4, the UE/MS will use NAPT44 resources to enable communication. During the co-existence phase of the IPv6 migration, more IPv4 traffic will be offloaded from the NAPT44 as more and more services/peers become dual stack reachable or complete the transition and become IPv6 only reachable.

7.1.2.2 NAPT44 deployment options

7.1.2.2.1 Basic deployments

Typically, a single physical PDN-GW can serve an order of few million UEs in maximum. As the amount of traffic per user increases it is not expected that there will be a major increase in this number. Therefore, we can expect a single PGW can hardly ever reach the point where a PDN-GW would need to hand out more than 16 million RFC 1918 [26] addresses. It looks evident for time being that a single or even a small cluster of PGWs implementing collocated NAPT44 functions should not run out of RFC 1918 [26] addresses for one APN. Figure 7.1.1 illustrates a deployment discussed here.

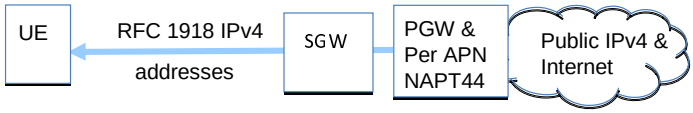

Figure 7.1.1: NAPT44 collocated in a PDN-GW for each APN

In a case multiple PDN-GWs serve a single RFC 1918 [26] addressed PDN identified by a single APN, the RFC 1918 [26] address space must be partitioned so that overlapping does not happen between PDN-GWs serving the PDN. This is a pure address management issue and illustrated in Figure 7.1.2.

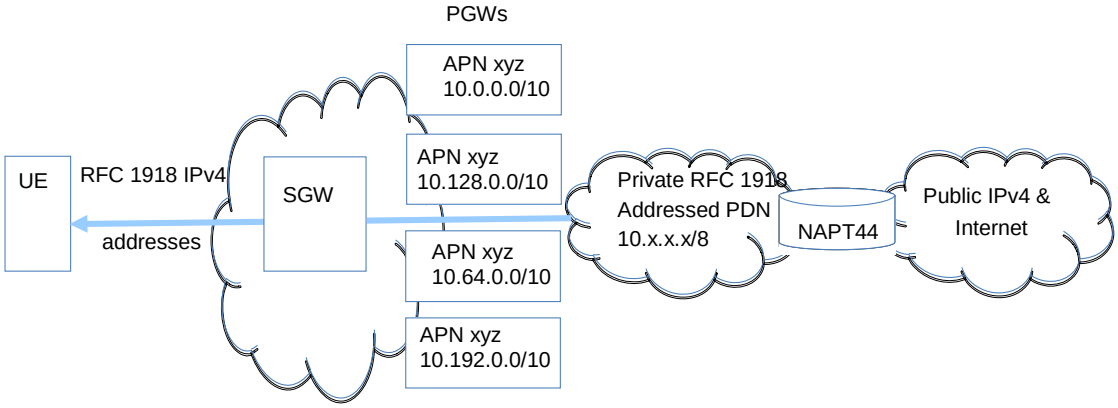

Figure 7.1.2: Private RFC 1918 [26] addressed PDN with multiple PDN-GWs and non-overlapping address spaces

7.1.2.2.2 Deployments with overlapping RFC 1918 [26] address spaces

However, it is also possible that multiple PDN-GWs serving the same APN would go beyond 16 million RFC 1918 [26] addresses. In this situation the APN has to be partitioned into independent PDNs with overlapping RFC 1918 [26] address spaces. This is a pure network deployment issue. In this deployment model the NAPT44 functionality can be located either in a PDN-GW or at the edge of the RFC 1918 [26] addressed PDN and the public Internet. The model where NAPT44 functionality is collocated in a PDN-GW is illustrated in Figure 7.1.3.

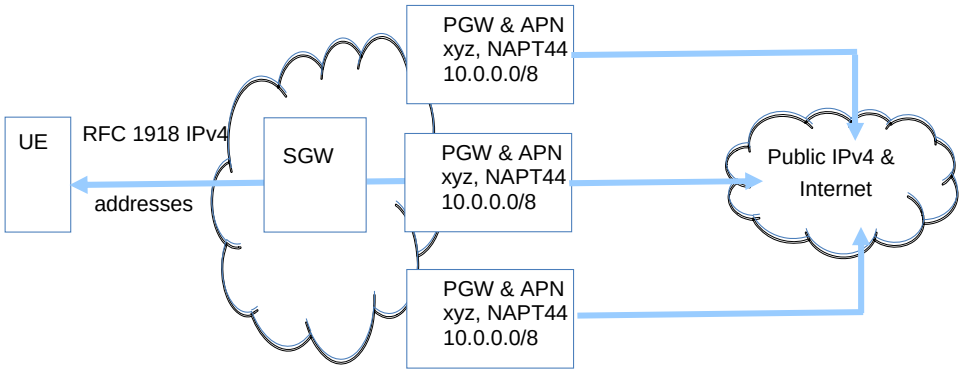

Figure 7.1.3: Overlapping RFC 1918 [26] address spaces for the same APN and NAPT44 collocated with a PDN-GW

The deployment model where the APN has been partitioned multiple independent PDNs is illustrated in Figure 4. Here the NAPT44 functionality is distributed between each independent & private RFC 1918 [26] PDN and the public Internet.

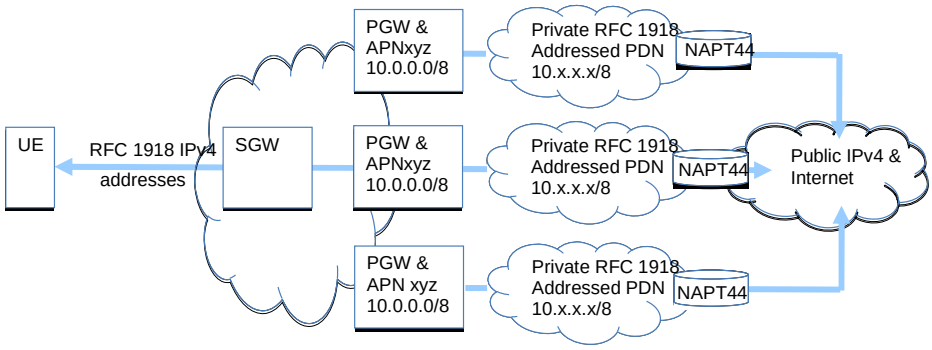

Figure 7.1.4: Overlapping RFC 1918 [26] address spaces for the same APN and distributed NAPT44

7.1.2.2.3 Identity considerations when using overlapping RFC 1918 [26] address spaces

If a PLMN needs to provide more than 16 million RFC 1918 [26] addresses to its own subscribers, and wants to correlate a private RFC 1918 [26] address to a specific UE, then additional functionality would be needed. This deployment model is actually the same what was already illustrated in Figure 7.1.3 and 7.1.4. If the service infrastructure needs to distinguish between subscribers with overlapping RFC 1918 [26] addresses but still only compare IP addresses, then the comparison has to include additional information or be context aware of the source of the used RFC 1918 [26] pool.

For example, beyond knowledge of the NAT binding state to derive the private IPv4 address, the comparison could also take the public IP address of the GGSN/PGW into account where the NAPT44 takes place. If there need to be traffic inspecting devices within each PDN of Figure 4 with overlapping RFC 1918 [26] addresses then the IP address of the traffic inspecting device could be used to identify the PDN where the RFC 1918 [26] pool belongs to.

Other solutions are also possible. If needed, a unique identity may be tied to the user packets on the outer/external/public address side of the NAT. This can be arranged, as an example, by using packet encapsulation where a unique identifier is included either in the packet encapsulation information or in the source address of the encapsulating packet.

7.1.3 Functional Description

The MS/UE need to obtain dual stack connectivity in order to be able to reach both IPv4 and IPv6 services/peers. This can be arranged either by using a dual stack connection by requesting a connection of PDP Type IPv4v6 or PDN Type IPv4v6 depending on radio access technology and MS/UE capability. If these dual stack are not possible to obtain it is also possible to request two separate connections, one PDP context/PDN connection Type IPv4 and one PDP context/PDN connection Type IPv6 in parallel to the same dual-stack APN. The preferred way would be to use only one

connection for both IP versions but the two connection approach could be used due when ether MS/UE or core does not allow for a single dual stack PDP context connection to be established.

The following table lists the basic requirements for this scenario in an IP version co-existence phase referencing the user plane capabilities only.

Table 7.1.1: IPv4 offload requirements

Basic Components Name	States	PDP/PDN Types
Terminal IP capability	Dual stack (IPv6 preferred over IPv4 if both can be used for a remote endpoint)	IPv4v6, IPv4 and IPv6 (NOTE 1)
Type of application program	Dual stack capable	not applicable
Type of assigned IP address,	IPv4 and IPv6	not applicable
Subscriber IP capability	Dual stack APN in subscriber data	IPv4v6, IPv4 and IPv6
Network IP capability	Dual stack network	IPv4v6, IPv4 and IPv6 (NOTE 2)
Service/peer capability	Dual stack (NOTE 3)	not applicable

The GGSN/PDN GW IPv4 Internet connectivity is provided over a NAPT44 solution either co-located with the GGSN/PDN GW or elsewhere placed in the operator network.

NOTE 1: To be able to use PDP/PDN Type IPv4v6 the MS/UE need to be Release 8 or later.

NOTE 2: To be able to serve PDP/PDN Type IPv4v6 the core nodes need to be Release 8 or later except for SGSN/GGSN using Gn/Gp need to be Release 9.

NOTE 3: If DNS is to be used to resolve the service/peer FQDN into an IP address the node DNS information need to contain both A and AAAA record entries for the service/peer.

7.1.4 Information flows

See TS 23.060 [11], TS 23.401 [9] and TS 23.402 [10] for the appropriate information flow details.

7.1.5 Evaluation

The solution assumes that Internet content/services start becoming dual-stack capable and thus available via IPv6. The presented NAPT44 based solution and few other considerations are sufficient to support the migration period, more specifically to address the problems of limited private IPv4 address space. The 3GPP community should consider influencing major Internet content/service providers to make their services available via IPv6 in a user friendly manner. Offloading some traffic to IPv6 reduces the amount of active connections required in the NAPT44. This reduces the scalability issues with NAT and the number of public IPv4 addresses/ports needed to serve the UEs.

Known issues of the solution:

- The session binding between private and public IPv4 address/port is not known to the PCC architecture. Therefore, depending on deployment and if the application is NAT aware and has access to the binding (as e.g. in the case of IMS), there may or may not be issues with applying PCC to the session.

- General NAT concerns, not specific to 3GPP networks, apply. For example, applications that embed IP addresses in the payload and are not NAT aware require additional functionality to work across NATs.

Known benefits of the solution

- This solution requires no changes to UEs, it can be used with legacy dual-stack UEs.

- This solution has no impact to the 3GPP network architecture, no new interface or network element is needed. It can be deployed without any additional normative specification within 3GPP. The limitations described under "known issues" above apply.

- This solution does not introduce any additional tunnelling overhead on any interfaces.

- This solution with the appropriate deployment supports UEs with overlapping address space, thus there is no limitation of the number of subscribers.

- Support for UEs with public, private, and overlapping private IPv4 addresses. If so desired, all the UE's in the mobility domain can be assigned the same IPv4 private address.
- No changes to the IPv4 / IPv6 address-assignment procedures required.
- No bearing on the type of transport network: Transport network can be IPv4 or IPv6.
- NAPT44 can be either co-located or separate from GGSN/PDN-GW.
- Solution to the public IPv4 address exhaustion problem through the use of NAPT44.
- Solution to the private IPv4 address exhaustion problem through the use of overlapping private IPv4 addresses.
- This solution does not have any impact on the UE's roaming support.
- No impact on QoS/bearer procedures between UE and PDN GW/SGW/GGSN.

7.1.6 Applicability

This solution applies to scenario 1.

This approach also suggests solutions to address scenario 2.

Given the solution description above, the described functionality can be configured in currently deployed mobile networks as well as in future deployments regardless of 3GPP access technology. When to deploy such a setup in an operator's network is more of a business and operational decision.

7.2 Transition Solution: Gateway-Initiated Dual-Stack Lite

7.2.1 GI-DS-lite Overview

Gateway-Initiated Dual-Stack Lite [13] (GI-DS-lite) is a modified approach of the DS-Lite concept. The GI-DS-lite concept applies to EPC as well as GPRS. For reasons of simplicity, this clause uses EPC nomenclature. GPRS applies in a similar way.

GI-DS-Lite builds on top of the current dual-stack deployment model of the 3GPP architecture which supports dual-stack UEs and uses tunnelling technology between the Serving Gateway and the PDN Gateway, over GTP or PMIPv6 based S5/S8 interfaces, and between the UE and the PDN Gateway over the S2c interface. GI-DS-Lite lifts some of the restrictions of the DS-lite solution:

- Carrier Grade NAPT (CGN) does not need to be co-resident with PDN-Gateway.
- No added overhead for IPv4 user plane traffic transport on the airlink.
- Support of IPv4 and IPv6 transport networks.
- Support for deployments with public, private, and overlapping IPv4 addresses on the UEs.
- No UE changes mandated for any of the deployment scenarios.

With GI-DS-Lite, UE and access architecture remain unchanged. PDN Gateway and CGN are connected through a "softwire" identified by a Softwire ID (SWID). The SWID does not need to be globally unique, i.e. different SWIDs could be used to identify a softwire at the different ends of a tunnel. A Context-Identifier (CID) is used to multiplex flows associated with the UE onto the softwire tunnel. Local policies at the PDN Gateway determine which part of the traffic received from an UE is sent via a softwire to the CGN. The combination of CID and SWID serves as common context between PDN Gateway and CGN to identify flows associated with an UE. The CID is typically a 32-bit wide identifier assigned by the gateway. It is retrieved either from a local or remote (e.g. AAA) repository. The CID ensures a unique identification (potentially along with other traffic identifiers such as e.g. interface, VLAN, port, etc.) of traffic flows at the Gateway and CGN. The embodiment of the CID and SWID depends on the technology used and the type of the network connecting PDN Gateway and CGN.

If, for example GRE (RFC 2784 [27]) with "GRE Key and Sequence Number Extensions" (RFC 2890 [28]) is used as softwire technology, the network connecting PDN Gateway and CGN could be either IPv4-only, IPv6-only, or a dual-stack IP network. The GRE-key field represents the CID.

In case of MPLS VPN, RFC 4364 [20] used between PDN Gateway and CGN, the softwire identification is supplied by the VPN identifier of the MPLS VPN, whereas the IPv4 address serves as CID. Depending on if the PDN Gateway and CGN are connected as CEs or PEs, the VPN identifier could e.g. be an attachment circuit identifier (e.g.. a VLAN tag), a representation of the VPN route labels pointing to routes within the VPN, the VPN route distingisher etc. The combination of CID and SWID ensures a unique identification (potentially along with other traffic identifiers such as e.g. interface, VLAN, port, etc.) for traffic flows at the CGN, which should be associated with a single NAT-binding. Deployment dependent, the CID can also be used as an identifier for traffic flows or UEs in backend systems: Deployments which use non-overlapping private IPv4 addresses for the UE could e.g. choose to map private IPv4 addresses 1:1 to the CID.

In a GI-DS-Lite deployment, the CGN combines DS-Lite softwire termination and NAT44. The outer/external IPv4 address of a NAT-binding at the CGN is either assigned autonomously by the CGN from a local address pool, configured on a per-binding basis (either by a remote control entity through a NAT control protocol or through manual configuration), or derived from the CID (e.g. the 32-bit CID could be mapped 1:1 to an external IPv4-address). The choice of the appropriate translation scheme for a traffic flow can take parameters such as destination IP-address, incoming interface, etc. into account. The IP-address of the CGN, which, depending on the transport network between the PDN Gateway and the CGN, will either be an IPv6 or an IPv4 address, is configured on the gateway. A variety of methods, such as out-of-band mechanisms, or manual configuration apply.

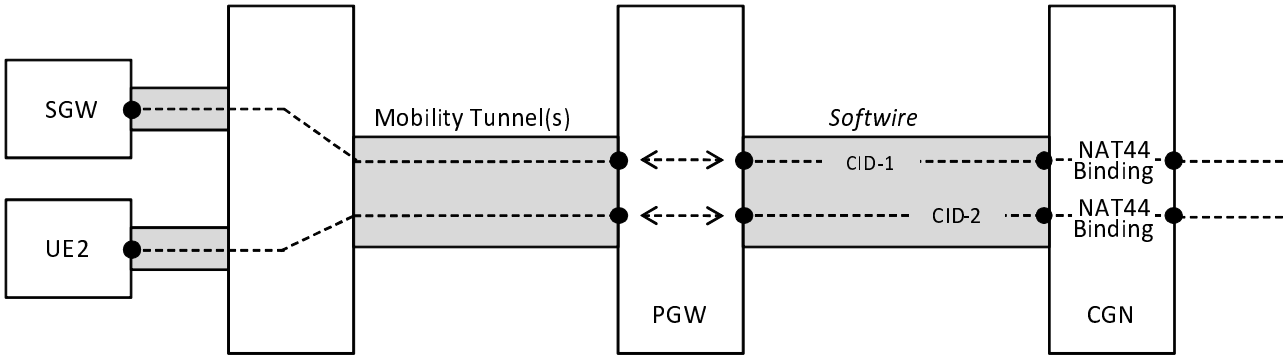

Figure 7.2.2a: Gateway-Initiated Dual-Stack Lite deployment scenario

Figure 7.2.2a shows an example of Gateway-Initiated DS-Lite applied to the EPC architecture when S5 or S8 interfaces are used. The PDN Gateway associates the mobility tunnels with the DS-Lite softwire to facilitate traffic forwarding to and from the CGN.

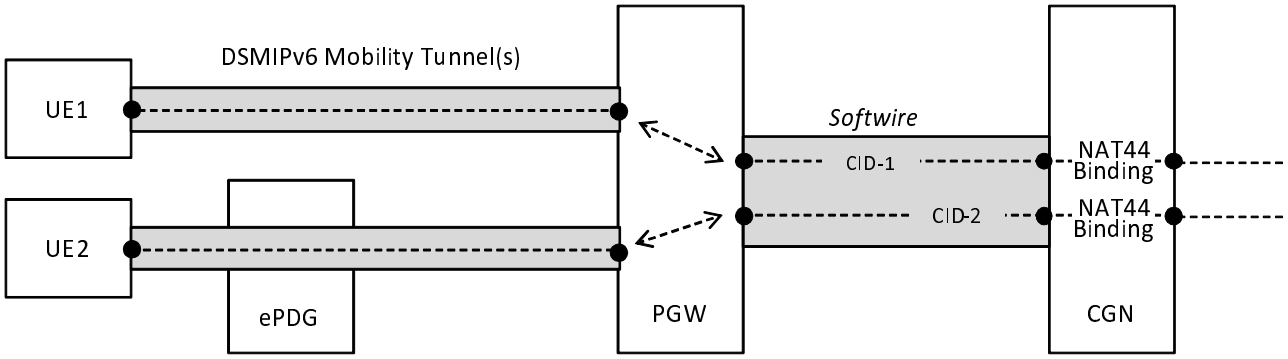

Figure 7.2.2b: Gateway-Initiated Dual-Stack Lite deployment scenario over S2c

Figure 7.2.2b shows an example of Gateway-Initiated DS-lite applied to the EPC architecture when the S2c interface is used. The PDN Gateway associates the mobility tunnels with the softwire to facilitate traffic forwarding to and from the CGN.

In its simplest form, there could be a 1:1 relationship between mobile access tunnels (e.g. identified by a TEID or the DSMIPv6 HNP) and a combination of CID and SWID identifying the softwire facing the CGN - resulting in a simple

tunnel-stitching operation on the PDN Gateway. Deployment dependent (e.g. for deployments which use non-overlapping private IP addresses on the UEs), the PDN Gateway could e.g. choose to only send Internet-bound traffic to the CGN - and route internal traffic locally.

7.2.2 GI-DS-lite Evaluation

Impact on the existing architecture:

The following capabilities are used to support GI-DS-lite:

- Softwire tunnelling on SGi, between the PDN Gateway and CGN, for instance:
  - GRE w/ GRE-key extensions (or alternative schemes, such as MPLS) tunnelling to/from the Carrier Grade NAT.
  - MPLS VPNs using attachment circuits according to RFC 4364 [20] between the PE devices and PDN Gateway, and between the PE devices and CGN.
  - MPLS VPNs using MPLS between PDN Gateway and CGN.

The following capabilities are used to support GI-DS-lite using GRE:

- Procedures for the PDN Gateway to support UE with overlapping IPv4 addresses
- A tunnel with the appropriate encapsulation mode needs to be setup between the PDN Gateway and the CGN. It is established at the system startup time and is enabled based on the configuration.
- PDN GW may assign overlapping private IPv4 addresses to all the UE's within that operational domain.
- when overlapping IPv4 address assignment is supported and used in the softwire tunnel, the PDN GW shall associate the UE session with a CID. This identifier will be unique to the UE's PDN connection.
- the PDN GW shall tunnel the IPv4 UE traffic using the appropriate encapsulation scheme on SGi to the CGN. It will use the CID associated with the UE's session.
- CID management on the PDN Gateway
- Maintenance of a CID key-space (possibly in conjunction with an external repository (e.g. AAA)).
- PCC enhancements (to cover cases where the GRE-key would need to be used to identify IP-CAN sessions, which would be the case for deployments which use non-unique IP-addresses within the mobile domain and use the IP-address as IP-CAN session identifier).

The following capabilities are used to support GI-DS-lite using MPLS VPNs:

- PDN Gateway support for:
  - if deployed as CE, at least one type of attachment circuit according to RFC 4364 [20]
  - if deployed as combined CE and PE, MPLS VPNs according to RFC 4364 [20]
- CGN support for:
  - if deployed as CE, at least one type of attachment circuit according to RFC 4364 [20]
  - if deployed as combined CE and PE, MPLS VPNs according to RFC 4364 [20]
- Support for MPLS VPNs by the IP transport network connecting PDN Gateways and CGNs
- Procedures for the PDN Gateway to support UE with overlapping IPv4 addresses
  - Support of different APNs with different routing/forwarding for each of them (different routing instances, or layer 2 binding to attachment circuits)

Known issues of the solution:

- If overlapping private IPv4 addresses are used within one operation domain for the UEs, all traffic needs to go through the CGN. This could potentially result in non-optimal communication patterns for the scenario of direct IPv4 communication between UEs that are attached to the same CGN.

- GI-DS-lite involves the usage of NAT and therefore potential PCC issues due to NAT apply. Additional PCC issues may need to be considered for cases where meaningless or overlapping IPv4 addresses are used.

Known issues of the GRE implementation:

- GRE encapsulation overhead between the PDN Gateway and the CGN.

- Deployment dependent (e.g. scenarios where all UEs are assigned the same address), enhancements to PCC may be required. Traffic between UEs connected to the same APN via different PDN-GWs needs to go through the CGN even if no overlapping IP addresses are used by those PDN-GWs, or alternatively PDN-GWs must be aware of each others UE IP address ranges used and tunnel traffic among each other.

Known issues of the MPLS VPN implementation:

- Overlapping IPv4 addresses are only supported between operational domains, i.e. using different MPLS VPNs, but not within the same VPN.

- MPLS encapsulation overhead in the backbone.

- MPLS has to be introduced in the IP network used as transport, if not deployed already.

- If the PDN-GW or the CGN are deployed as CEs, attachment circuit overhead between PDN-GW or CGN and the backbone provider edge routers.

- If the PDN-GW or CGN are deployed as CE, they may (deployment dependent) need to implement an appropriate routing protocol for CE-PE peering.

- If the PDN-GW or the CGN are deployed as PEs:

  - MPLS overhead encapsulation overhead between each such node and the rest of the backbone

  - MP-BGP peering to all other PEs or to route reflectors, as well as MPLS encapsulation support in such nodes according to RFC 4364 [20]

Known benefits of the solution:

- Support for UEs with public, private, and overlapping private IPv4 addresses.

- No changes to the UE required.

- No changes to the IPv4 / IPv6 address-assignment procedures required.

- The CGN can be placed on the service provide IPv4 network edge and is not required to be collocated with the PDN Gateway.

- This solution does not introduce any additional tunnel overhead on the air-link, or on the access network for carrying the UE's IPv4 traffic. It leverages the tunnelling infrastructure existing between the UE and the PDN gateway.

- Solution to the public IPv4 address exhaustion problem through the use of NAT44. The NAT44 function is only required at a single location within the architecture.

- Solution to the private IPv4 address exhaustion problem through the use of overlapping private IPv4 addresses and softwires.

- This solution does not have any impact on the UE's roaming support.

- No impact on QoS/bearer procedures between UE and PGW/SGW.

Known benefits of the GRE implementation:

- If so desired, all the UE's in the mobility domain can be assigned the same IPv4 private address.

- No bearing on the type of transport network: Transport network can be IPv4 or IPv6.

This solution requires only a single IPv4 or an IPv6 transport tunnel between the PDN Gateway and the Carrier Grade NAT, with the GRE encapsulation mode. This single GRE tunnel is used for carrying all the IP traffic belonging to all the UEs supported on that PDN Gateway (i.e. the GRE-key is used to multiplex and differentiate traffic from multiple UEs onto the very same GRE-tunnel (which is identified by the addresses of the end-points)).

Known benefits of the MPLS VPN implementation:

- No additional tunnel overhead between the PDN-GW and CGN if MPLS is already deployed or only MPLS encapsulation on the backbone if not previously deployed.
- By sharing the same MPLS VPN for the same APN by several PDN Gateways and CGNs, traffic between end-users can be sent directly without going via any CGN.
- Different CGNs can be deployed in an MPLS VPN providing CGN redundancy by relaying on basic routing protocol mechanisms.

7.2.3 GI-DS-lite Applicability

Gateway-initiated Dual-Stack Lite applies to the following IPv6 migration scenarios outlined in clause 5:

- Scenario 1: Dual-stack connectivity with Limited Public IPv4 Address Pools
- Scenario 2: Dual Stack connectivity with Limited Private IPv4 Address Pools

7.3 Solution 3 - MS/UE IPv6-only deployment with stateful NAT64 support

7.3.1 Overview

When deploying an MS/UE with IPv6-only connectivity it will be able to communicate with other IPv6 reachable servers and peers if they are either dual-stack or IPv6-only connected. Since the decision to deploy an IPv6-only communications model in many cases will be a unilateral decision, there may be a need for an IPv6 transition mechanisms designed to enable transition and to support IPv6-enabled hosts and routers that need to interoperate with IPv4 hosts and utilize IPv4 routing infrastructures. Introducing transition tools such as the functional elements DNS64 and NAT64 into the network will enable an IPv6-only MS/UE to communicate with IPv4-only reachable servers and peers.

7.3.2 Description

MS/UE attaches to network APN(s) using applicable procedures described in TS 23.401 [9], TS 23.402 [10] and TS 23.060 [11] in order to get IPv6 connectivity to Internet. The operator assigns IPv6 prefixes to the MS/UEs allowing native IPv6 access to IPv6 networks. The MS/UE is provisioned with DNS server address of the DNS64 server which is used to create and return synthetic AAAA records for a queried FQDN that would only return A records in a regular DNS lookup. The synthetic AAAA record is used by the MS/UE as a destination address effectively sending the packets to the NAT64 function which translates IPv6 packets to IPv4 packets and vice versa to enable the communication between the MS/UE and the IPv4-only destination.

7.3.3 Functional Description

The MS/UE need to obtain IPv6 connectivity in order to be able to reach IPv6 services/peers including making queries to DNS and sending packets to NAT64.

The DNS64 function needs dual-stack connectivity in order to perform DNS lookups and answer MS/UE DNS queries.

The NAT64 function needs IPv6 and IPv4 connectivity to enable translated packet flows.

The IPv4 Internet connectivity is provided over a NAT64 function either co-located with the GGSN/PDN GW or elsewhere placed in the network as show in the following figures.

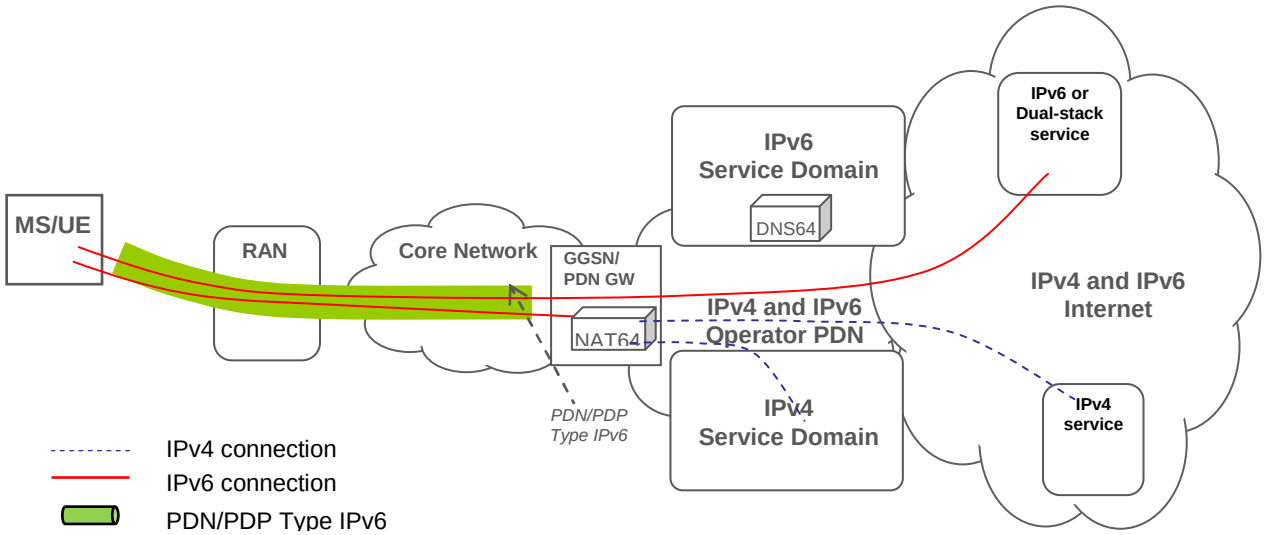

Figure 7.3.1: Example DNS64 and NAT64 functionality network placement with co-located NAT64 and GGSN/PDN-GW

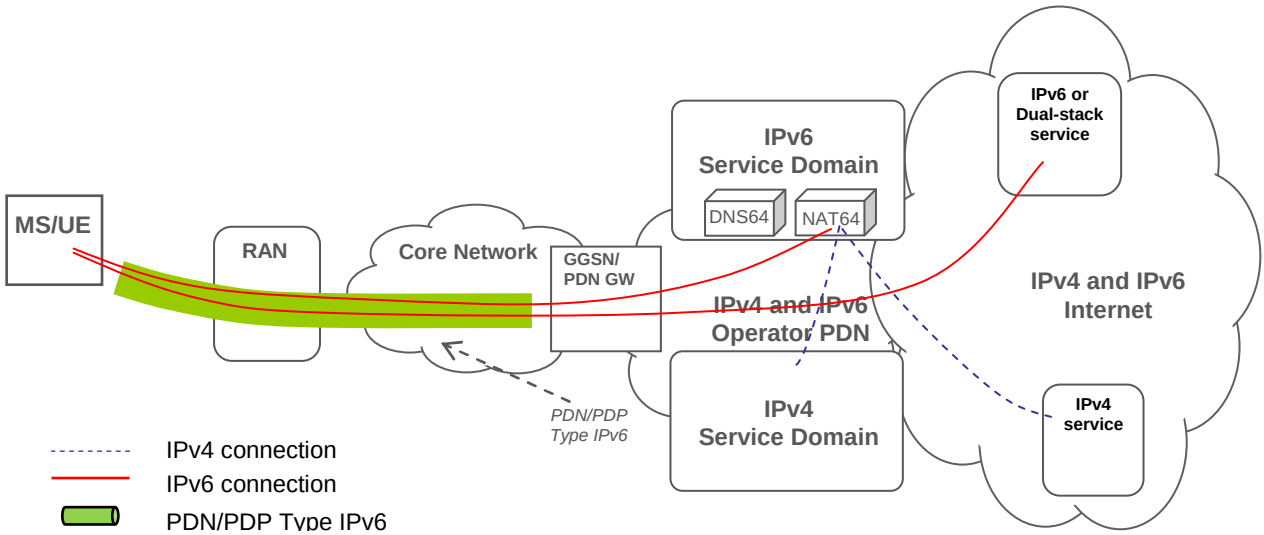

Figure 7.3.2: Example DNS64 and NAT64 functionality network placement with standalone NAT64

7.3.4 Server Flow Example

Suppose an IPv6 only MS/UE's IPv6 address is Y, the IPv4 only Server's IPv4 address is X, the DNS64 selects NAT64A as the NAT64 gateway for this service. The main procedures for the MS/UE visiting an IPv4 only server with IPv4 address X is illustrated in Figure 7.3.2. More details can be found in IETF RFC 6146 [17].

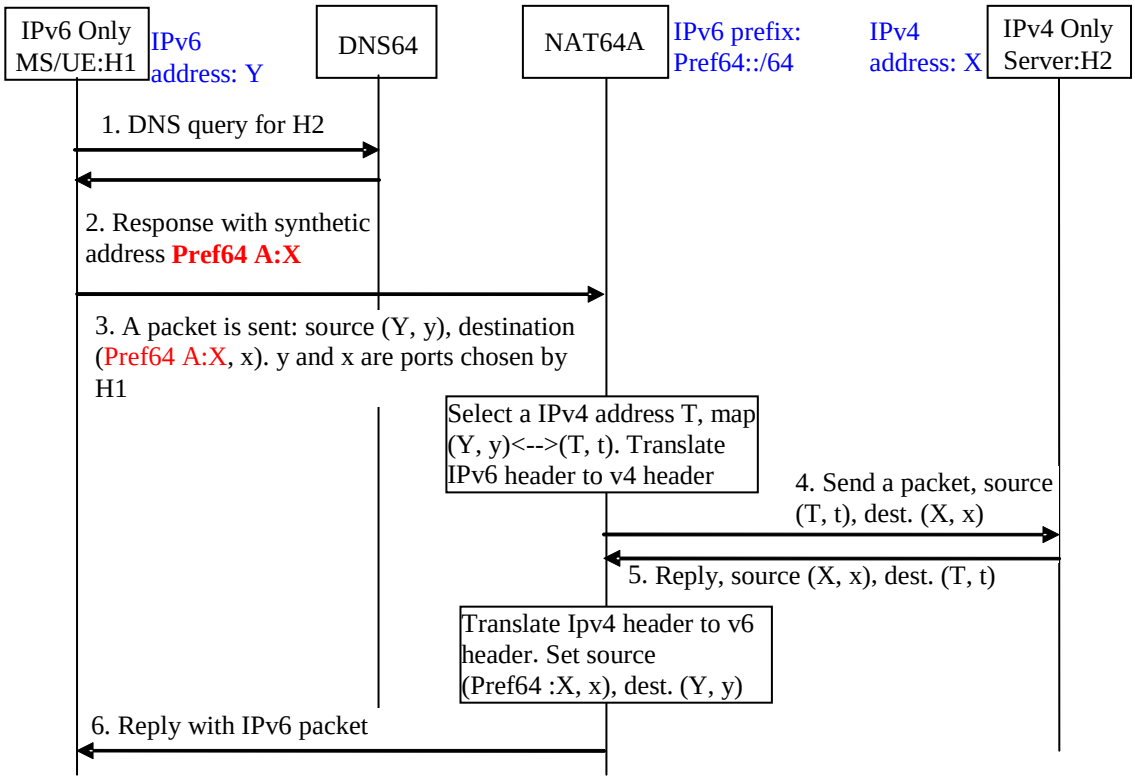

Figure 7.3.3: Message flow of NAT64

7.3.5 Evaluation

Besides end to end IPv6 communication the solution also allows communication between IPv6-only MS/UE and IPv4-only services/peers using NAT64 functionality.

Known issues of the solution:

- The session binding between MS/UE IPv6 address and public IPv4 address/port is not known to the PCC architecture. Therefore, depending on deployment and if the application is NAT aware and has access to the binding (as e.g. in the case of IMS), there may or may not be issues with applying PCC to the session.
- General NAT concerns, not specific to 3GPP networks, apply. For example, applications that embed IP addresses in the payload and are not NAT aware require additional functionality to work across NATs.
- In case of roaming with Local Breakout (PGW/GGSN in VPLMN), if there is a need to reach IPv4-only services, the VPLMN operator would be required to deploy NAT64/DNS64 in order to provide the same user experience using IPv6-only connections for IPv4-only services as in non-roaming scenarios.
- It is unclear how IPv4 literals are supported.

Known benefits of the solution

- This solution requires no changes to the MS/UE.
- The existing bearer and session management procedures can be used without any change.
- This solution does not introduce any additional tunnelling overhead on any interfaces.
- This solution has no impact to the 3GPP network architecture, no new interface or network element is needed. This solution can be deployed without any additional normative specification within 3GPP. Limitations described under "known issues" above apply.
- No changes to the IPv6 address-assignment procedures required.
- No bearing on the type of transport network: Transport network can be IPv4 or IPv6.

- NAT64 can be either co-located or separate from GGSN/PDN-GW.
- Solution to the public IPv4 address exhaustion problem through the use of stateful NAT64.
- No impact on QoS/bearer procedures between UE and PDN GW/S GW.

7.3.6 Applicability

This solution applies to scenario 3.

Given the solution description above, the described functionality can be configured in currently deployed mobile networks as well as in future deployments regardless of 3GPP access technology. When to deploy such a setup in an operator's network is more of a business and operational decision.

8 Evaluation

TBD

9 Summary

Recommended strategies for 3GPP specifications based on the solutions developed above are provided in clause 10 "Recommendations".

In addition, operators may consider some of the solutions described in Annex B for deployment. As these mechanisms are not considered further in 3GPP, they are not expected to be further updated or maintained in 3GPP documentation.

10 Recommendations

3GPP specifications recognize two main strategies to provide IPv6 connectivity to UEs.

For the first strategy, the operator may provide IPv4 and IPv6 connectivity for the UE. According to the scenario considered, the operator will assign a public IPv4 address or a private IPv4 address in addition to an IPv6 prefix. The operator can select one of the technical solutions described in clause 7 of this document.

The second strategy, consisting of providing the UE with IPv6-only connectivity, can be considered as a first stage or an ultimate target scenario for operators. The operator can use NAT64/DNS64 capability to access to IPv4-only services if access to IPv4 services is needed.

Annex A:
Reference Scenarios for NAT in the EPC

IPv6 migration may involve the use of NAT. The use of NAT in the EPC raises several issues, in particular as it relates to interactions with dynamic Policy and Charging Control (PCC). In this clause examines these issues. This annex only covers non-roaming scenarios. NAT related to roaming scenarios is not being considered in this TR. Note that while this annex focuses on NAT in EPC, the considerations in this annex apply for GPRS networks in a similar way.

A.1 UE and AF In the same Address Realm

Figure A.1 and figure A.2 show scenarios where no NAT function exists between the AF and the UE, which is the typical deployment for e.g. IMS. NAT is employed beyond the AF. UE, PCC, and AF are all within the same addressing domain, hence NAT does not impact PCC services.

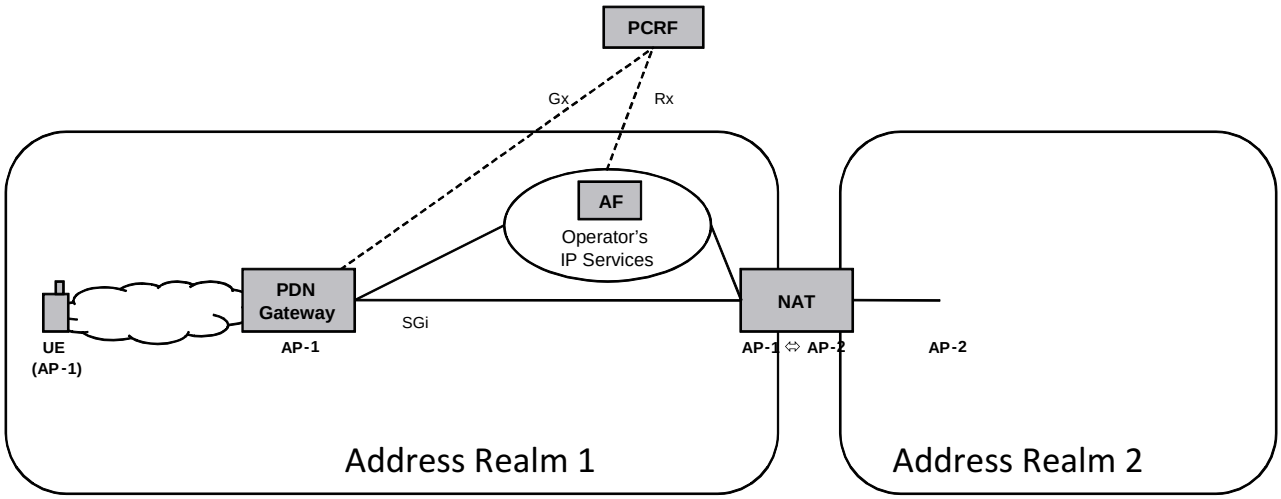

Figure A.1: EPC with standalone NAT-function, Services in front of NAT

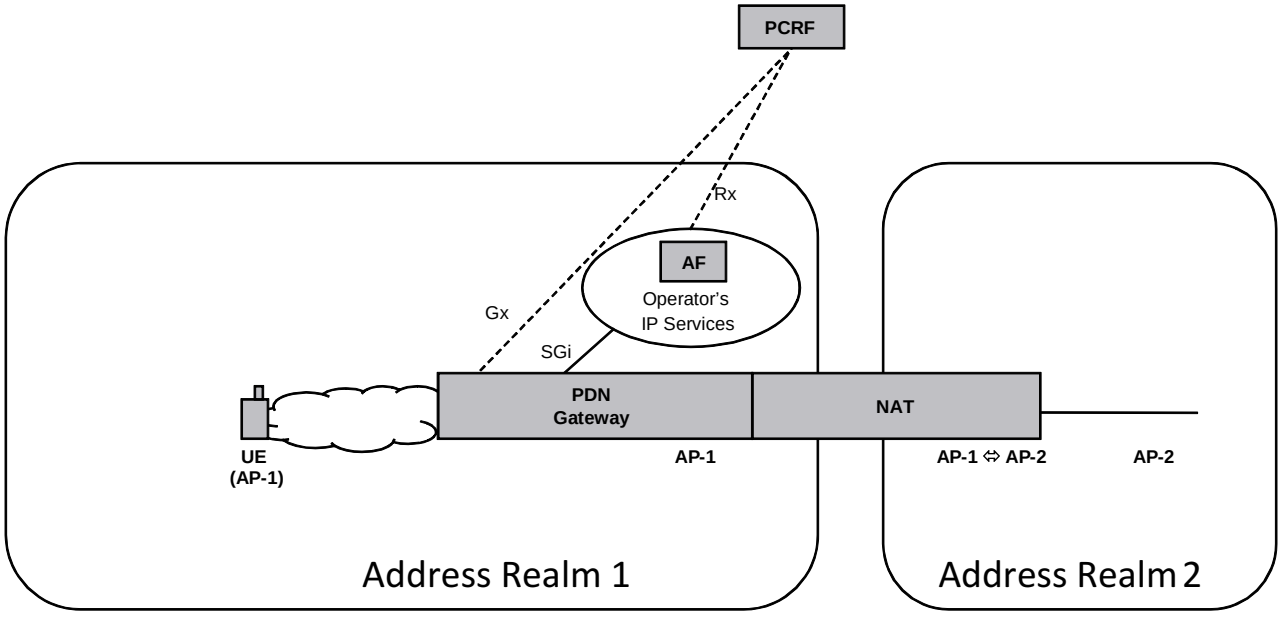

Figure A.2: EPC with NAT-function integrated with the PDN-Gateway, Services in front of NAT

A.2 NAT between UE and AF

Scenarios in this clause cover non-roaming cases where the NAT function resides between the PDN Gateway and the Application Function. The basic scenario is illustrated in figure A.3 with a standalone NAT function, and in figure A.3 with a NAT function integrated with the PDN-Gateway.

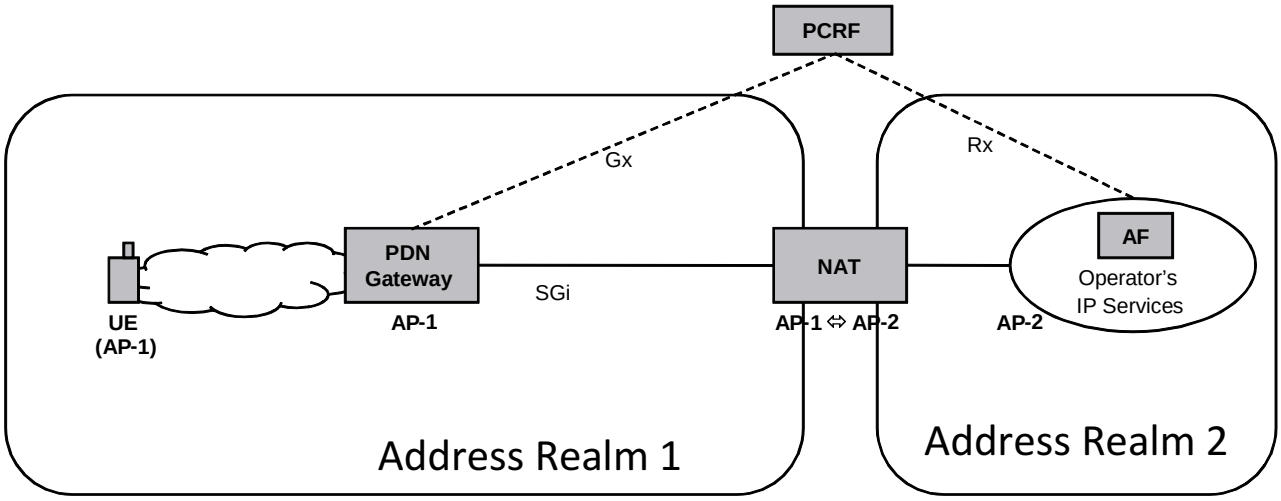

Figure A.3: Application Function Realm Traversal with Standalone NAT

When the UE performs an Initial Attach, it is assigned an IP-address, A-1, by the PDN-Gateway in address realm 1. As part of the attach procedure, the PCRF is informed of the A-1 IP-address (or IPv6-prefix) assigned to the UE.

Subsequently, the UE may invoke an Application Function, that resides in Address Realm 2. The UE uses an IP-Address and Port pair AP-1 for this application invocation.

NOTE: There may be more than one Address and Port pair, however for simplicity, we consider only one for now.

The associated IP datagrams traverse the NAT function, where they are translated to a new IP-Address and Port pair, AP-2, and the NAT-function performs the corresponding mapping. In case of a stateful NAT operation, a corresponding mapping is stored in the NAT function. The IP datagrams that reach the application function will thus contain AP-2 as the source IP-Address. AP-2 is possibly unknown to the UE. The AF may in turn interact with the PCRF to perform Policy and Charging Control (PCC) for the application.

Several issues can be seen at this point:

- The Rx and Gx interface are crossing address realms and hence a mapping between the two will be necessary somewhere.
- The mapping between AP-1 and AP-2 is not done until the application is invoked.
- The mapping between AP-1 and AP-2 may not happen until sometime after the application has been invoked. For example, for applications that establish media streams (such as IMS using SIP/SDP), the corresponding mappings for the media streams may not be established until after the session is established, yet PCC interactions for the media streams are needed before that.
- Application Functions that use embedded IP-address information (e.g. AP-1) in application signalling need to deal with the implications of such address information crossing address realms.

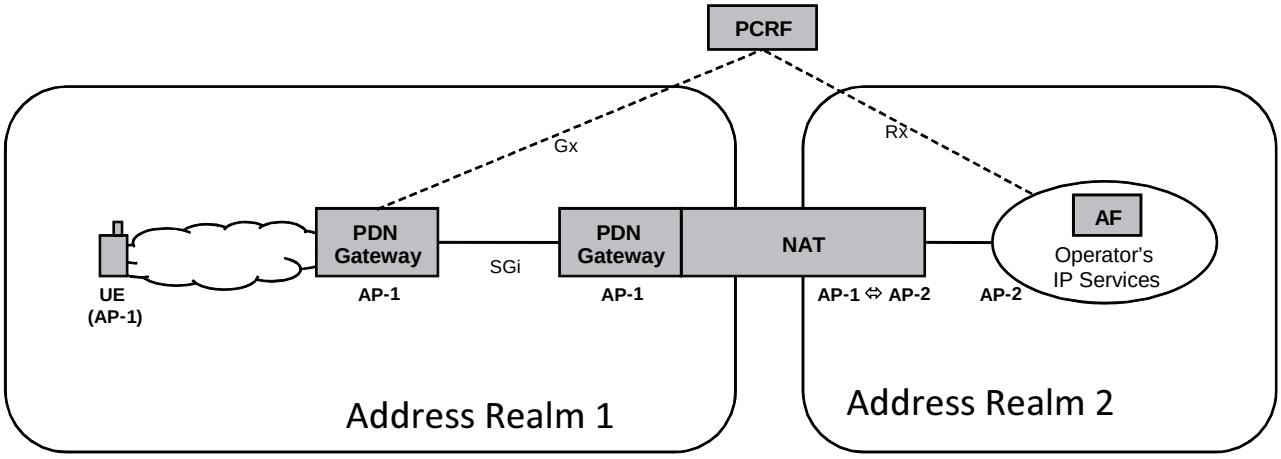

Figure A.4 - Application Function Realm Traversal with Integrated NAT

The scenario with integrated NAT shown in figure A.4 is very similar to figure A.3, and hence it raises many of the same issues.

A.2.1 Overlapping IPv4 address realms

The limited number of IPv4 addresses offered by even the largest private IPv4 address space (i.e., 10/8, which holds less than 2^24 or 16.8M addresses) implies a need for large operators with IPv4 enabled UEs to employ overlapping IPv4 address realms. The considerations within this clause build on top of those discussed above, i.e. the issues identified above apply to the scenarios with overlapping IPv4 address realms as well.

- Two scenarios are considered: **Overlapping address realms exposed to a single PCRF**
  figure A.5 expands on the basic scenarios by introducing the notion of overlapping address realms (1a and 1b), where AP-1 (or A-1) may be assigned to and used by multiple UEs in the network, as long as those UEs reside in different address realms. This adds the issue that PCC may need to traverse address realms, however with overlapping IP address realms, there is now also a need to identify the correct address realm. Similarly, if the same NAT function is used between multiple address realms with overlapping IPv4 addresses, the NAT function (and possibly the network) needs to incorporate functionality to differentiate these different NAT address realms.

  Figure A.3 shows one possible scenario where the overlapping address realms are exposed to the PCRF. PCC-related issues of the overlapping address realms are shielded from the Application Function. If PCRF receives requests from the AF it needs to figure out which address realm (and P-GW) to interact with for Gx.

- **Overlapping address realms with individual PCRFs:** Figure A.6 shows a scenario with two address realms with overlapping addresses. There is a PCRF for each address realm. From a NAT perspective, this scenario resembles the scenario shown in figure A.3. Additional considerations for the overlapping address realms are not required. The AF and/or DRA needs to follow the appropriate PCRF selection procedures to identify the appropriate PCRF.

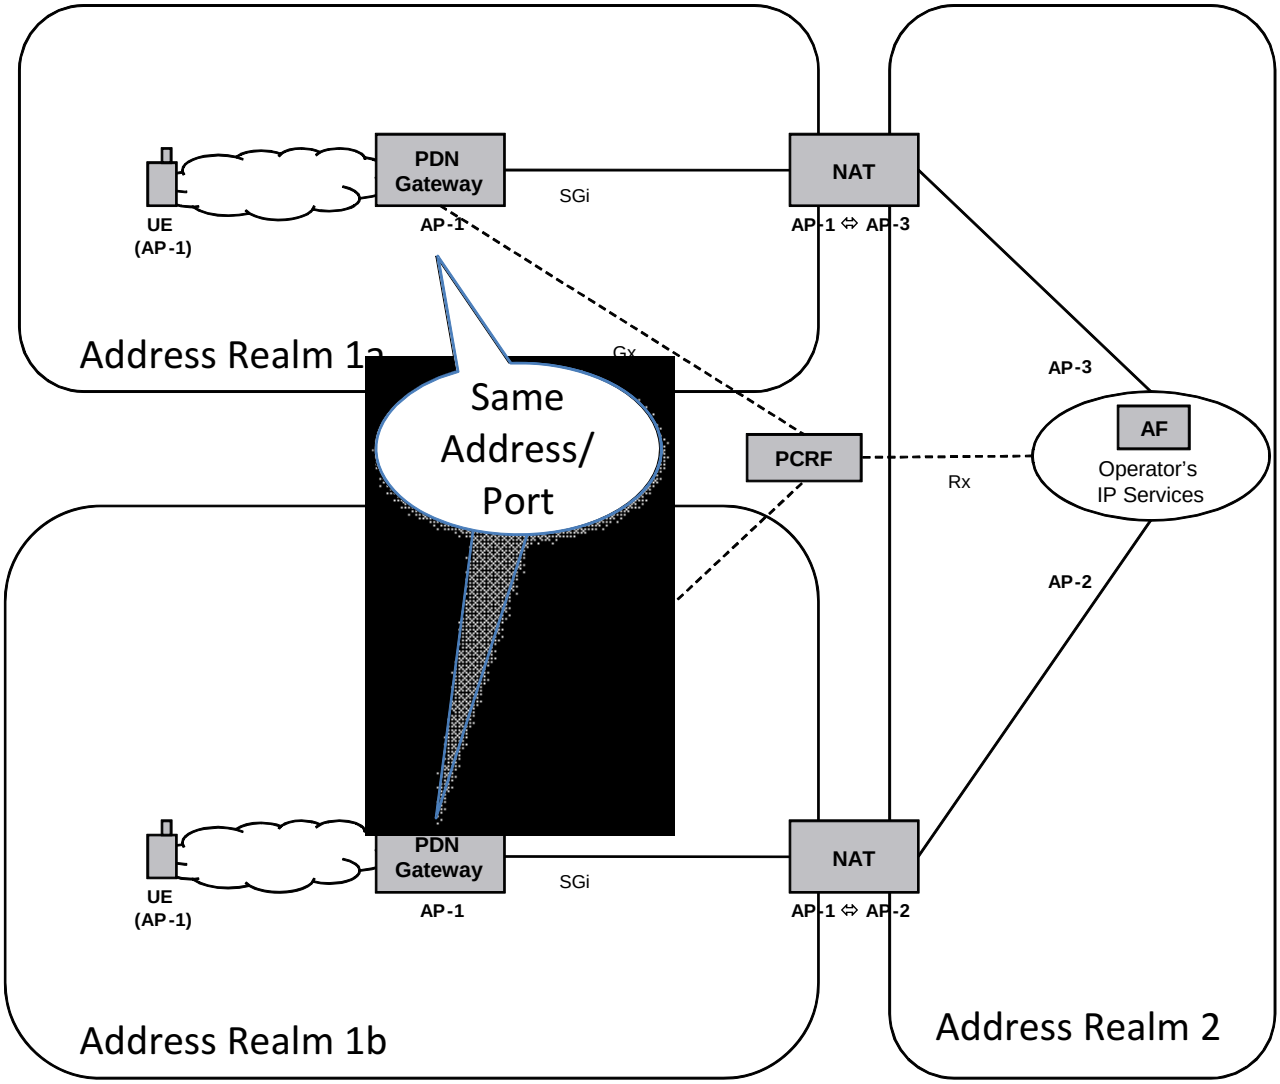

Figure A.5: Application Function Realm Traversal with Overlapping Address Realms on Gx

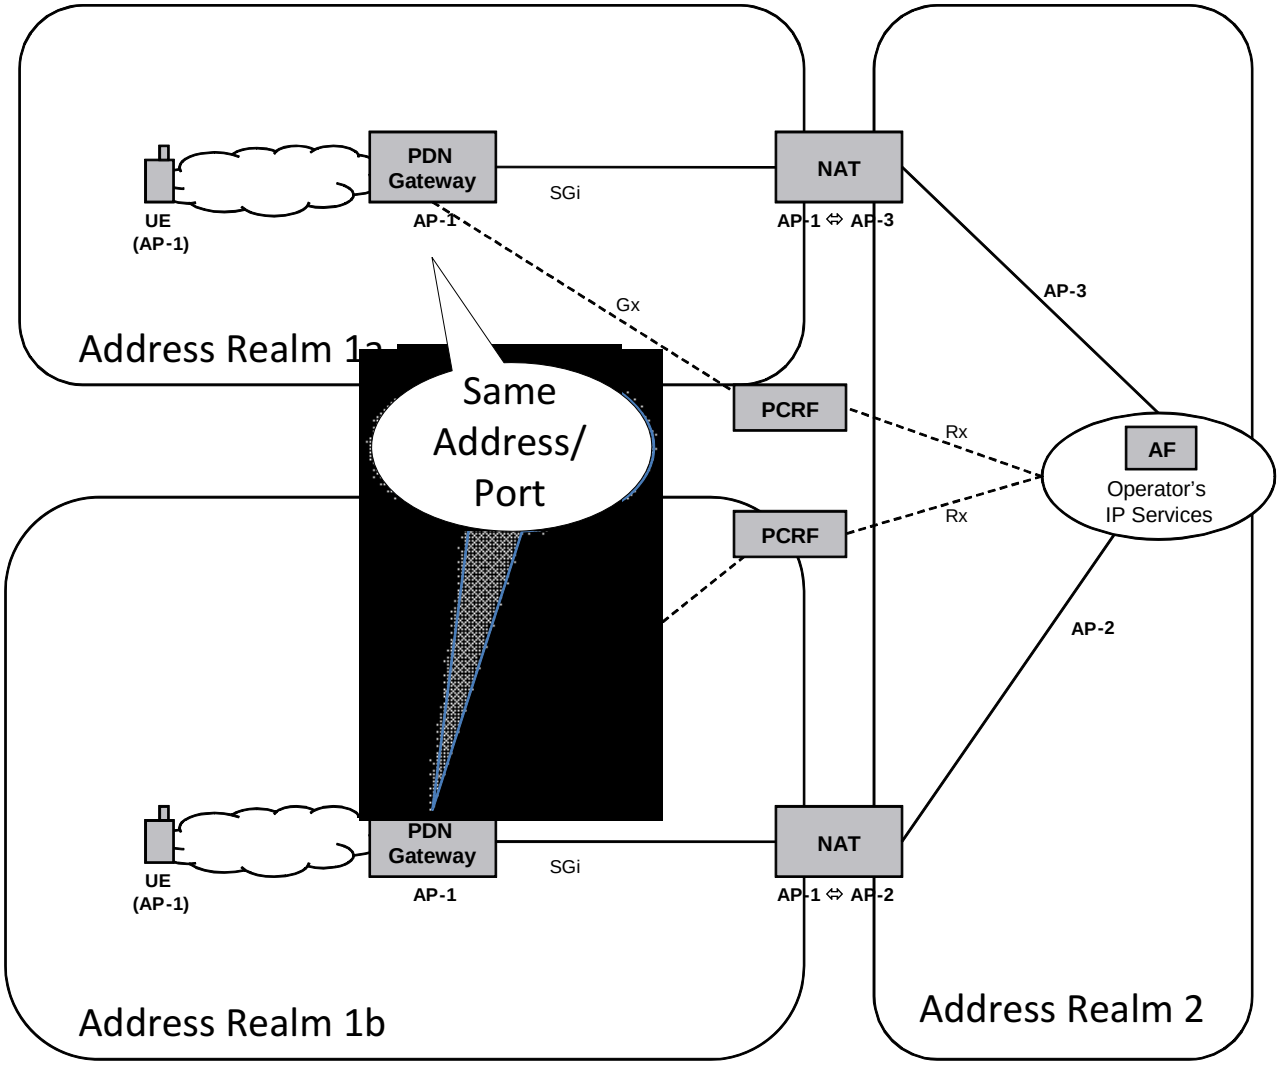

Figure A.6: Application Function Realm Traversal with Overlapping Address Realms on Rx

Annex B:
Overview of Solutions for IPv6 Transition

B.1 Solution 1 - Dual-Stack Lite Architecture

B.1.1 Solution 1 Description

Dual-Stack Lite architecture [2] can be understood as IPv4 packets being encapsulated using either IPv6 or some L2 technology. The tunnel endpoint is usually the Carrier Grade NAT (CGN). Since the hosts are not provisioned with an IPv4 address, they have to self-generate their own IPv4 address from the private IPv4 address pool. Thus, these self-generated IPv4 addresses may overlap, and packets from different hosts may arrive to the CGN with the same private IP address. The CGN differentiates hosts with same private IPv4 address based on information provided by encapsulation technology. When packets are destined to the IPv4 Internet, CGN will act as a NAT. Several options exist for deploying DS-Lite.

The encapsulation method can be chosen at least from the following set:

- Plain IPv6: IPv4-in-IPv6 is the basic DS-Lite encapsulation scenario. In this scenario the UE encapsulates IPv4 packets into IPv6. The CGN can be a separate entity or integrated to e.g. PDN GW. Only an IPv6 bearer is needed.

- GRE: When PMIP6 is used, the MAG can encapsulate IPv4 into GRE tunnel. CGN has to be implemented in LMA. No UE impact. A dual-Stack bearer is needed.

- GTP: When GTP is used, PDN GW must implement CGN. No UE impact. A dual-Stack bearer is needed.

- DSMIP6: The HA must implement CGN. Only an IPv6-bearer is needed. The UE must implement standard DSMIP6 support.

There are also other encapsulation methods, such as L2TP, but those are not included in this study.

The common feature of DS-lite is that all IPv4 communication from UEs will have to go through NAT functionality, even if traffic is destined to the operator's own services (no hairpinning is possible, as there is no IPv4 address allocation). Consequently DS-lite is best suited for IPv4 Internet access by legacy applications, which are able to initiate communication and connections. In such a deployment scenario, the majority of new applications and operator services would be accessed with IPv6.

B.1.1.1 Plain IPv6 encapsulation in 3GPP architecture

When plain IPv6 encapsulation is used, DS-Lite can be deployed independently over existing 3GPP IPv6 access. The UE is required to be able to discover the CGN's IPv6 address (for example by using stateless DHCPv6), and then to encapsulate IPv4-over-IPv6 to the CGN, which does the decapsulation and network address translation. The CGN can be a stand-alone entity, or integrated into the PDN GW. The CGN differentiates UEs with same IPv4 address based on their globally unique IPv6 address. When using IPv6 encapsulation, it is enough to establish IPv6-only bearers to between the UE and PDN GW.

Known issues:

- MTU: to avoid fragmentation and dropped packets MTUs must be configured properly. For IPv6 communication, the UE will use the default MTU of the bearer or the MTU advertised in Router Advertisements, while for IPv4 communication, the UE will use an MTU of (IPv6_MTU-20) bytes.

- Tunnelling overhead: an IPv6 header (128 bits) is added to each IPv4 packet.

- IPv4 P2P communication: all IPv4 based communication, including P2P, must traverse through CGN.

- QoS: 3GPP TFTs are limited in such a way that it is not possible to differentiate traffic based on information in the inner headers of a tunnel.

Known benefits:

- Simple UE side implementation.
- Can be deployed over existing 3GPP networks, with the known issues.

B.1.1.2 GRE encapsulation

When PMIP6 is used for network based mobility, it is possible for the LMA to use GRE identifiers to differentiate between UEs. The CGN function must reside in the LMA, as it is the only entity capable of differentiating between UEs having the same IPv4 address. The MAG will need to differentiate UEs with same IPv4 address by some other identifier (such as the default bearer id). UEs do not need to be modified, as they are provided with native dual-stack connectivity. When using GRE encapsulation, a dual-stack bearer (or two single stack bearers) needs to be established between UE and MAG.

Known issues:

- Requires support on the MAG and the LMA.
- Cannot be deployed into existing 3GPP networks.
- IPv4 P2P communication, all IPv4 based P2P communication must traverse through CGN.

Known benefits:

- No UE changes mandated (but UEs may need to support some other encapsulation for other access technologies than 3GPP access)
- Interworks with the existing QoS schemes.
- No tunnelling overhead over the air interface

B.1.1.3 GTP encapsulation

A special case is the GTP based solution, where the PDN GW implements CGN and differentiates UEs based on the TEID It allows allocation of the same IPv4 address for all hosts. When using GTP encapsulation, a dual-stack bearer (or two single stack bearers) needs to be established between UE and MAG.

Known issues:

- Requires support on the PDN GW.
- Cannot be deployed into existing 3GPP networks.
- IPv4 P2P communication, all IPv4 based P2P communication must traverse through CGN.

Known benefits:

- No UE changes mandated (but UEs may need to support some other encapsulation for other access technologies).
- Interworks with existing QoS schemes.
- No header overhead over the air interface.

B.1.1.4 DSMIP6

With DSMIP6, it is possible to provide session continuity during inter-technology handovers and at the same time provide an IPv6 transition solution. DSMIP6 can, by definition, always provide dual-stack connectivity independently of the address family of care-of address(es) obtained within the visited network. In case public IPv4 addresses are scarce, and private IPv4 address space is too small for ordinary IPv4 Network Address Translation to suffice, the DSMIP6 Home Agent could implement the CGN function and thus be able to allocate the same private IPv4 address for multiple UEs. A DSMIP6 HA behaving as a CGN can be seen as instance of Dual-Stack Lite architecture.

Known issues:

- Tunnelling overhead from the DSMIP6 header.
- IPv4 P2P communication, all IPv4 based P2P communication must traverse through a CGN.

Known benefits:

- The UE does not need to implement anything special over standard DSMIP6 support
- Can be deployed over existing 3GPP networks, with the known issues
- QoS can be provided as currently.

B.2 Solution 2 - A+P architecture

The main principle of the IPv4 A+P solution (see [3][4]) is to assign the same IPv4 address (called Primary IPv4 Address) to several end-users' devices and to constraint the source port numbers to be used by each device. In addition to the assigned IPv4 address, an additional parameter, called Port Range, is also assigned to the customer's device.

This allows the allocation of the same public IPv4 address to multiple UEs, as they all will use different sets of ports. By doing so, the need for having NAT functionality in the network disappears.

As the IPv4 address is shared among multiple hosts, A+P addresses can only be used in point-to-point links (not in shared medium) and routing must be based on both the IP address and the port number. The entity that routes IP packets based on the port number is called a Port Range router (PRR).

The link between the UE and the PRR can use any encapsulation method, i.e. IPv6, GRE, GTP and DSMIP6.

As the UE has a limited set of ports which it is allowed to use, the UE must be modified to use allowed ports only. This can be realized e.g. by modifying the applications to deal with shared addresses, or having an internal NAT within the UE which translates between a self-generated private IPv4 address shown to the internal applications and the port restricted public IPv4 address received from network.

For outbound communications, a given port-restricted device proceeds according to its classical operations except for the constraint to control the source port number assignment so as to be within the assigned Port Range. The traffic is then routed without any modification inside the PLMN and delivered to its final destination.

For inbound communications, in the base IPv4 A+P variant, the traffic is trapped by the Port Range Router (PRR).

B.2.4 Port Range Router (PRR) function

B.2.4.1 General

As mentioned above, a PRR function is required to be enabled in the data path so as to deliver incoming packets to the appropriate UE among those having the same IPv4 address.

This function may be embedded in current nodes or hosted by new nodes to be integrated in the PLMN, in particular a PRR function can be embedded in a GGSN, PDN GW, WIMAX ASN GW, 3GPP2 PDSN, etc.

B.2.4.2 PRR in binding mode

In binding mode, a PRR associates an IPv4 address and a port range with a specific identifier called routing identifier (e.g., GTP tunnel identifier, IPv6 address). This identifier is used to forward the packets to the suitable device among all those having the same IPv4 address. These associations are stored in a table referred to as port range binding table.

The routing identifier may be an IPv6 address belonging to the prefix assigned to the UE, or the GTP tunnel identifier.

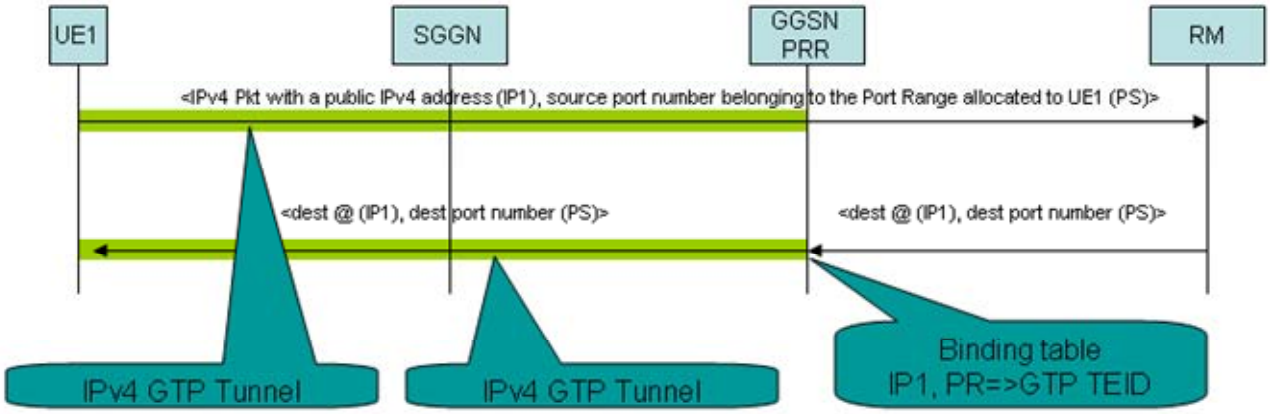

Figure B.2.5: A+P Flow Example

If an IPv6 address is used as a routing identifier, the PRR is not required to be co-located with GGSN. In such case, once the encapsulation is undertaken by the PRR, the appropriate GGSN/P-GW will receive the IPv4-in-IPv6 packets. Then, the GGSN/PDN GW proceeds to its "normal" operations in order to forward the IPv4-in-IPv6 packets to the appropriate UE: in particular, it uses the maintained IPv6 PDP context to relay those packets.

B.2.4.3 PRR in stateless mode

In addition to the binding mode defined in clause B.2.4.2, a stateless IPv6 A+P mode can be implemented as defined in IETF RFC 6346 [3]. In such mode, no port range binding table is required.

In this case, for incoming packets, the PRR encapsulates a received IPv4 packet in an IPv6 one using the following information:

- The destination IPv6 address is constructed using the shared IPv4 destination address and port number plus the IPv6 prefix which has been provisioned to the PRR. To do so, the PRR retrieves the destination IPv4 address and destination port number from the received IPv4 packet.

NOTE: To illustrate this behaviour, assuming that the PRR is provisioned with 2a01:c0a8::/29 as a prefix to build IPv4-Embedded IPv6 addresses, the IPv4 destination address equal to 193.51.145.206 and the port number equal to 19039 (0100101001011111), then the corresponding IPv6 address (which falls into a prefix assigned to the UE) is 2a01:c0aE:099C:8E72:52F8::/128

```
2a01:c0a 1 11000001001100111001000111001110 0100101001011111 ::
           --------193.51.145.206---------  -----port-------
```

- The source IPv6 address is one of the global IPv6 addresses of the PRR.

For more information about the stateless mode, the reader is invited to refer to IETF RFC 6346 [3].

B.2.6 Requirement on UEs

Mobile UEs must be able to constrain their source port numbers and to use only source port numbers within the allocated Port Range. If an IPv4 packet is received by a given port-restricted UE, with a destination port number outside the assigned Port Range, the packet must be discarded. Furthermore, port-restricted UEs must be able to enforce configuration data received from the PLMN so as to constrain its Port Range.

According to the enforced routing identifier mode (GTP tunnel or IPv4-in-IPv6 tunnel), an encapsulation / de-capsulation function may be required. However, when GTP is used, no extra tunnelling technique is required to be supported by the UE. Nevertheless, if IPv4 packets are transported over IPv6, then the IPv4-in-IPv6 encapsulation/de-capsulation function is required.

As a conclusion:

- it is mandatory to support port-restriction feature;
- it is optional to support another tunnelling technique in addition to GTP.

B.2.7 Updating legacy UEs

For the efficiency of public IPv4 address sharing, some of the legacy UEs may be updated to be port-restricted UEs owing to a software update. No operation is required in the hardware. For instance, some Linux-based mobile OSs such as Maemo supports the Iptables capabilities; the implementation of port restriction on the UE is then done with two command lines.

Nevertheless, from an operational perspective, updating UEs may not be obvious.

B.2.8 Co-existence with other transition techniques

A+P can be deployed jointly with other IPv6 transition techniques such as DS-Lite. In particular, a DS-Lite AFTR (Address Family Translation Router) can delegate a set of port numbers to the UE to be used for "push" services. No NAT operation would be achieved by the CGN for delegated port numbers.

B.2.9 Applicability

A+P can be deployed in several configuration schemes:

- Dual-Stack PDP context / bearer with a shared IPv4 address. The GTP tunnel identifier will be used by the PRR for forwarding incoming IPv4 packets;
- Single IPv6 PDP context / bearer: IPv4-in-IPv6 encapsulation is used to exchange IPv4 packets between the UE and the PRR,
  - with a binding table in the PRR (binding mode)
  - without any binding table in the PRR (stateless mode)

B.2.10 Evaluation

The hard-partitioning of the port space reduces the efficiency of the A+P architecture. Ports-ranges assigned to a UE are no longer available for other UEs - even if these ports are not used. In consequence, the efficiency of A+P wrt IPv4 address utilization is less than with a centralised NAT functionality.

Known issues:

- The UE needs to be modified to support A+P scheme
- The gateway needs to forward not only based on IP address but based on address plus port. The network needs to implement PRR in similar places as CGN in the DS-Lite approach
- The backend RADIUS system needs to be changed as subscribers can no longer be identified by IP address only, but by IP address and port
- In the IPv6 tunnelling approach QoS differentiation between bearers cannot be provided easily
- The solution works only with applications using transport protocols, which have concept of port numbers (such as UDP and TCP). There will be challenges with protocols which use plain IP.
- The solution sets restrictions to applications within in the UE, as the allocation of fixed port numbers becomes more complicated.
- For ICMP messages, the UE must use an ICMP query identifier within the allocated port range, otherwise the response will not be received by this UE.

Known benefits:

- The UE has access to public IPv4 address, which simplifies the behaviour for P2P applications such as VoIP.
- Allows IPv4 lifetime extension if used with GTP/GRE.

- Legal requirements for tracing which traffic flow was originated from which UE is simpler than in CGN solutions, as the operator does not need to store each flow but only A+P allocation information.
- In GTP/GRE/DSMIP6 based solutions QoS can be provided.
- A+P allows for an incremental migration to IPv6-only network
- A+P can be fully stateless when used together with IPv6
- No per-state sessions are maintained in the PLMN realm
- Unlike NAT44/NAT64, no dynamic state synchronization is required to ensure service robustness
- No ALG is required to be implemented in the service/PLMN realm
- No extra-cost on the UE to support NAT traversal techniques is required
- Unlike double NAT solutions, peer-to-peer services can be delivered with one exception as documented in [23]
- No Keep-alive messages are required to maintain NAT entries. This characteristic mitigates battery consumption issues induced by "Always-on" services. Especially, the use of short intervals between keep-alive messages has a big impact on the battery consumption (See [24] for more details about complications to tweak UDP timers in NAT devices and also keep-alive intervals used by UE). Moreover, the network load is more optimized since the load induced by keep-alive messages in the context of CGN solutions is avoided.
- Latencies and related problems of NATs are avoided.

B.3 Solution 3 - Protocol translation

Translation of IPv6 communication to IPv4 communication, and vice versa, is one way of providing connectivity between IP address families, see [5], [6], [7], [8]. If an UE would be strictly IPv6-only, it would be enough to have stateless or stateful NAT64 function in a network to provide access to IPv4-only destinations. However, as the UE is probably going to be running IPv4-only applications as well, a fully network based solution is not possible.

A host based translation approach enables the usage of IPv4-only applications on a UE which only has IPv6 access connectivity. Essentially, the UE implements protocol translation from IPv4 to IPv6 (NAT46), and thus all communications sent by the UE is IPv6-only. An IPv6-to-IPv4 translation (NAT64) is needed in the network for those cases where the destination happens to be in IPv4-only domain. However, if the destination has IPv6-connectivity, only NAT46 translation is needed within the UE.

Known issues:

- Requires protocol translation implementation within the UE
- ALGs are required in the UE to allow IPv4-embedding IPv4-only applications to communicate (such as FTP/SIP).

Known benefits:

- Direct point-to-point connectivity is possible, as IPv6 packets do not need to traverse via CGN
- Allows IPv4-only applications to access IPv6-only destinations without any translation taking place within the network.
- Less MTU problems due to the avoidance of a tunnel header
- Can be deployed in current 3GPP networks/technologies, as 3GPP network would consider all traffic IPv6-only (IPv4 awareness is only at edges)
- QoS can be provided as currently

B.4 Solution 4 - Per-interface NAT44

B.4.1 Overview

Per-Interface NAT44 (also known as Dual Stack Extra Lite) is a technique that relies on Layer 2 information for de-multiplexing NAT operations, as described in IETF RFC 6619 [19]. When applied to 3GPP networks, a NAT function can be embedded in the GGSN and/or P-GW, and UEs may be configured with the same IPv4 address. If NAPT44 is co-located with GGSN/PDN GW the NAT state can use the identity of the MS/UE mobility tunnel instead of the MS/UE assigned IPv4 address for managing the NAT session bindings. Packets are then translated and forwarded to their destination (either internal or external). Distinct pools may be configured on the GGSN depending on the APN. This procedure can be implemented in single stack PDP context / EPS bearers or in dual stack PDP contexts / EPS bearers. Since this approach allows UEs to be configured with the same IPv4 address, it resembles GI-DS-lite, as described in clause 7.2, but with the CGN collocated with PDN GW and without softwire tunnel. The approach is illustrated in Figure 7.Y.1 allows each MS/UE to use the same private IPv4 address range.

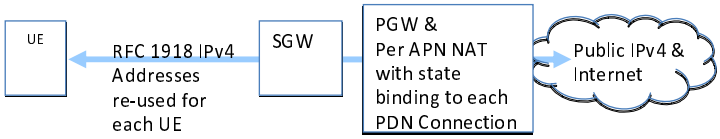

Figure B.4.1: Per-Interface NAT. Overlapping RFC 1918 [26] address space for the same APN with per-interface (PDN Connection) basis NAT binding

B.4.2 Evaluation

Known issues of the solution:

- The session binding between private and public IPv4 address/port is not known to the PCC architecture. Therefore, depending on deployment and if the application is NAT aware and has access to the binding (as e.g. in the case of IMS), there may or may not be issues with applying PCC to the session.
- General NAT concerns, not specific to 3GPP networks, apply. For example, applications that embed IP addresses in the payload and are not NAT aware require additional functionality to work across NATs.

Known benefits of the solution

- This solution requires no changes to UEs, it can be used with legacy dual-stack UEs.
- This solution has no impact to the 3GPP network architecture, no new interface or network element is needed. It can be deployed without any additional normative specification within 3GPP. The existing bearer and session management procedures can be used without any change. The limitations described under "known issues" above apply.
- This solution does not introduce any additional tunnelling overhead on any interfaces.
- This solution with the appropriate deployment supports UEs with overlapping address space, thus there is no limitation of the number of subscribers.
- Support for UEs with public, private, and overlapping private IPv4 addresses. If so desired, all the UE's in the mobility domain can be assigned the same IPv4 private address.
- No changes to the IPv4 / IPv6 address-assignment procedures required.
- No bearing on the type of transport network: Transport network can be IPv4 or IPv6.
- Solution to the public IPv4 address exhaustion problem through the use of NAPT44.
- Solution to the private IPv4 address exhaustion problem through the use of overlapping private IPv4 addresses.

- This solution does not have any impact on the UE's roaming support.
- No impact on QoS/bearer procedures between UE and PDN GW/SGW/GGSN.

B.4.3 Applicability

This solution applies to scenario 1.

This approach also suggests solutions to address scenario 2, based on similar considerations as described in 7.1.2.2.3.

Given the solution description above, the described functionality can be configured in currently deployed mobile networks as well as in future deployments regardless of 3GPP access technology. When to deploy such a setup in an operator's network is more of a business and operational decision.

B.5 Void

B.6 Void

B.7 Solution 7 - BIH/NAT64

B.7.1 Overview

During IPv6 migration, the network may only provide IPv6 only for reasons such as simplify the maintenance and reduce the management cost. However, it is not easy to mandate all the applications in the UE to update to support IPv6 in the first place. Therefore, the IPv4 applications in the UE are expected to still be able to access IPv4/IPv6 services.

The BIH/NAT64 solutions described here combines the BIH [7] module in the host and a NAT64 GW on the network side. The solution addresses the following scenarios:

1) The legacy IPv4 applications which reside in the host could continue to run and access IPv4 services through an IPv6 only network.
2) IPv6 applications access IPv4 services/peers.
3) The IPv4 applications which reside in the host could access the IPv6 servers/Peers.

The BIH in the host translates the IPv4 application's packets into IPv6 and the NAT64 translates IPv6 packets into IPv4 and vice versa. Such process can be depicted in Figure B.7.1.

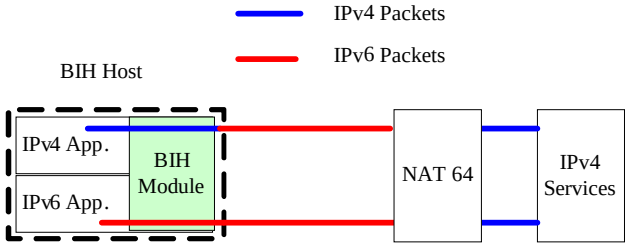

Figure B.7.1: Scenarios that BIH/NAT64 can address

B.7.2 Solution Description

The network architecture of deploying BIH/NAT64 in EPS is illustrated in Figure B.7.2. There are mainly two entities involved, e.g., BIH host/UE and NAT64.

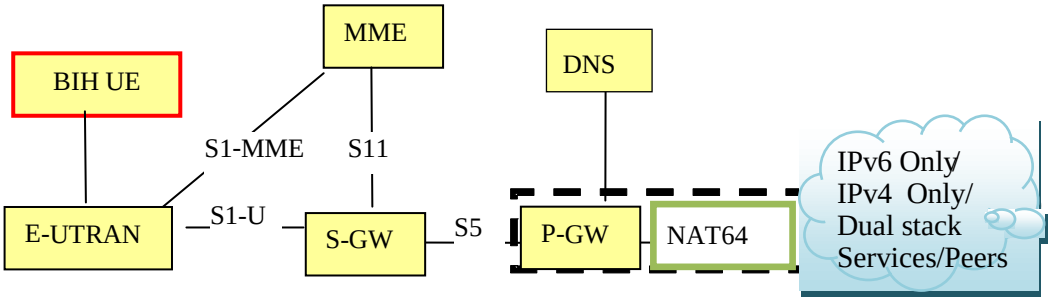

Figure B.7.2: The architecture of BIH/NAT64

The BIH host is a dual stack with the BIH module in it. It is could be either enhanced BIA (RFC 3338 [29]) or BIS (RFC 2767 [30]) and, in both cases, the Extension Name Resolver (ENR) [7] is introduced to intercept and properly synthesis DNS queries.

NAT64 can be standalone or be collocated with PGW. DNS64 is not required, while, BIH has ENR module inside the host which does similar thing to DNS64.

From the above, BIH/NAT64 is essentially a solution which leverages existing BIA/BIS and combines NAT64.

B.7.3 Service Flow Example

The procedures of an IPv4 application in a BIH enabled host accessing IPv4 services through IPv6 only connection are illustrated in Figure B.7.3.

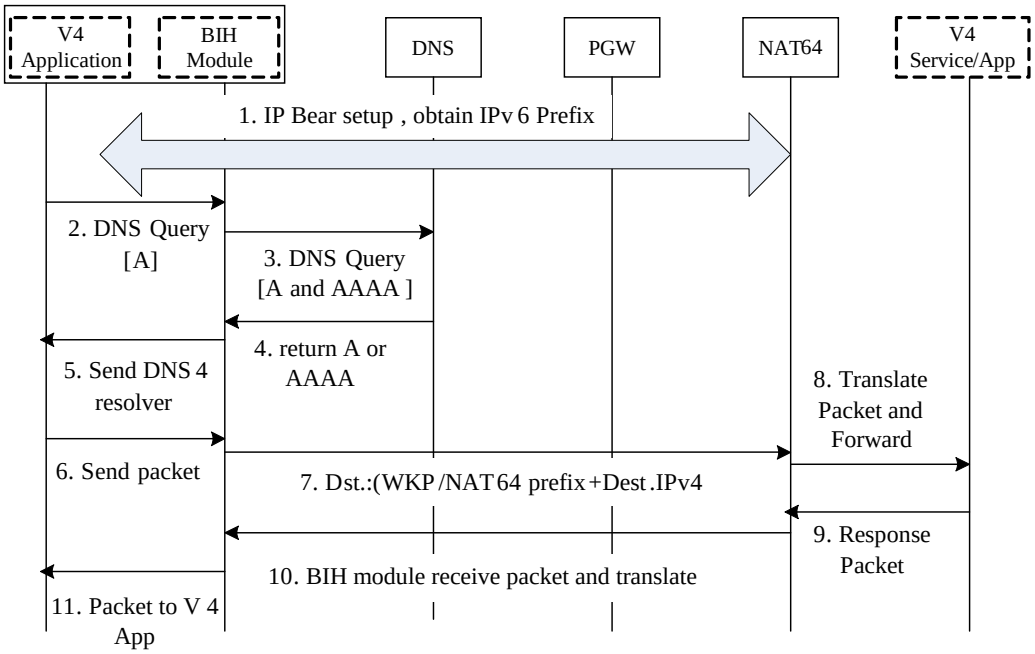

Figure B.7.3: IPv4 application in a BIH enabled host access IPv4 services

The detail information flows are described as below.

1) After bearer activation the UE was assigned an IPv6 address.

2) The IPv4 application in the BIH UE would like to start the communication. The DNS resolver in the UE may send a type A DNS query.

3) The BIH module will intercept that query and converts it to both type A and type AAAA queries and send it the DNS system.

4,5) The DNS system may return a type A or type AAAA DNS query response. BIH module will intercept the DNS response message and may need to convert the type AAAA DNS response to type A DNS response or just return the DNS response to the IPv4 application(in case of received DNS response is type A).

6,7) BIH module translates the application's IPv4 packets into IPv6 packets, BIH module creates the destination address by combine either WKP (Well Known Prefix) or NAT64 prefix together with a 32 bit IPv4 address. The source address will be network assigned IPv6 prefix.

8) The NAT64 receives IPv6 packets. It will translate it into IPv4 packet and sent it to the IPv4 network.

9,10,11) The IPv4 server response with IPv4 packet, the NAT64 gateway translates it into IPv6 and the IPv6 bearer carry the IPv6 packet to the UE. The BIH module translates will then translate it into IPv4 and forward it to the IPv4 application, I.

To let IPv6 applications access IPv4 services, the ENR in a BIH enabled host will compose an AAAA response based on a type A response (e.g., WKP/NAT64 prefix + the IPv4 address of the service). The IPv6 packet will be sent out with the source address being the UE's IPv6 prefix and the destination address being the one generated by ENR. The packet will be routed to the NAT64, translated to IPv4 packet and forwarded to the IPv4 servers/peers.

B.7.4 Evaluation

The Application Level Gateway (ALG) function only needs to invoke in the NAT64 to support applications embedding IP address in the payload.

Since BIH/NAT64 employs NAT64 and enhances BIA/BIS, therefore, the solution shares similar issues and benefits with NAT64.

Known issues of the solution:

- UE is required to install the BIH module, which may increase the complexity of the UE.
- The impact of BIH to the UE's CPU utilization and its OS is FFS.
- The impact to PCC with the deploying BIH/NAT64 is for FFS.
- In case of roaming with Local Breakout (PGW/GGSN in VPLMN), if there is a need to reach IPv4-only services, the VPLMN operator would be required to deploy NAT64 in order to provide the same user experience using IPv6-only connections for IPv4-only services as in non-roaming scenarios.

Known benefits of the solution:

- This solution allows IPv4-only applications running on an IPv6 network to communicate with IPv4, dual stack and IPv6 servers.
- The packet overhead is less compared with DS-Lite plain IPv6 encapsulation solutions.

B.7.5 Applicability

The solution only requires IPv6 only connection and can support IPv4 legacy applications. Therefore, it applies to the following IPv6 migration scenarios outlined in clause 5:

- Scenario 1: Dual-stack connectivity with Limited Public IPv4 Address Pools
- Scenario 2: Dual Stack connectivity with Limited Private IPv4 Address Pools
- Scenario 3: UEs with IPv6-only connection and applications using IPv6

- Scenario 4: IPv4 applications running on a Dual-stack host with an assigned IPv6 prefix and a shared IPv4 address and having to access IPv4 services

Annex C:
Building Block: Dual-Stack EPS Bearer Contexts in EPS/GPRS

C.1 Description

Release 8 specifications TS 23.401 [9], TS 23.060 [11] introduce dual-stack EPS bearer contexts to the EPS and GPRS networks, offering a basic cellular layer feature, which not only enables connectivity to IPv4 and IPv6 PDNs but also simplifies the process of migrating from IPv4 to IPv6 in the network. Dual stack bearer contexts are able to transport native IPv4 and native IPv6 packets within one PDN connection/PDP context. Dual-stack bearer contexts are identified in EPS/GPRS signalling by PDN/PDP type 'v4v6'.

The usage of dual-stack bearer contexts omits the need for opening parallel PDN connections/PDP contexts for different IP address family types. This is an advantage during a phased transition to IPv6 within networks, where PDNs need to support legacy applications using IPv4 whilst other applications have already been upgraded to support IPv6.

From Release 8 onwards, the support for dual-stack bearer contexts is mandatory for E-UTRAN/UTRAN/GERAN terminals, which support both IPv4 and IPv6 addressing.

C.2 Functional Description

It is specified in Release 8 EPS and GPRS specifications TS 23.401 [9], TS 23.060 [11], a Release 8 UE, which has both IPv4 and IPv6 capability, shall always initiate the activation of a PDN connection/PDP context by requesting for a dual-stack (PDN/PDP type v4v6) bearer. The UE is not assumed to have knowledge of the IPv4 and/or IPv6 capabilities of a given PDN. The UE also has no awareness of whether dual-stack bearer contexts are supported by the network to which it is attaching.

The EPS/GPRS network is required to handle requests for dual-stack EPS bearer contexts from the UE and to enforce the type of bearer contexts that are allocated to it. The network may downgrade the request for PDN/PDP type v4v6 for one of the following reasons:

- A given PDN supports/allows only one of the address types i.e. IPv4 or IPv6. This limitation may stem from operator policy.
- All GnGp SGSNs in the operator's network have not been upgraded to support PDP type v4v6. In this case, parallel v4 and v6 bearers contexts to a PDN need to be used instead, so that inter-RAT mobility to/from GnGp SGSNs is possible.

In Release 8, all EPS control plane entities (MME ,S4-SGSN) and user plane entities (SGW, PGW) are able to identify and handle requests to activate a dual-stack bearer context. Dual stack bearer context support for the GPRS core network (GGSN, Gn/Gp SGSN) is specified in Release 9. A pre-Release 9 Gn/Gp SGSN handles PDP type v4v6 as an 'unknown' PDP type, meaning that it handles a request for PDP type v4v6 as if it were a request for PDP type 'v4'. A pre-Release 9 GGSN does not support dual-stack bearer contexts, but dual-stack usage requires the activation of parallel IPv4 and IPv6 bearer contexts to a PDN.

If the UE fails to activate a dual-stack bearer context, and it receives a single-stack IPv4 or IPv6 bearer context, it may attempt to activate a parallel single-stack bearer context for the other IP address type to the same PDN. The Release 8 network may explicitly signal to the UE an error cause that parallel single stack bearers are allowed to the same PDN.

Parallel PDP contexts to a single PDN may also be supported in GPRS networks where PDP type v4v6 is unknown. Therefore, in order to ensure dual-stack connectivity for this case, a UE which first attempted to open a dual-stack bearer context should attempt to open parallel single-stack v4/v6 PDP contexts to the same PDN even without receiving an explicit error cause.

C.3 Information flows

The information flows depicting the activation and mobility of dual-stack bearer contexts are included in specifications TS 23.401 [9], TS 23.060 [11].

C.4 Applicability

In many network deployments, the usage of dual-stack bearer contexts in the network will be the initial method used to begin the transition from IPv4 to IPv6.

The usage of dual-stack bearer contexts has the advantage of offering parallel support of IPv4 and IPv6 addressing within one bearer context. This is a simple solution for end hosts in comparison to handling the activation and mobility of a parallel bearer. Importantly, dual stack bearer contexts offers simplified handling of parallel IPv4 and IPv6 traffic within the network after early EPS deployment phase, when upgraded GPRS core network elements can also be expected to support dual-stack bearer contexts.

The usage of dual-stack bearers during IPv6 transition does not address the shortage or IPv4 addresses, which has been identified as challenge in some IPv6 migration scenarios. However, the usage of dual-stack bearer contexts is an integral part of several IPv6 transition solutions, which also address IPv4 address conservation/re-use. An advantage of using dual-stack bearers within the context of IPv4 address conservation/re-use is that full support for QoS differentiation is already available in Release 8 based UEs.

Annex D:
Change history

Change history						
Date	**TSG #**	**TSG Doc.**	**CR**	**Rev**	**Subject/Comment**	**New**
2011-06	SP-52	-	-	-	MCC Update to version 11.0.0 after TSG SA approval	11.0.0
2014-09	SP-65	-	-	-	Update to Rel-12 version (MCC)	**12.0.0**
2015-12	-	-	-	-	Update to Rel-13 version (MCC)	**13.0.0**
2017-03	-	-	-	-	Update to Rel-14 version (MCC)	**14.0.0**
2017-09	SP-77	SP-170723	0001	1	Update IETF IPv6 related reference of 23.975	14.1.0
2018-06	SP-80	-	-		Update to Rel-15 version (MCC)	**15.0.0**
2020-07	SP-88E	-	-	-	Update to Rel-16 version (MCC)	**16.0.0**
2022-03	SP-95E	-	-	-	Update to Rel-17 version (MCC)	**17.0.0**
2024-03	-	-	-	-	Update to Rel-18 version (MCC)	**18.0.0**
2025-09	-	-	-	-	Update to Rel-19 version (MCC)	**19.0.0**

History

Document history		
V19.0.0	October 2025	Publication