# ETSI TR 103 829 V1.1.1 (2022-08)

TECHNICAL REPORT

# Lawful Interception (LI);
# IP address retention and traceability

Reference

DTR/LI-00204

Keywords

CGNAT, IP address resolution, IP retention, NAT,
PAT, traceability

*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

## *Important notice*

The present document can be downloaded from:
http://www.etsi.org/standards-search

The present document may be made available in electronic versions and/or in print. The content of any electronic and/or print versions of the present document shall not be modified without the prior written authorization of ETSI. In case of any existing or perceived difference in contents between such versions and/or in print, the prevailing version of an ETSI deliverable is the one made publicly available in PDF format at www.etsi.org/deliver.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at https://portal.etsi.org/TB/ETSIDeliverableStatus.aspx

If you find errors in the present document, please send your comment to one of the following services:
https://portal.etsi.org/People/CommiteeSupportStaff.aspx

If you find a security vulnerability in the present document, please report it through our Coordinated Vulnerability Disclosure Program:
https://www.etsi.org/standards/coordinated-vulnerability-disclosure

## *Notice of disclaimer & limitation of liability*

The information provided in the present deliverable is directed solely to professionals who have the appropriate degree of experience to understand and interpret its content in accordance with generally accepted engineering or other professional standard and applicable regulations.
No recommendation as to products and services or vendors is made or should be implied.
No representation or warranty is made that this deliverable is technically accurate or sufficient or conforms to any law and/or governmental rule and/or regulation and further, no representation or warranty is made of merchantability or fitness for any particular purpose or against infringement of intellectual property rights.
In no event shall ETSI be held liable for loss of profits or any other incidental or consequential damages.

Any software contained in this deliverable is provided "AS IS" with no warranties, express or implied, including but not limited to, the warranties of merchantability, fitness for a particular purpose and non-infringement of intellectual property rights and ETSI shall not be held liable in any event for any damages whatsoever (including, without limitation, damages for loss of profits, business interruption, loss of information, or any other pecuniary loss) arising out of or related to the use of or inability to use the software.

## *Copyright Notification*

# Contents

# 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 Web server (https://ipr.etsi.org/).

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™** and **LTE™** 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.

# Foreword

This Technical Report (TR) has been produced by ETSI Technical Committee Lawful Interception (LI).

# 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.

# Executive summary

The present document acts as a guide for policy makers, communication service providers and law enforcement agencies, regarding the retention of IP addresses for law enforcement purposes.

# Introduction

The present document provides information regarding typical IP usage and Network Address Translation within a Communication or Internet Service Providers network. The present document does not attempt to describe all possible implementation variations of NAT but instead focuses on the key underlying principles of IP Communication and typical NAT implementation patterns.

Through understanding these concepts, the reader can gain an appreciation of the impact these techniques have on traffic attribution as observed from outside the Communication Service Providers' private network.

# 1 Scope

The present document considers the following aspects of IP address retention and traceability:

1) Basic Internet Protocol principles, highlighting specifically how IP addresses and ports are used to access the internet.

2) Key differences between IPv4, IPv6, Dual Stack and other relevant layer 3 protocols.

3) How IP addresses are allocated within networks, including EPC and 5GC, documenting any differences in approach.

4) The role of Network and Port address translation within a CSP network.

5) The impact address translation technologies on IP address attribution as observed from outside the CSP network. This includes a discussion on the different translation technologies commonly used by CSPs (e.g. NAT, PAT, CGNAT, NAT64).

6) Description of the key elements which define user and IP address association and therefore make up the minimal set of stored attributes for a viable IP retention solution.

7) Methods for accessing records of IP and port allocation from within a CSP network.

8) Methods for retaining and querying stored IP association records; consideration is given to the storage volumes, durations and accuracy.

The present document does not consider TOR®, VPN services or over top identity protection services, which may impact the ability to attribute observed IP addresses to a specific User Equipment.

# 2 References

## 2.1 Normative references

Normative references are not applicable in the present document.

## 2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] Wikipedia: "IPv4".

NOTE: Available at https://en.wikipedia.org/wiki/IPv4.

[i.2] Wikipedia: "IPv6".

NOTE: Available at https://en.wikipedia.org/wiki/IPv6.

[i.3] IETF RFC 8799: "Limited Domains and Internet Protocols".

NOTE: Available at https://datatracker.ietf.org/doc/html/rfc8799.

[i.4] ETSI TS 103 280: "Lawful Interception (LI); Dictionary for common parameters".

[i.5] Wikipedia: "Regional Internet Registry".

NOTE: Available at https://en.wikipedia.org/wiki/Regional_Internet_registry.

[i.6] Wikipedia: "IPv6-deployment".

NOTE: Available at https://en.wikipedia.org/wiki/IPv6-deployment.

[i.7] ETSI TS 123 501: "5G; System architecture for the 5G System (5GS) (3GPP TS 23.501)".

[i.8] Recommendation ITU-T E.164: "The international public telecommunication numbering plan".

[i.9] IETF RFC 7422: "Deterministic Address Mapping to Reduce Logging in Carrier-Grade NAT Deployments".

NOTE: Available at https://datatracker.ietf.org/doc/html/rfc7422.

# 3 Definition of terms, symbols and abbreviations

## 3.1 Terms

Void.

## 3.2 Symbols

Void.

## 3.3 Abbreviations

For the purposes of the present document, the following abbreviations apply:

| | |
|---|---|
| 4G | 4th Generation Mobile Technology |
| 5G | 5th Generation Mobile Technology |
| 5GC | 5G Core Network |
| AAA | Authentication, Authorization, and Accounting |
| APN | Access Point Name |
| CGNAT | Carrier Grade Network Address Translation |
| CPE | Consumer Premises Equipment |
| CSP | Communication Service Provider |
| DNS | Domain Name Server |
| DHCP | Dynamic Host Configuration Protocol |
| EPC | Evolved Packet Core |
| FTP | File Transport Protocol |
| GB | Gigabyte |
| IANA | Internet Assigned Numbers Authority |
| ICANN | Internet Corporation for Assigned Names and Numbers |
| ID | Identifier |
| IMEI | International Mobile Equipment Identifier |
| IMEISV | International Mobile Subscriber Identity and Software Version |
| IMSI | International Mobile Subscriber Identity |
| IP | Internet Protocol |
| IPv4 | Internet Protocol version 4 |
| IPv6 | Internet Protocol version 6 |
| ISP | Internet Service Provider |
| LEA | Law Enforcement Agency |
| LI | Lawful Intercept |
| LSN | Large Scale Network Address Translation |
| MAC | Media Access Control |
| MSISDN | Mobile Station International Subscriber Directory Number |

| | |
|---|---|
| NAT | Network Address Translation |
| NIC | Network Interface Controller |
| OTT | Over The Top |
| PAT | Port Address Translation |
| PCF | Policy and Charging Function |
| PCRF | Policy and Charging Rules Function |
| PDU | Protocol Data Unit |
| PGw | Packet Gateway |
| RADIUS | Remote Authentication Dial-In User Service |
| RIR | Regional Internet Registry |
| SBI | Service Based Interface |
| SMF | Session Management Function |
| SPAN | Switched Port Analyser |
| SQL | Structured Query Language |
| SSL | Secure Sockets Layer |
| TB | Terabyte |
| TOR | The Onion Router |
| TR | Technical Report |
| UDM | Unified Data Management |
| UE | User Equipment |
| UPF | User Plane Function |
| URL | Uniform Resource Locator |
| UTC | Coordinated Universal Time |
| VPN | Virtual Private Network |
| Wi-Fi | Wireless Fidelity |

# 4 IP address retention and traceability

## 4.1 Basic Internet Protocol principles, highlighting specifically how IP addresses and ports are used to access the internet

**What is an IP?**

The Internet Protocol (IP) is a fundamental building block for communication between entities over a computer network, such as the internet. The Internet Protocol provides a mechanism for passing information between connected end points through the use of IP addresses. Within the context of the internet an IP address can be considered analogous to a postal address, it is owned by a specific entity that can use this address to send and receive information to and from other similar end points that are connected to the same network.

IPv4 addresses are expressed as a set of four numbers separated by the '.' character, an example address might be 192.168.10.1. Each of the four numbers within the IP Address string can range from 0 to 255 allowing the full IP addressing range to go from 0.0.0.0 to 255.255.255.255.

IPv6 addresses consist of eight blocks of 16 bits each. Each group is written as four hexadecimal digits (sometimes called hextets or more formally hexadectets and informally a quibble or quad-nibble) and the groups are separated by colons (:). An example of this representation is 2001:0db8:0000:0000:0000:ff00:0042:8329.

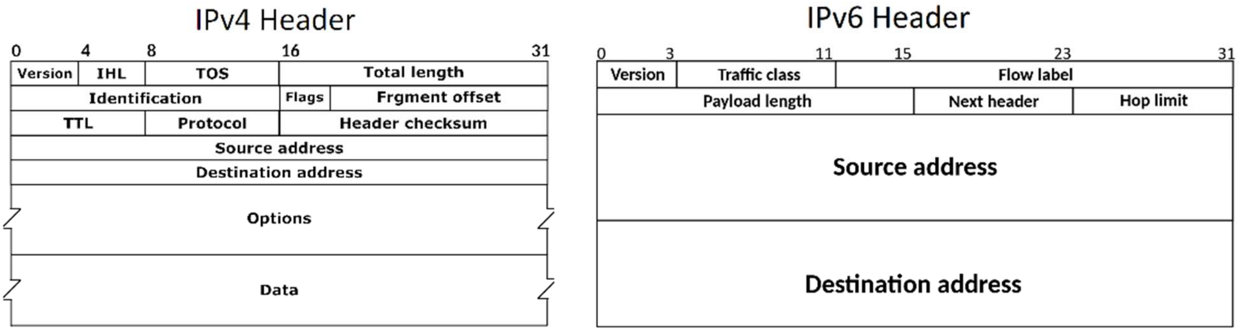

**Figure 1: IPv4 and IPv6 Headers (see [i.1] and [i.2])**

IP addresses are produced and allocated by the Internet Assigned Numbers Authority (IANA), which is part of the not for profit Internet Corporation for Assigned Names and Numbers (ICANN). Each time anyone registers a domain (website address) on the internet, they go through a domain name registrar, who pays a small fee to ICANN to register the domain.

**Are there different types of subscriber IP address?**

All IPv4 or IPv6 addresses are expressed in the same way but there are different categories of IP addresses, and within each category, different types.

**Private IP addresses**

Devices that connect to the internet usually belong to a private network and as such are allocated a private IP address, this includes laptops and mobile devices. Private IP addresses are generated and allocated for each device within the private network allowing them to communicate.

For further details regarding private IP addresses and their used in limited domains see IETF RFC 8799 [i.3].

**Public IP addresses**

A public IP address is an IP address that can be accessed directly over the internet, in the case of a broadband connection this could be the Public IP which is allocated to your router by the ISP. In the case of a mobile network, this is likely to the IP address that will be used to route IP traffic onto the internet after Network Address Translation (NAT).

Public IP addresses come in two forms - dynamic and static.

**Dynamic IP addresses**

These are IP addresses that change automatically and regularly. CSPs buy large contiguous IP address ranges and assign individual IP addresses automatically to their customers or to their customers individual internet sessions. Unallocated addresses are held in a pool to be used for other customers. The rationale for this approach is to generate cost savings for the CSP through better utilization of their available IP addresses space. There are security benefits, too, as a dynamic IP address makes it harder for targeted cyber-attacks.

**Static IP addresses**

Unlike dynamic IP addresses, static addresses are allocated to a specific customer and remain consistent for the lifetime of the agreement. Most individuals and businesses do not need a static IP address, but for businesses that plan to host their own web server, it is typically a requirement to have one. This is because a static IP address ensures hosted services (e.g. websites and email servers) can be reached with a consistent IP address, this is necessary if they want other devices to be able to find them on the internet.

**Subscriber IP address allocation**

Service Provider Subscribers will be allocated either a private IP address which is translated to a public address at the perimeter of Service Providers network or are allocated a Public IP address which is routable and addressable from the public internet, in either cases the IP address can be statically or dynamically allocated.

**Website IP addresses**

If a subscriber is allocated a static public IP address by its Service Provider, then it is possible that this IP address can be used to host (run) a web server which would be addressable through this IP address.

For website owners who do not host their own server and instead rely on a web-hosting provider, which is the case for most websites, there are two types of website IP addresses, shared and dedicated.

**Shared IP addresses**

Websites that rely on shared hosting from a web-hosting provider will typically be one of a number of websites hosted on the same server. This tends to be the case for personal and small business websites, where traffic volumes are low. Websites hosted in this way will share the same IP addresses.

**Dedicated IP addresses**

Some web-hosting providers offer the option to use one or more dedicated IP addresses. This can give more control over things such as SSL certificates and make the support of some services such as FTP possible. A dedicated IP address also allows the website to be reached using the IP address alone rather than the domain name.

## 4.2 Introduction to Network Address Translation

When observing internet activity from outside of a Communication Service Provider's private network, whether this be in real time or through transaction logs generated by web servers, the Source IP and Source Port are key values that identify the origin of the communication session.

While the combination of these identifiers will be unique for a specific data session as observed by the destination host server, due to the implementation of Network Address Translation (NAT) and Port Address Translation (PAT) at the premier edge of the CSP's network, these identifiers may not be unique to a specific connected device within the CSP's private network:

- **Private Source IP** - the private (internal) IP address assigned by the CSP to the user device for communication within their private network.
- **Public Source IP** - the public (external) IP address assigned by the CSP to the user's communications for outside of their network (i.e. the internet).
- **Destination IP** - the external (public) IP address that the user is trying to reach (often resolved by DNS) on the internet.
- **Terminology** - Source and Destination IP addresses are defined as viewed from the User Equipment (UE) point of view.

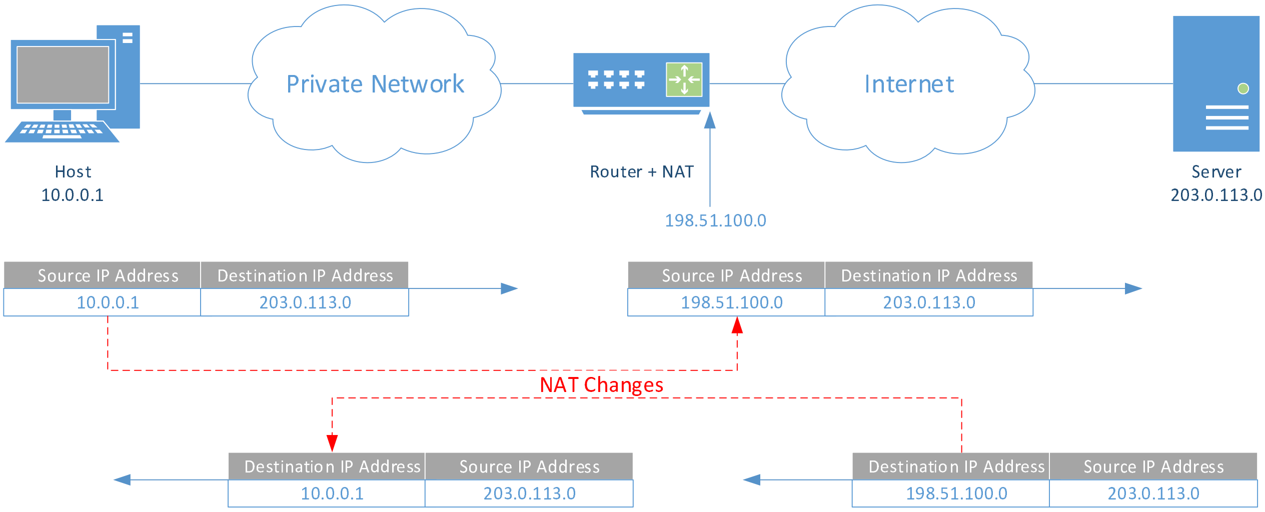

| Source IP Address | Destination IP Address |
|---|---|
| 10.0.0.1 | 203.0.113.0 |

| Source IP Address | Destination IP Address |
|---|---|
| 198.51.100.0 | 203.0.113.0 |

| Destination IP Address | Source IP Address |
|---|---|
| 10.0.0.1 | 203.0.113.0 |

| Destination IP Address | Source IP Address |
|---|---|
| 198.51.100.0 | 203.0.113.0 |

**Figure 2: IP communication with Network Address Translation**

For Law Enforcement Agencies (LEAs) the attribution of a permanent subscriber identifier such as a Mobile Station International Subscriber Directory Number (MSISDN) to observed internet transactions can be a critical part of an investigation and therefore the ability to trace data sessions across a CSP's address NAT or PAT capability is a necessary requirement if attribution is to be achieved.

The IP address to user attribution problem is typically largest in mobile telephony networks, where NAT is universal, although it is also a common problem for public Wi-Fi® networks. However, it is generally not a problem for fixed line or broadband services where each subscriber is typically allocated a static or long held IP address.

The present document focuses on the attribution of observed identifiers, as seen by the external destination, to permanent subscriber identifiers within the CSP's network, such as MSISDN or MAC address, for the use case where IP communication traverses a Carrier Grade NAT (CGNAT).

**Why do CSPs use NAT?**

The implementation of NAT within CSP network is most often through the use of CGNAT technology. CGNAT differs from standard NAT primarily through the introduction of predefined deterministic NAT and the support of large volumes of NAT transactions per second. The primary reason for CSPs to implement a CGNAT is to maximize the use of their available IP public address space as CGNAT allows the same public address to be used by multiple, potentially thousands, of subscribers simultaneously.

A secondary benefit of CGNAT is that it provides a layer of security for CSP subscribers by obscuring the source address, preventing the observed Public IP address from being tracked or attributed from outside the CSP domain.

**How is CGNAT implemented?**

When a device, allocated with a private address, needs to communicate with a publicly addressable entity (e.g. a website), in order for the response to be routed back to the initiating source, the private address needs to be changed for a publicly routable address. This is termed Network Address Translations (NATs) and is normally transparent to both the internal and external hosts.

CSP deployed address translation capabilities are typically implemented at the gateway between the CSP's private network and the public internet. This translation allows the CSP to define their own private address ranges, allocate IP addresses within these ranges to devices connected to the CSP's private network, and support internal traffic routing.

Similarly, to support bidirectional communication, sessions initiated from an internet host towards the UE, the sessions will pass through the CSP NAT device with the public source address being translated into a private source address for the purposes of routing.

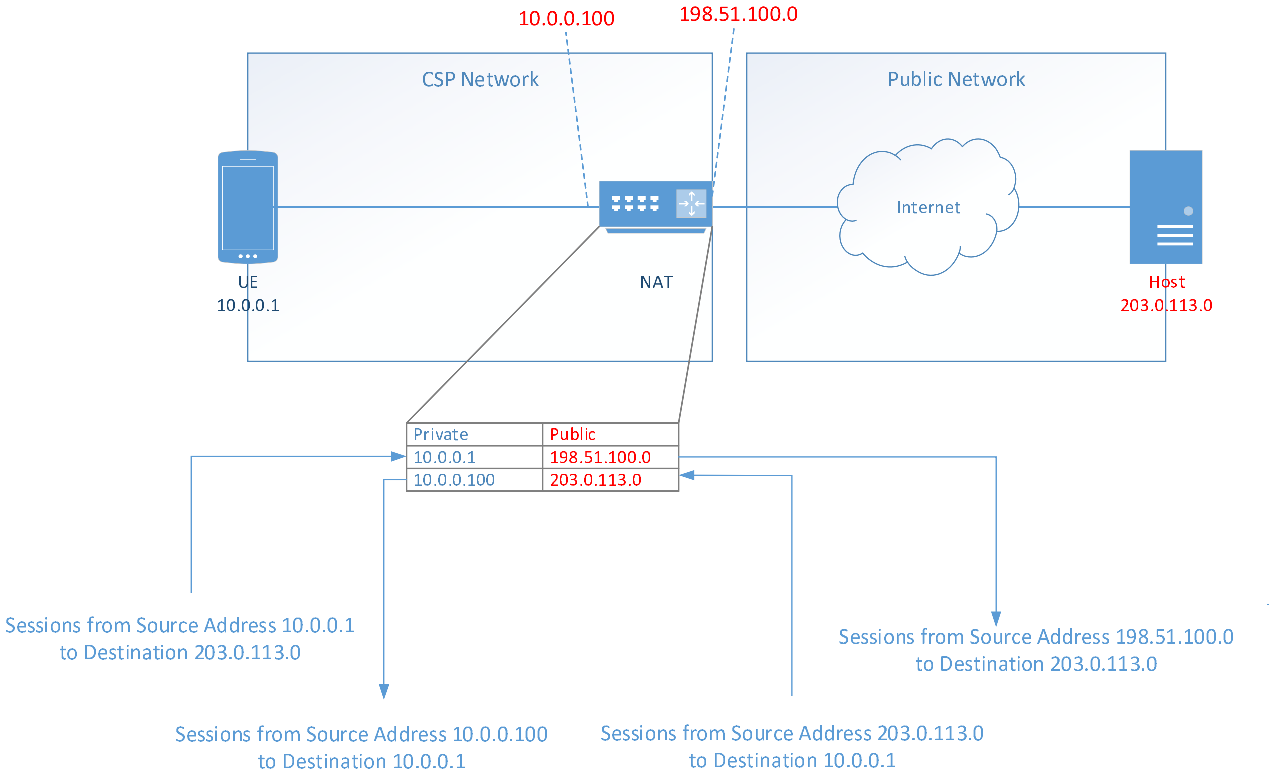
10.0.0.100
198.51.100.0
CSP Network
Public Network
Internet
UE
10.0.0.1
NAT
Host
203.0.113.0
Sessions from Source Address 10.0.0.1
to Destination 203.0.113.0
Sessions from Source Address 198.51.100.0
to Destination 203.0.113.0
Sessions from Source Address 10.0.0.100
to Destination 10.0.0.1
Sessions from Source Address 203.0.113.0
to Destination 10.0.0.1

| Private | Public |
|---|---|
| 10.0.0.1 | 198.51.100.0 |
| 10.0.0.100 | 203.0.113.0 |

**Figure 3: Bidirectional NAT**

**What NAT variants are there?**

NAT can be implemented as:

- **Static** - Internal and external IP addresses have a 1:1 mapping; this is typically used at a CSP where a subscriber has requested or has paid for a static public IP address.

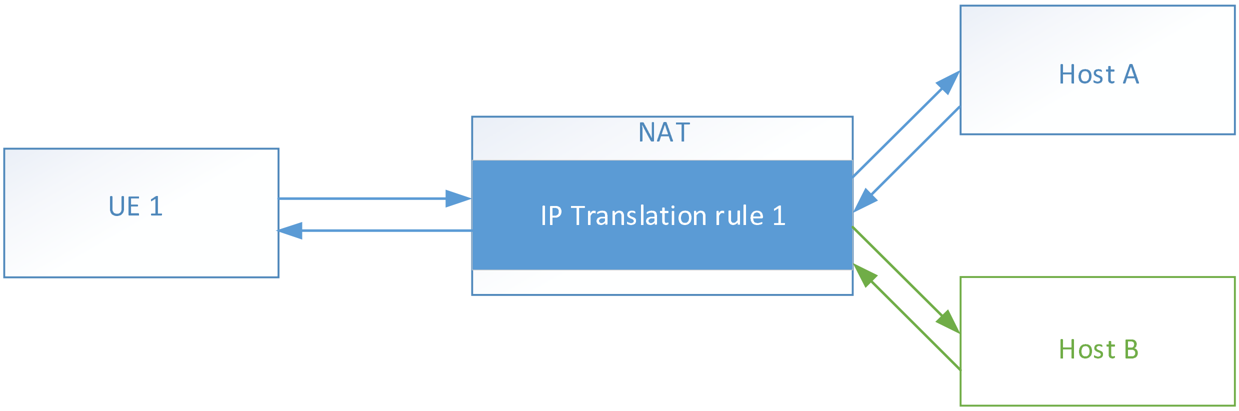

**Figure 4: Static NAT**

- **Dynamic** - Internal IP addresses are dynamically allocated a public IP address from an available public address pool; this is commonly used to maximize the CSP's available public address space.

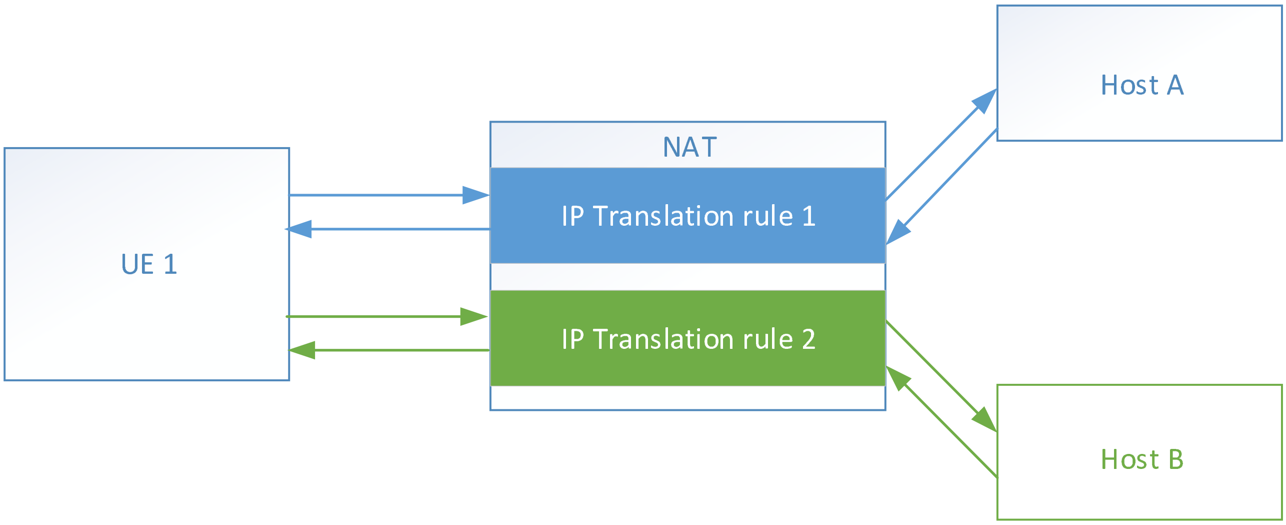

**Figure 5: Dynamic NAT**

- **PAT** - This is the most popular of the three types and is a variant of dynamic NAT, but maps multiple private IP addresses to a single public IP address by making use of source ports to distinguish between individual sessions.

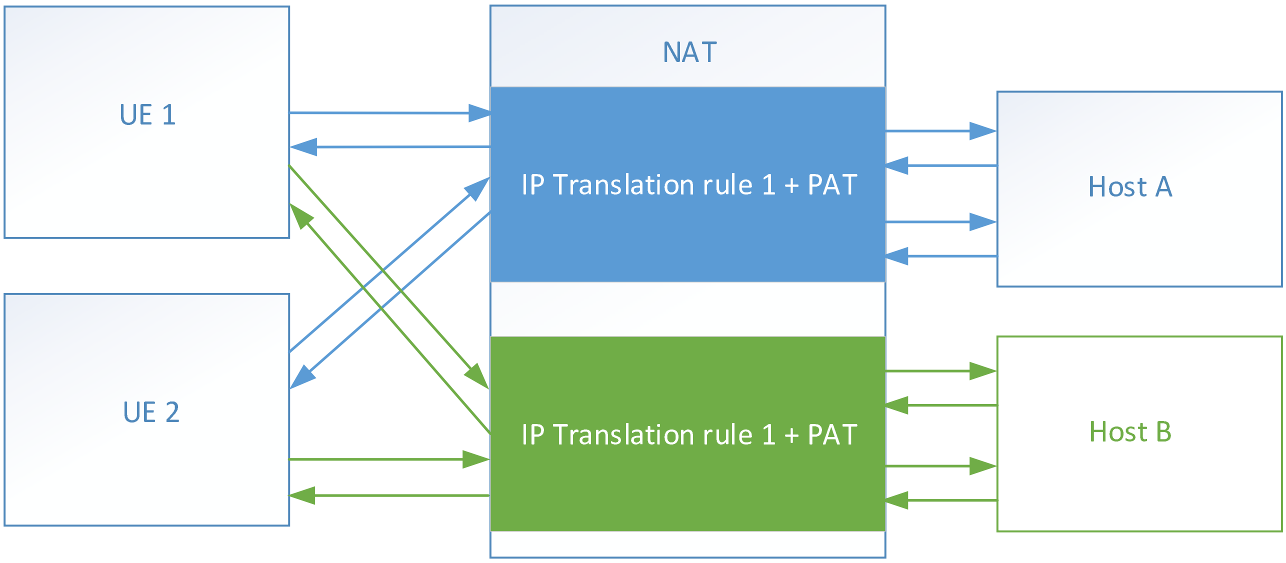

**Figure 6: PAT**

- **Predefined/Deterministic NAT** - In this mode of operation, private IP addresses and source ports are translated using a predefined logic which selects a public IP address and a source port from a subscriber specific range. This avoids the need for full session-based NAT log retention since the public IP address, and source port can be mapped back to a source IP address using the CGNATs deterministic logic, which may be logged or provided as part of a static configuration file.

Note that IP addresses are not assigned to users but to user devices. In some specific cases IP addresses are assigned to subscriptions (e.g. where a fixed public IP address can be provided as part of the service offering of the ISP).

**Why is IP address retention and traceability difficult?**

- IP addresses are often dynamically assigned and only for short periods of time.
- The private IP address to customer identifier mapping is held separately to the private/public IP address mapping, and both sessions are managed independently.
- Correlation of the records requires high accuracy and precise timestamps.

- Determining session end records requires an understanding of session duration.
- Creating a complete picture involves joining sources of data from multiple systems (typically AAA and NAT logs).
- Dynamic IP address allocations are not typically retained for business purposes (transient).
- The capture and processing of large volumes of NAT data can be storage and computing power intensive.

**How do over the top identifier protection mechanisms affect IP attribution?**

While the use of an Over The Top (OTT) service such as a VPN, TOR® or Private Relay™ will not affect the private and public IP address allocations made by the CSP, the use of these type of OTT-services is likely to limit the ability of law enforcement to identify the server or service accessed, or limit law enforcements ability to identify the source of the observed traffic. Further details of the use of OTT identity protection services are out of scope of the present document.

## 4.3 How IP addresses are allocated within networks, including EPC and 5GC, documenting any differences in approach

Within a fixed line network, static or temporary private or public IP addresses are allocated to Customer Premises Equipment (CPE), such as modems, by the ISP as part of the authentication process using services such as DHCP (Dynamic Host Configuration Protocol). Allocated temporary IP addresses will remain assigned to the CPE until either the device is restarted, or the validity period as defined by the ISP expires. At this point, the ISP's IP address allocation service issues the CPE with a new temporary IP address.

Within an EPC (4G) or 5GC (5G Core) mobile network, temporary private IP addresses are allocated to User Equipment (UE) as part of the respective bearer establishment or PDU session establishment procedures. When UEs are in idle mode, in 5G Suspend or have no active data sessions, it is possible for no IP address to be allocated. The allocated IP addresses can be IPv4, IPv6 or a combination of IPv4 and IPv6 and remain with the UE as long as the UE's bearer or session remain active, although depending on network configuration, timeouts can be implemented which will force a new internal IP address to be allocated.

In 4G, UE IP address allocation is performed by the PGw (Packet Gateway) function in conjunction with an internal or external DHCP server. For 5G it is the SMF (Session Management Function) which acts as the DHCP server towards the UE although address allocation can be performed by the UPF (User Plane Function) via the N4 interface with support from the SMF, or through an external DHCP, again with support from the SMF [i.7].

In both the EPC and the 5GC dynamic IP allocation models, IP addresses are selected by the DHCP sever based on the UE's requested session or bearer type (i.e. IPv4), the UE's applicable IP address pool and address availability within the pool.

Where the UE has been allocated a static IP address the same EPC and 5GC functions are used to deliver this IP address to the UE but rather than this IP being sourced from an address pool it is provided through configuration on functions such as the DHCP sever or 5GC UDM (Unified Data Management).

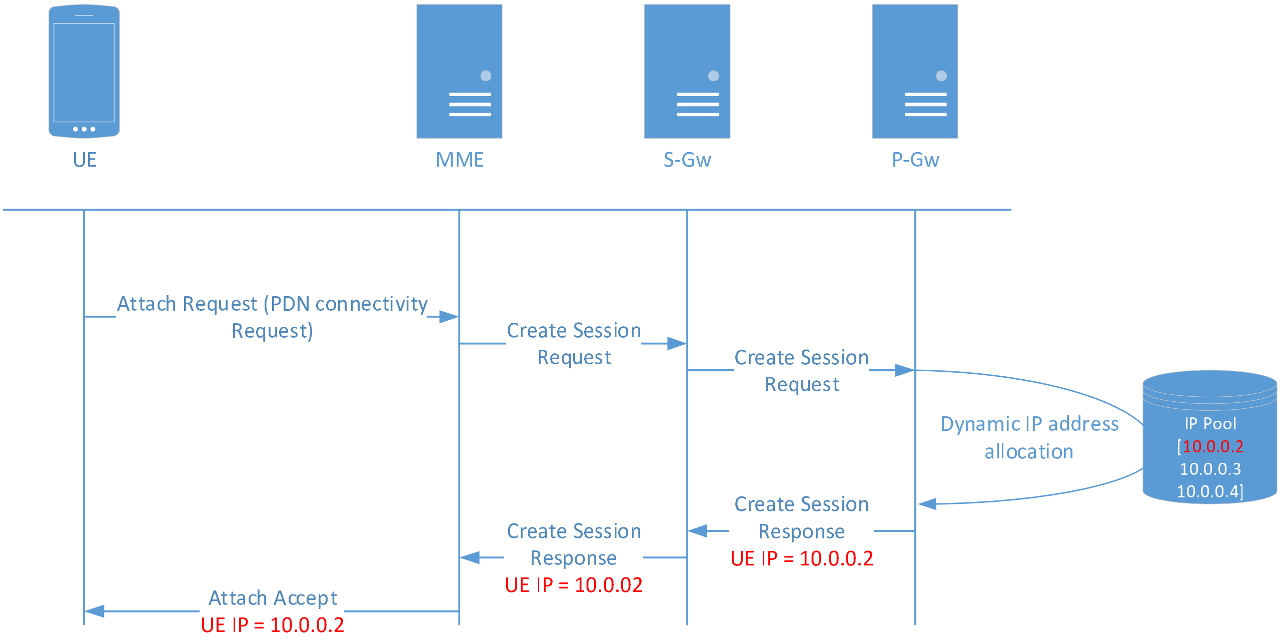

Figure 7: IP address allocation in an EPC

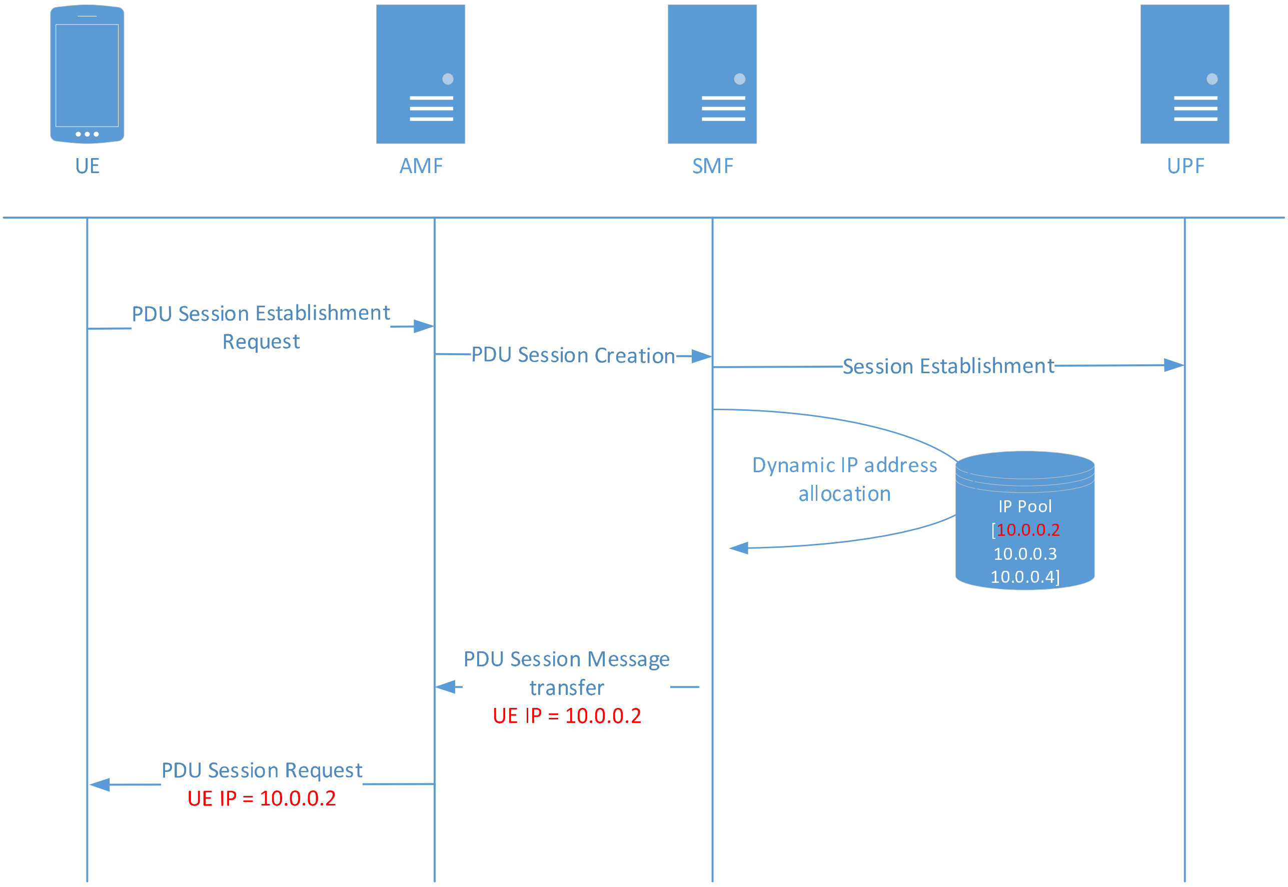

Figure 8: IP address allocation with 5GC

### How is a private IP address linked to a UE?

In a fixed line network, the IP address allocation service, such as the DHCP service, will maintain a record of the association between the unique CPE identifier (e.g. MAC address) and the allocated public or private IP address. While only the most recent association is required for the purposes of service operation, it is possible that historic association information is retained by the ISP for a limited time period for business purposes.

In the EPC the link between permanent subscriber identifiers (e.g. MSISDN) and the private IP address is maintained by a PGw for the lifetime of the data session, during which the user may connect to the internet with one or more internet sessions. This data session will have a maximum time limit, e.g. 24 hours after which the private IP address allocation will be renewed, although many sessions can be shorter. In the event of a time-out where a subscriber has an active internet session the data session is terminated, and a new data session is automatically started, which is typically transparent to the user.

To support functions such as policy and charging, AAA accounting messages are typically made available to downstream consumer functions such as the EPC Policy and Charging Rules Function (PCRF).

These accounting records provide session details such as:

- Mobile Station International Subscriber Directory Number (MSISDN).
- International Mobile Subscriber Identity (IMSI).
- International Mobile Equipment Identifier (IMEI).
- Access Point Name (APN).
- Private IP Address.
- Session Start.
- Session End.
- Bytes Up.
- Bytes Down.

In the 5GC the SMF is the primary source of subscriber to IP allocation information, unlike the EPC with uses protocols such as RADIUS or DIAMETER for communication with the PCRF, the SMF needs to inform the 5GC PCF about allocated and de-allocated IP addresses over the 5GC Service Based Interface (SBI).

**How are public IP addresses allocated?**

Within a CSP the CGNAT functions will reside at the perimeter of the network and allocates public IPv4 or IPv6 addresses to internet bound data sessions.

The CGNAT acts as a transparent proxy for a UE's data session, replacing the CSP allocated private IP address with a public IP address taken from the available pool of CSP public addresses. Each UE internet browsing session or application transaction will consist of one or multiple data sessions, which, depending on the CGNAT configuration, may each be allocated a different public IP address. In practice, this means that as a minimum, one CGNAT session will be created for each unique URL browsed, and some websites or applications may require tens or even hundreds of unique connections to be made to display all content or provide the requested service.

The CGNAT manages the allocation of public IP addresses and their association to private IP addresses. Depending on the configuration used it may be necessary for the CGNAT to store in memory or log the public, private and destination IP addresses, although the CGNAT at no point is aware of the permanent identifiers associated with the UE.

In a standard architecture model, there are no control messages exchanged between the core network functions (e.g. PGw, or SMF) and a perimeter CGNAT. This results in different session termination timings for internal (user to CGNAT) sessions and external (CGNAT to the internet server) sessions.

For example, when a UE data session ends, e.g. through the closing of a web browser on a mobile device, some or all of the subscriber's data sessions may be terminated on the anchoring network gateway. Since there is no direct feedback to the CGNAT device, the established connections between the CGNAT and the internet web server may continue until a time out limit is reached, (e.g. configuration on the CGNAT for active session timeout has been reached).

**Is CGNAT still required with IPv6 address?**

Despite of the introduction of IPv6, which in itself supports the requirement for each connected device to receive a unique public IP address, CGNAT remains necessary for CSPs allocating IPv6 addresses to their UE in order to provide support for internet servers and services which are still running the legacy IPv4 protocol.

According to a report by Google® [i.6], the IPv6 adoption rate is 30 % - 35 % as of April 2021, subsequently CGNAT can be expected to be present within CSP network for some years to come from the date of publication of the present document.

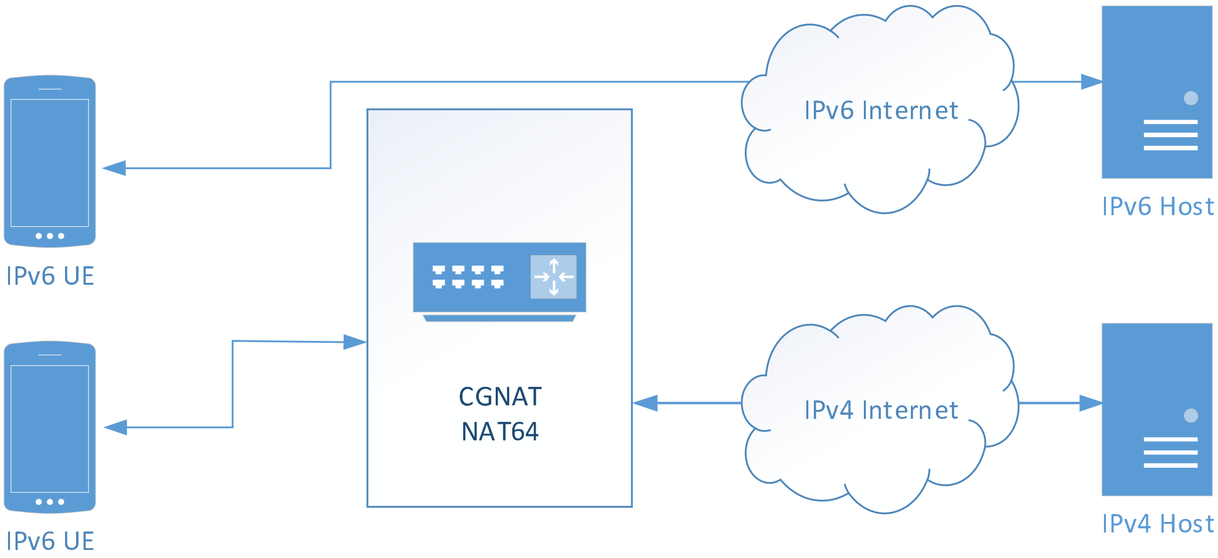

**Figure 9: NAT64 for IPv6 UEs**

**What is the role of source port allocation within CGNAT?**

In addition to the public IP Address allocation, CGNAT may also allocate a public source port. This can be used to establish sessions which share the same unique source IP address; this is also known as Port Address Translation (PAT).

## 4.5 Location from public IP addresses

While public IP addresses themselves do not relate to a specific location, they can be associated with a geolocation through querying of an IP-geolocation database. It should be noted that there is no guarantee on accuracy for location information associated with an IP address when provided by an IP-geolocation database.

Geolocation IP databases attempt to provide a mechanism to determine the approximate location of a UE. The database will not necessarily provide the exact location of the UE but will provide an estimation of the location based on a variety of data sources.

These geolocation databases are created and maintained by third parties, and thus the location data of IP addresses is only as good as the company that manages the data. Owners of these databases acquire their data through a variety of mechanisms including access to Regional Internet Registry [i.5] data sets, ISP provided data sets, data mining of user submitted geographic location data (e.g. weather websites) and publicly available network routing information.

**How accurate are Geolocation Databases?**

Since the issuance of public IP addresses to ISPs is managed by five Regional Internet Registries, IP geolocation data is typically accurate at country level although even this cannot be guaranteed. At the city level, that accuracy sharply declines especially if the user is not located in a major metropolitan area. For example, an IP address used in a suburb, might be geo-located to the centre of the nearest major metropolitan area, although the wider region and country will possibly be accurate.

## 4.6 The role of Network and Port Address Translation within a CSP's network

Network and Port Address Translation serve three primary purposes within a CSP's network:

1) Network security.

2) IP address sharing (with limited number of available addresses).

3) IP version translation.

**Network security** - CGNAT provides network security by hiding the details of the internal network. NAT hides subscribers IP addresses by making it difficult to target a specific individual through IP address allocation alone.

**IP address sharing** - Each CSP will own or have access to a limited range of internet facing IP addresses. It is typical that the number of connected devices will exceed the number of available IP addresses. With readily available IPv4 addresses exhausted, it is increasingly difficult for CSPs to significantly expand their IP address allocation. CGNAT allows a CSP to enable many UEs to share the same public IP address and with the introduction of PAT, these IP addresses can be shared simultaneously.

**IP version translation -** CGNAT can provide a capability to translate IPv6 addresses to public IPv4 addresses. This is required where an IPv6 public source address is being used to communicate with a web server that only supports the older IPv4 address format.

**Are there different types of CGNAT?**

CGNAT or Large-Scale NAT LSN can be implemented in a variety of ways in support of different business requirements. The common implementations at a CSP are:

- NAT44 - This describes the situation where a CSP's private IPv4 address is translated to an IPv4 address from the CSPs public address pool.
- NAT444 (Private-Private-Public) - This describes the situation where an enterprise customer's traffic is translated from its private IP address to the CSP's private IPv4 address and then finally translated again to an IPv4 address from the CSPs Public address pool.
- NAT64 - This allows UEs with only an IPv6 address to access legacy IPv4-only services by mapping the IPv6 address to a public IPv4 source address.
- NAT66 - This is used to translate a private IPv6 addresses into public IPv6 addresses at the edge of the network.

## 4.7 The impact of address translation technologies on IP address attribution as observed from outside the CSP's network

When observing traffic generated by a UE that is passing through a CGNAT, the source IP and possibly the source port are allocated by the CGNAT.

Different NAT technologies impact the behaviour of the CGNAT, table 1 defines the possible translations.

**Table 1: Behaviour of different NAT Technologies**

| Translation Type | Private Source IP | Private Source Port | Public Source IP | Public Source Port |
|---|---|---|---|---|
| NAT44 | CSP allocated IPv4 | UE allocated | CGNAT allocated IPv4 | UE allocated<br>CGNAT allocated when in conjunction with PAT |
| PAT | CSP allocated IPv4 | UE allocated | CGNAT allocated IPv4 | CGNAT allocated |
| NAT64 | CSP allocated IPv6 | UE allocated | CGNAT allocated IPv4 | UE allocated<br>CGNAT allocated when in conjunction with PAT |
| NAT46 | CSP allocated IPv4 | UE allocated | CGNAT allocated IPv6 | UE allocated<br>CGNAT allocated when in conjunction with PAT |
| NAT66 | CSP allocated IPv6 | UE allocated | CGNAT allocated IPV6 | UE allocated<br>CGNAT allocated when in conjunction with PAT |

**Are there different IP address and port allocation algorithms?**

In its basic form, NAT44, IP addresses are allocated in a 1-1 ratio, where the public IP address is randomly allocated from the next available IP address within the configured address pool. Address pools can themselves be configured to support specific ranges and specific source addresses can be configured to receive their public addresses from one or more specific address pools.

Port Address Translation can be considered as a subset of NAT with IP allocation as per NAT44 but with the source port also being allocated. As with IP address allocations, available source ports are put into pools with rules created to define which source IP addresses can use which source ports.

**What is a Deterministic NAT?**

Deterministic NAT or Fixed NAT is a CGNAT feature that allocates source ports for each UE from a predefined (fixed) set of ports for the specific source IP address. Where UE devices use Fixed-NAT to receive a deterministic set of ports, UEs can be identified based on knowledge of the deterministic logic, deterministic parameters and Public IP address allocation, this removes the need for session based CGNAT logging, which in turn reducing the logging overhead on CGNAT devices.

Where deterministic NAT is used, it is necessary that the CGNAT device makes available, via logs or CGNAT configuration dumps, the Public to Private IP address allocation and the logic used for Source Port range allocation. Further details of the CGNAT logging requirements for Deterministic port allocation can be found in IETF RFC 7422 [i.9].

When using deterministic NAT, the number of configured deterministic blocks (public IP and source port allocations) has to be greater or equal to the number of available private IP addresses, this constraint may introduce some limitations when implemented in networks with a large number of subscribers.

## 4.8 Description of the key elements which define user and IP address association and therefore make up the minimal set of stored attributes for a viable IP retention solution

Two key data feeds are required to resolve a subscriber from an IP Address: AAA and CGNAT logs.

**AAA data feed**

To browse the internet, a customer/subscriber needs to be both authenticated, and assigned an IP address for use within a session. The AAA service typically logs session allocations and the details from those logs can be used in searches.

In the example where AAA is provided by a RADIUS service, RADIUS accounting logs would normally contain the following fields:

- Event date and time(s).
- Subscriber/Device identifiers (e.g. MSISDN, IMSI, IMEISV, MAC).
- Private Source IP Address (Internal IP Address).

AAA mechanisms create sessions with a defined start and end point, during which time CGNAT events for a given private source IP address can be connected to a given subscriber or device.

Where the AAA logging is configured to provide a single record session, these are typically supplied at the end of a session and include the start and end date and times. In this scenario, the accounting would not be available for query until the session had ended, and a session could last 24 hours, several days or even longer, depending on the network configuration.

If AAA logs were supplied at the start of a session, the end point of that session would not be known for certain, although it could be known that it did not continue past the start of the next session with the same IP address (and possibly port). Accounting logs delivered in this way would mean that CGNAT records cannot be guaranteed to have occurred within that session, particularly if there was no subsequent accounting log for the same IP.

Alternatively, information on a session can be provided for various session-based events as and when they happen, for example:

- Session start.
- Intermediate point.
- Session end.

These are usually connected to a unique identifier, e.g. session ID, but these IDs may be reused over time, so care should be taken over their uniqueness.

This accounting approach, while richer in information, should be implemented with care to ensure that robust logic is used to deal with edge case scenarios or missing records as this can lead to a distorted view of session states.

The data in each type of record may vary by set-up, e.g. the IMEI may be supplied on only one event type by default.

To understand a complete session, the date and time on the session start record needs to be combined with the date and time on the session end record.

The intermediate records can be used for two purposes: to act as an extension past a maximum time limit without a record, or to make it more certain that the correct session is being used when either a Start or Stop record is missing.

**Table 2: AAA Accounting records**

| Data Element | Further information |
|---|---|
| Subscriber ID(s) | A unique identifier(s)belonging to the subscriber, MSISDN, Account ID, IMEI, IMSI |
| Source IP Address (including port) | The internal/private IP address assigned to the subscriber for the RADIUS session |
| Timestamp | The timestamp of the event |
| Event Type | Any session event information (start, stop, interim) |

**CGNAT data feed**

NAT logs provide the link between source IP address and port (Internal) and source IP address and port (External).

The logs will typically contain information about the data exchange that has taken place, including the following fields:

- Event date/time(s) - e.g. date time at start/stop.
- Duration.
- Private source IP address.
- Private source port.
- Public source IP address.
- Public source port.
- Destination IP address.
- Destination port.
- Protocol.
- Bytes uploaded.
- Bytes downloaded.

CGNAT or network probe generated logs can either be event-based or session-based, depending on the network configuration, but as a NAT session is often brief, logging can be withheld until the end of a session, and information about that session is aggregated into one record. However, a CGNAT event will always start within an AAA session (where this network combination is used), but not always finish within one, especially if the AAA session terminates and the CGNAT event is not finished - it will then timeout after a given period of time, which will fall outside the AAA session's boundaries.

End records containing aggregations of the session, including duration and bytes, will require the source CGNAT service to hold data during each session. This can require significant resources on the CGNAT. However, without it, either the session would need to be treated as a point-in-time event, or records would need to be generated at the start and end and joined as per clause 4.10. With the former approach of point-in-time events, as CGNAT events can end after a AAA session has completed, only start sessions should be used. NAT session end times would not be available with this approach.

**Table 3: CGNAT event record**

| Data Element | Further information |
|---|---|
| Internal IP address (and port) | The internal/private IP address assigned to the subscriber for the lifetime of the session as allocated within the AAA |
| External IP address (and port) | The external/public IP address and port assigned to the data session by the CGNAT |
| Timestamp | The timestamp of the Event |

The join between the AAA and CGNAT data sets is made through the only field present in both data sets: source IP address and port (internal).

## 4.9 Methods for accessing records of IP and port allocation from within a CSP's network

Where deterministic NAT is used it eliminates the need for full session based CGNAT IP address and port allocation logging. This is because in this mode the subscriber's IP address is always mapped to the same public IP address and port range. Through access to the allocation algorithm or table, it is possible to identify from the observed public IP address and source port the corresponding private IP addressed.

Where deterministic NAT is not used, there are two primary methods of data acquisition, which apply to both of the AAA accounting and CGNAT logs discussed in the clause above. Data acquisition approaches can be defined as active, which requires close integration with network functions such as CGNAT devices or passive, which does not require network function integration but instead acquires data through passive monitoring of network links.

Both approaches are subject to national requirements and agreements that these records can be accessed and stored.

**Passive data acquisition**

Since the use of active logging on a CGNAT device has the potential to affect customer services, through the introduction of additional workloads within the CGNAT device, mitigations such as additional CGNAT compute resources may be required to ensure consistent and accurate log generation.

Passive solutions, which do not rely on CGNAT device generated logs, have the potential to avoid these risks, but introduce different complexities around the accuracy of the data, an enduring requirement to track and respond to network changes and capacity expansion requirements in response to growth in data rates across the interfaces.

Passive solutions rely on the ability to passively acquire all network traffic that ingresses and egresses a CGNAT device. This can be done through the use of switch technologies such as Mirror and SPAN or through the use of physical devices such as optical splitters.

Once acquired, data is then processed to determine a correlation between the ingress and egress IP packets with the end purpose being the ability to trace all IP sessions through the CGNAT. There are a number of ways to achieve this. Although the evaluation of the IP payload is a typical method, fundamentally any algorithm that is implemented needs to generate accurate correlations across the pre and post CGNAT packet streams, generate synthetic CGNAT logs and then combine this with accounting logs to identify the UE responsible for each connection/session to which NAT was applied.

There are two established ways in which the UE's private IP address can be passively identified – either by placing the pre CGNAT passive acquisition device at an appropriate location where the MSISDN or similar UE identifiers are available, or by taking an AAA feed (e.g. RADIUS or DIAMETER) and correlating the users private IP address against the sessions. Both of these methods are complimentary and may be required in conjunction to identify UE associated with an IP session.

In some cases, passive acquisition devices such as probes will be able to maintain a state table of AAA (subscriber to internal IP address) mappings and combine this with a suitable CGNAT feed to perform the join in real-time, outputting a single simple data feed making the ingest storage and query side simpler.

**Log acquisition - Active acquisition**

Data can be acquired directly from the network devices that provide the access and translation for user sessions. Data is typically sent as either a data stream or as batched log files.

Where the CSP's network is already configured to provide CGNAT logs and AAA logs for business purposes, new IP address resolution collection systems can often subscribe to these existing data feeds to receive the logs.

**Network coverage**

It should be noted that the potentially distributed nature of CGNAT, which can occur at each of a CSP's Internet Access Point, means that all of the above acquisition approaches will need to be deployed at each of the CSP's Points or Presence where address translation takes place.

**Data standardization**

In order for data to be searched effectively, some form of standardization should be implemented during the data acquisition and processing stage. That can be provided through transformation of several known format variations into a single standardized format, as well as validation to spot records that are of an unknown format that cannot be transformed using the existing logic.

Field validation and transformation is important for two main reasons:

1) So that records can be searched for in a known format.

2) So that records can be ingested into the format of the data store.

Should records be stored as provided in the source format, and that format does not match the format of the data store, the record could be rejected.

Should the format not match the query format, data would not be returned in searches.

**Dates and Times**

One of the most commonly standardized fields within acquired data sets are the date and time fields, which can be affected by time zones and daylight savings.

On ingest, dates and times should be standardized. When a query is made, it is known what times are being compared. The most common standardization is UTC, since it removes ambiguities caused by clock-changes. However, a local time could be used as long as it can be trusted to have altered accurately along with any national daylight-saving time-changes.

To standardize times, a known time zone of the source data is required. This can either be supplied or implied. Using local time is often problematic as this assumes the local time adjusts accurately, when required.

On query, either a time zone field should be supplied, or it should be known and fixed to what time zone the user will be supplying. Local time is reasonable, although for one hour per year there is a repeating hour leading to an ambiguous input. The approach to this should be agreed as part of the solution design.

When results are returned, the formatting approach will often require tailoring to the specific working practices of the CSP. For example, dates and times could include a time zone, they could have different time zones present in separate columns (e.g. local time of the event/customer, local time of team running the query, UTC, etc.), or they could have a single time zone presented in one column (often local time).

Additionally, time synchronization should be considered - where two or more systems are supplying a time, they could differ unless they are synchronized to the same source. This is important, in the example when an AAA session starts and a CGNAT event happens immediately - if the CGNAT clock is recording an earlier time than the AAA system time, it may appear that the CGNAT event occurred before the AAA session started, and therefore there would be no match to a UE device.

**IP Address**

When storing an IP address within a data store there are several possible representation options. When used to support querying an IP address for an exact match, the use if these representation needs to be standardized within a solution.

For IPv4, the two common stored formats for an IP address are:

- Each octet is supplied as an integer, e.g. `10.2.3.4.`
- Each octet is zero padded to three digits, e.g. `010.002.003.004.`

Combinations of these formats across octets is also valid, but less frequently used.

For IPv6 addresses, there are three common formats:

- Each octet is written in full format, zero padded:E.g. fc00:0000:0000:0000:0000:0100:ffff:ffff.

Each octet is written without leading zeros:

- E.g. fc00:0:0:0:0:100:ffff:ffff.
- Fully zero octets are stripped and replaced with "::" where there are two or more missing octets (this can only occur once in an address).
- E.g. fc00::0100:ffff:ffff.

Combinations of these formats are also valid, but less frequently used.

As long as a standard approach is used and documented, and any value entered into the query input fields is standardized to the same approach, it does not necessarily matter which standard is selected.

It would also be up to the CSP whether the original, unmodified IP address would also need to be stored and/or returned, or whether the modified value would be sufficient.

To search or filter on an IP address using a range, see clause 4.10.

**MSISDN**

A subscriber's phone number within a mobile CSP is called MSISDN. An MSISDN takes the form:

**MSISDN** = CC + NDC + SN, where:

CC = Country Code.

NDC = National Destination Code, identifies one or part of a PLMN.

SN = Subscriber Number.

The max length of an MSISDN is recommended to be 15 digits.

The CC part of a MSISDN should be standardized, with agreements made on whether the full international dialling format [i.8] be used in preference to the shortened MSISDN.

**IMEI/IMEISV**

An IMEI, which uniquely identifies a mobile device, can be supplied in three possible formats:

- A 14-digit IMEI, which uniquely identifies the handset.
- A 15-digit IMEI, which adds a check-digit to the above 14-digit IMEI; this is fixed and calculable using the previous 14 digits.

- A 16-digit IMEISV, which adds two digits, which relate to the software version of the handset and that can vary over time.

A search for a specific mobile device involves a search on IMEI.

To uniquely identify a handset, only the first 14 digits of the IMEI are required. However, the full source value should usually be returned, whichever is supplied.

Therefore, the most common approach is to store in an indexed field over the first 14 digits of the IMEI field, and only the first 14 digits entered in the query input field should be used for a query over this field. Any matching IMEI over those first 14 digits should be returned to the query as the original, unmodified value.

**MAC Address**

A Media Access Control (MAC) address is a unique identifier assigned to a Network Interface Controller (NIC) for use as a network address in communications within a network segment.

As typically represented, MAC addresses are recognizable as six groups of two hexadecimal digits, separated by hyphens, colons, or without a separator.

MAC addresses are primarily assigned by device manufacturers and are therefore often referred to as the hardware address, or physical address. Each address can be stored in hardware, such as the card's read-only memory, or by a firmware mechanism. Many network interfaces, however, support the changing of their MAC address.

Network nodes with multiple network interfaces, such as modems, routers and multilayer switches, need to have a unique MAC address for each NIC in the same network. However, two NICs connected to two different networks can share the same MAC address.

For further details of common parameters including additional information on those used and described in the present document see ETSI TS 103 280 [i.4].

## 4.10 Methods for retaining and querying stored IP association records

Consideration is to be given to the storage volumes, durations and accuracy requirements.

**Data Storage**

As noted above, CGNAT and AAA data acquisition is distributed across the CSP network at the various points of presence. There are two approaches to data storage generally taken:

- Separate data stores locally at each capture point.
- Centralized data storage.

For the centralized data storage approach, records acquired from the network are backhauled to a central data store. This can require sizeable network bandwidth to transfer the data records, even if compressed prior to transport.

In the distributed architecture, some care is required when querying data sets. Additionally, this architecture can be more complex when considering system resilience and redundancy.

**Distributed vs Centralized Storage**

The large volumes of data required for IP address retention and traceability mean that due consideration is necessary for transport and storage of this data in any system design.

It is likely that a distributed architecture suits large CSPs with many internet access points/nodes since it avoids transporting large volumes of data across the network, and as sites can generally be treated separately there is normally no data constraint making centralization a requirement.

A centralized architecture may suit smaller CSPs as is likely to cost less in terms of compute resource and storage allocation because extra capacity for resilience can be shared centrally. This does however have a large impact on a single data centre and relies on the CSP network architecture allowing for the transport at these volumes.

**Indicative Volumes**

Any such system will have large ingest and flow rates (records per second). To help to provide an indication of numbers, a system retaining AAA in the form or RADIUS and actively generated CGNAT logs will require roughly 25 TB of storage per 1 million customers for a 12 month retention period for a single copy of the data.

**Indicative volumes per million customers per day**

- Raw data: 80 - 100 GB.
- Record count: 650 - 750 million.

**Searching Data**

The underlying principles of searching IP address retention data is that there is an initial query on the input field(s), and then data is joined with other data stores to enrich it and return the required output fields.

However, IP address retention data is not often event-based data, which would involve records relating to point-in-time events. Instead, IP address retention data is usually session-based, which means a session has a start and an end and is active between these points.

Records relating to sessions will usually involve either a single record describing the entire session, or records at the start and end and optionally intermediate keep-alive records at regular intervals. More information on how to create sessions from these records can be found in clause 4.8.

In either case there needs to be a method of identifying sessions active at a point in time from records dated before and/or after that timestamp.

Where data needs to be joined, the same principle then needs to be applied to the data it joins with, and ensuring that only sessions from both sources partly or entirely active within the query window are returned.

Figure 10 shows how data should be queried in a store. In this example, it is assumed that there are two session-based data sources that need to exist together for the data to be returned. An example is RADIUS data (green) and CGNAT data (orange). It is assumed that the end points of each line represent a record having been received, or that the start point can be determined from the full information generated by the end record (e.g. end minus duration).

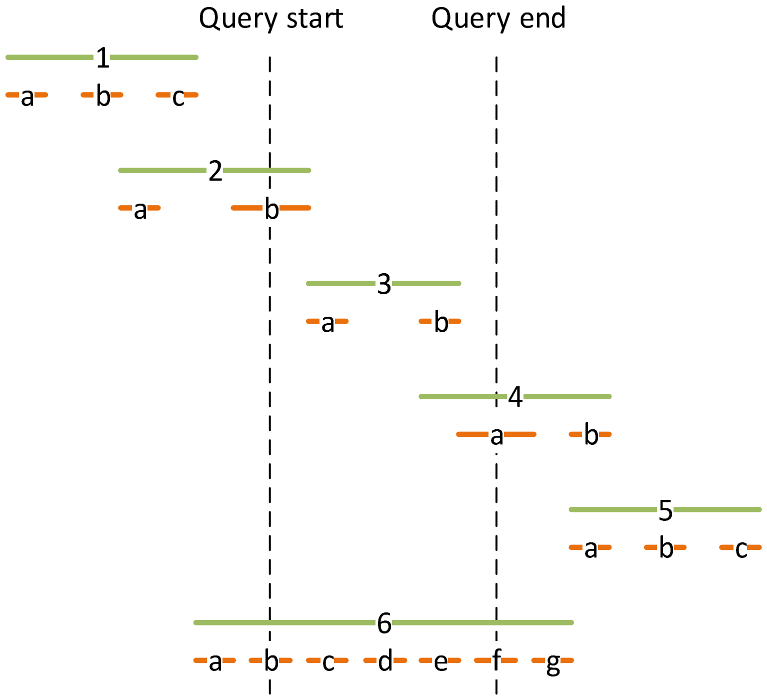

**Figure 10: Data store query**

If the user was searching only on the green records, 2, 3, 4 and 6 should be returned, because they occur at least partly within the query period.

If a second source is then joined in, the orange records that should be returned are shown in the table 4.

**Table 4: Data store query results**

| Record | Returned | Reason |
|---|---|---|
| 1a, 1b, 1c | No | All occurred outside the query window |
| 2a | No | As above |
| 2b | Yes | Occurred *partly* in the query window |
| 3a, 3b | Yes | Occurred *entirely* in the query window |
| 4a | Yes | Occurred *partly* in the query window |
| 4b | No | All occurred outside the query window |
| 5a, 5b, 5c | No | As above |
| 6a, 6g | No | As above |
| 6b, 6f | Yes | Occurred *partly* in the query window |
| 6c, 6d, 6e | Yes | Occurred *entirely* in the query window |

### Open-ended sessions

A complication of receiving session-based data as multiple records that mark events within that session, e.g. start and end, rather than single records at the end of a session describing the whole session, is that record delivery could go wrong. Records can go missing for various reasons.

Missing records highlight a problem that should be investigated and fixed. If identified, the system will need to know what to do in such an event.

### Missing end records

If a session start record is found, but the session end record is not available, the system will need to know how to determine where the session could have ended, or it will not be able to return the data at all. This would include records for which a session end has not yet been received because the session is still active and in use.

As can be seen from the following example, where the system needs to decide which CGNAT records occur within a RADIUS session, the end of that session is not known, so it is indeterminate which CGNAT records will have occurred within the session.

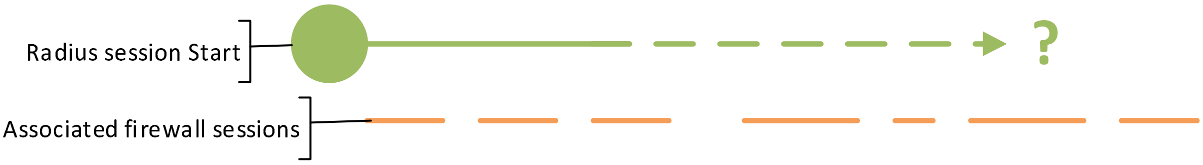

**Figure 11: Open ended sessions**

There are two options to identify the end of a session:

1) Assume a session terminates at the start of the next session.

A session could be deemed to have ended at the start of the next matching session.

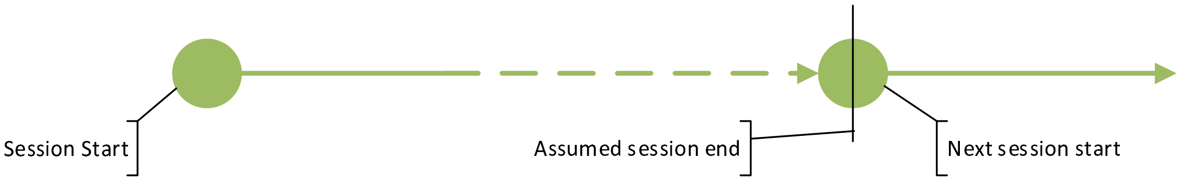

**Figure 12: Ending sessions at the start of the next session**

Sessions cannot overlap, e.g. with traditional NAT the same private IP address (and optionally port) cannot be in use by two subscribers at the same time, therefore a session needs to end before the next matching session starts. An assumed end record could therefore be placed at the same time as the next matching start record and all connected (CGNAT) events that happened in between the first start and the assumed end would be attributed to the identified session.

The risk of this approach is that the identified session is likely to have ended earlier than the next session start, and there could even have been another session in between that was missed entirely.

If two consecutive events found are a start and an end, but they do not match the same session, then the first session's end record and the second session's start record would have been missed and it would be indeterminate where one session had ended and the next had begun. CGNAT records associated with these sessions would therefore also be indeterminate as to which subscriber they belonged.

2) Assume a session terminated at the maximum session length.

A session could be determined to have ended when it has reached the maximum possible length.

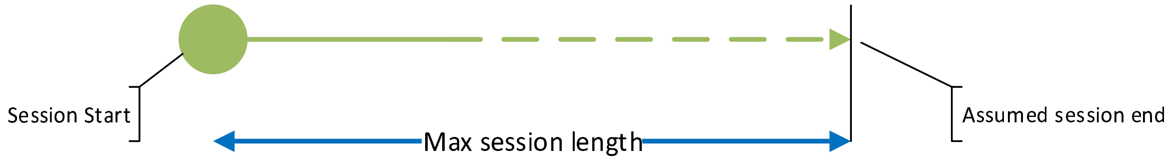

**Figure 13: Ending sessions at their maximum length**

A session cannot go past the maximum length set by the source system's configuration. Therefore, it can be assumed that any CGNAT event that happened after this is not connected to this session.

The downside to this approach is that the session could have ended much earlier, and CGNAT events found and assumed to be connected to it are actually connected to a different session entirely, and therefore were made by a different subscriber.

A combination of the two approaches may be the best approach to take - i.e. to extend the session to the maximum possible length, but only to the point of the next matching session start.

It should also be noted that if querying on the session data itself, the first record would be revealed by the search on the input field, e.g. MSISDN, but the second record would not be revealed by this search. The max session length approach could then take the session past the next session start record, which was not revealed. A secondary search on the session data store would have to be made on the join key with the secondary data store, e.g. IP address, because the second session would relate to a different MSISDN.

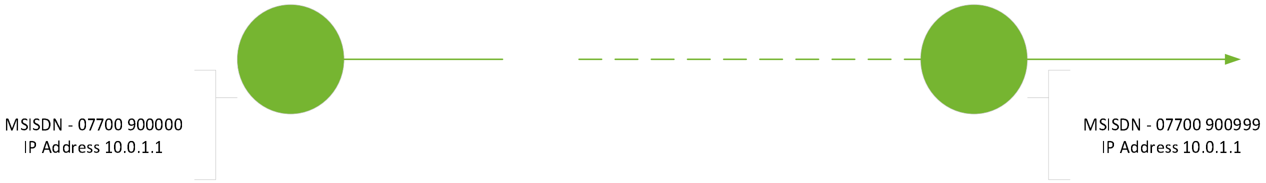

**Figure 14: Identifying the next session**

As can be seen from the above diagram, searching this store on MSISDN would only reveal the first record. To find the next session start, the query would have to be made on the IP address. If the end-user is searching on MSISDN, the purpose of searching this store is to find what IP addresses are in use, so a secondary search would be required if an open-ended session was found.

If CGNAT records were the primary search point, the session data would be joined in a secondary search on IP address already, and therefore all records would be returned.

**Missing start records**

The same principle above applies where a session start record is missing but a session end record is present.

#### Systems with intermediate records

For systems that produce intermediate records, the same principle again applies, but there can be confidence of joining records at least up to the last intermediate event found. However, these systems could have a much longer maximum session length and therefore the risk of getting the assumed end record wrong is increased.

## Searching data types

### Difference between searching and filtering

While the present document has thus far predominantly focused on mechanisms for *searching* a specific field, the same principles can apply when *filtering* the records returned.

A search and a filter are slightly different concepts. Searching is conducted against the data store to identify possible matches. Filtering is conducted against returned results to further narrow them down.

In a standard SQL type database, a select statement can include a *where* clause across multiple fields. The first field in that list is the primary search, and the following fields then narrow that set of results down. It is therefore more efficient to put an indexed field with high cardinality (low numbers of duplicates) as the first field in the *where* clause. E.g. filtering a list of people for those with a first name of John rather than filtering first on those who are male.

The fields chosen for searching, and therefore indexing, and those chosen only for filtering, will depend on query types allowed and what fields are indexed. Ideally, searching will be conducted against field(s) with high levels of cardinality (lots of unique values) and filtering can then be conducted on a smaller set of possible results.

An example would be to search on a source IP address, date and time, and filter by port number, or search on MSISDN, date and time and filter by destination IP address. By the nature of the data stored, i.e. the way it is partitioned, and query types in use in IP address retention, a date and time range will almost always form part of the search criteria.

### IP Address searching

*Exact match searching*

Using an IP address in its original format, e.g. a.b.c.d, only allows an exact match type search. i.e. from an address in this format, it is not immediately clear how to mathematically place it in a range provided by other IP addresses.

To search on an IP address by direct match, however, it needs to be stored and indexed in a standardized format - see clause 4.9. Provided a standard is adhered to, an exact string match search should be straightforward.

*Range searching*

Mathematically, it is not straightforward to compare an IP address as written to an IP address range. To do this, an IP address needs to be converted into a 4-byte integer for IPv4 or a 16-byte integer for IPv6.

## Joining data

### Creating sessions

#### Joining session boundary records

Within a data source, where records are point-in-time events, to create a session, a record indicating the Start will need to be joined with a record indicating the End of the session.

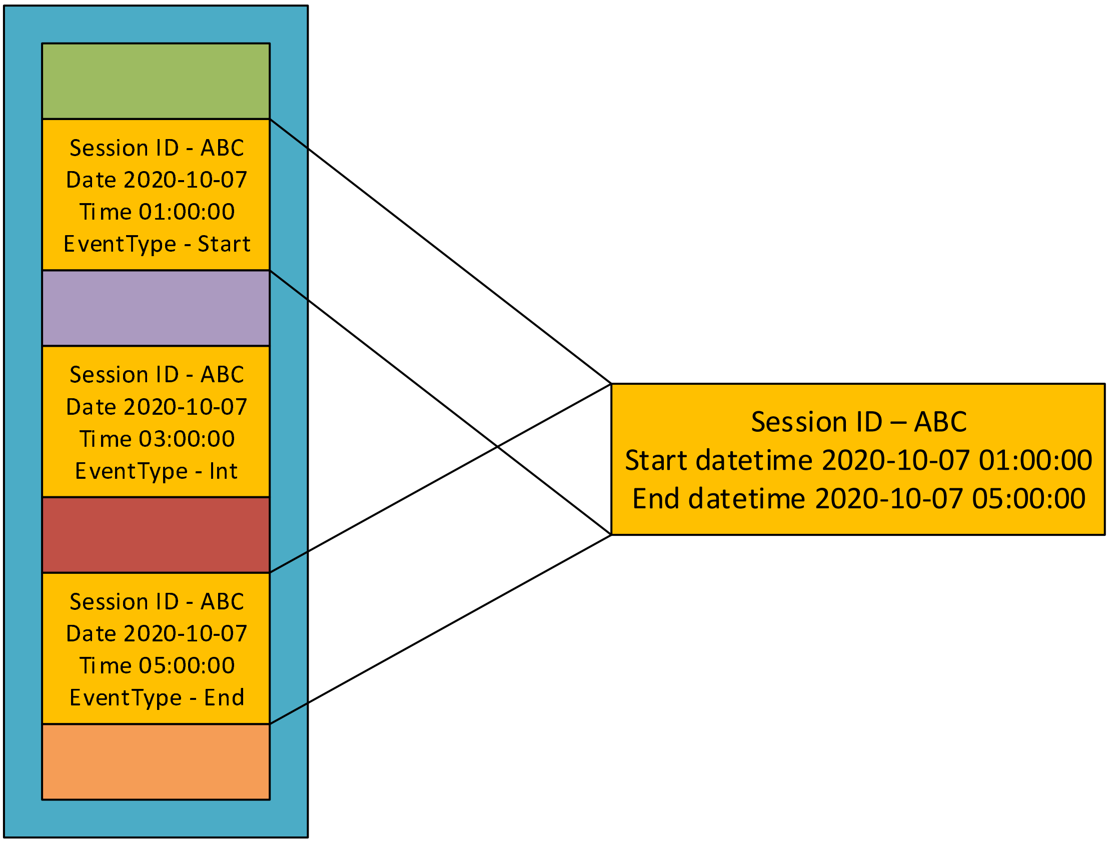

**Figure 15: Joining sessions**

Once the unique ID of a session is identified, records from that session can be joined to make a single session. If the search ID, e.g. MSISDN, is present in both, start and end records, then this join can take place in search results. If the search ID is not present in one of the record types, another option would be to perform a primary search on the user-provided search key, e.g. MSISDN, followed by a secondary search on the session IDs from those records that are found to find the other records that make up the session boundary.

Intermediate records are not required to make up a session, but they can be useful if one of the start or end records are missing, or to give a reason to extend the search for the session boundary record further.

**Extending the search time boundaries**

As a session exists in between the start and end records. A query range could cover a part of this period that has no records within it.

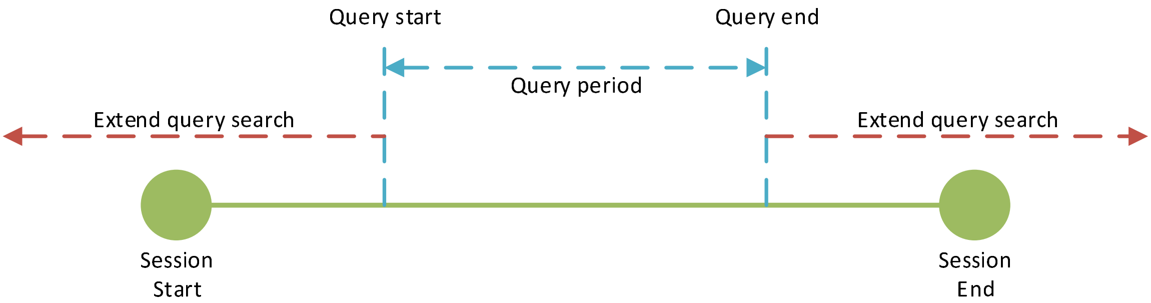

**Figure 16: Running a query within a session**

To ensure that a session is captured, the query period should be extended to capture any record that occurs within a given period before and after the original query search. The time period of the extension should be at least equal to the maximum length of a session (e.g. 24 hours), but not long enough for a session ID to become non-unique (e.g. 72 hours). Care should also be taken to ensure this takes into account any instance where the clocks may change.

Once this is done, any record that exists entirely outside the query period should then be discarded.

There are various scenarios that could be included, and the figure below shows why a query needs to be extended to find valid sessions. Each line represents a sample session. The session timeout duration is the same as the extension of the query period either side and is also shown in Figure 17.

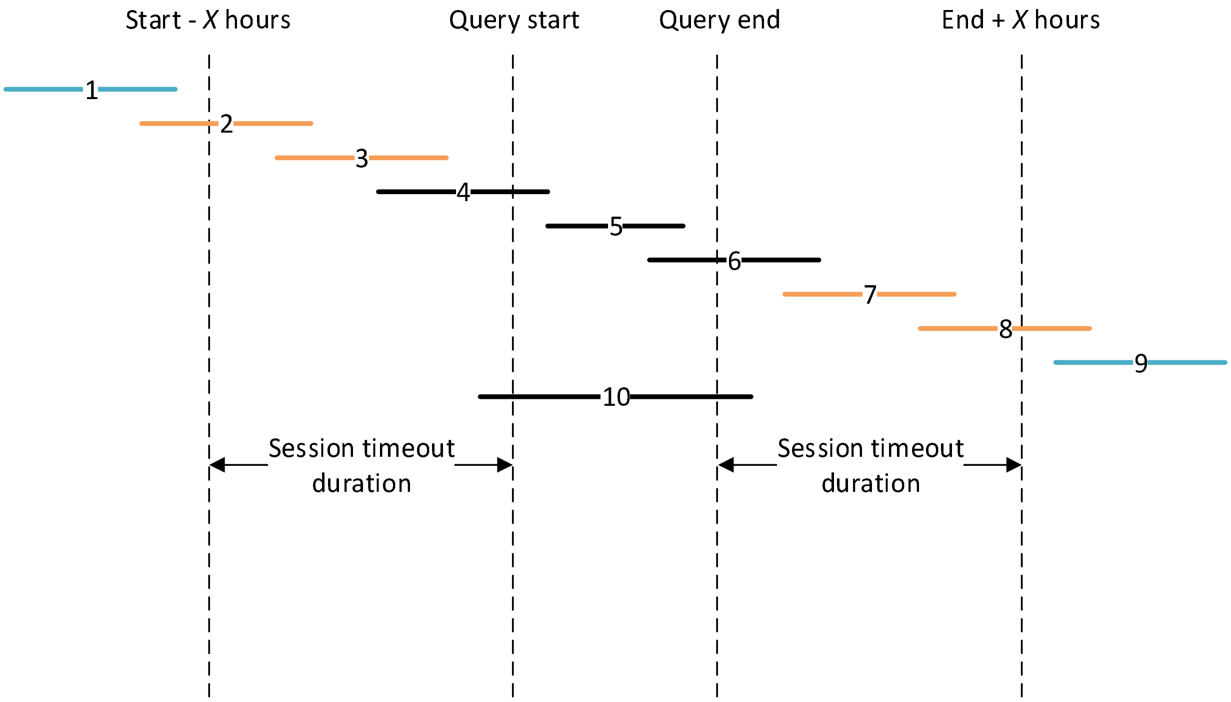

**Figure 17: Example sessions**

The key to the colours above:

**Black sessions** would be returned.

**Orange sessions** (2, 3, 7 and 8) would be found in extended search, but not returned.

**Blue sessions** would not be found in extended search and would not be returned.

Once valid sessions have been formed, one record per session can then be used in secondary queries and joins with other data if required, e.g. further subscriber details and account information such as required for subscriber name and address disclosure.

**Joining data**

Although AAA data is presented in terms of a session, CGNAT records could either be reported as events or sessions.

*Joining event-based records*

If the first stage of the query logic goes to the RADIUS store, sessions found can be used to define the time range(s) for the second stage of the query logic on CGNAT events. One query will be needed per AAA session, with a start and end of the AAA session being used as the query ranges. Any CGNAT events that occur outside the original input query range should then be discarded.

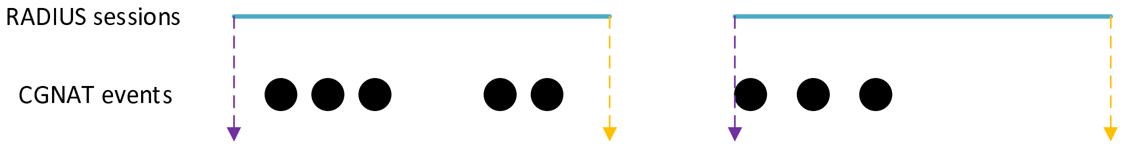

**Figure 18: Querying CGNAT events from a RADIUS session**

In Figure 18 above, the query start/end ranges for secondary queries can be seen. There are two RADIUS sessions, and therefore two secondary queries that will together return the CGNAT events shown.

If starting the query is from the CGNAT record store, a large number of results will often be returned. If each of the returned CGNAT events result in a secondary query against the AAA store, and each of those queries required the session building as previously described, the impact on query performance can be large.

Instead, some CSPs will be able to guarantee that if two CGNAT events occur within a given time period, then they are guaranteed to be from the same session.

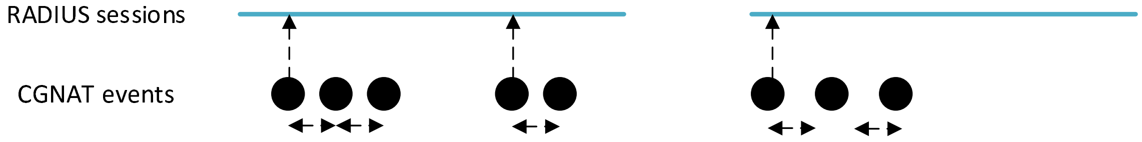

**Figure 19: Searching Radius sessions from CGNAT events**

In Figure 19, the first three CGNAT events occur within a given time period and are therefore guaranteed to be from the same RADIUS session, so only one query is required to find the Radius session. The next two CGNAT events occur outside this time period, so a second query against the RADIUS store is required, although it finds the same RADIUS session. The last three CGNAT events have occurred in a different RADIUS session. In this case, the secondary queries are point-in-time queries to check they exist within a given session as previously shown.

*Joining session-based records*

Where CGNAT logs are session-based, the search should be based on the start date in the CGNAT record, which should always occur within a given AAA session. The end date on the CGNAT record can then be used to check the validity of the returned record - it should either occur within the AAA session, or within a timeout period of the session ending (e.g. 30 minutes).

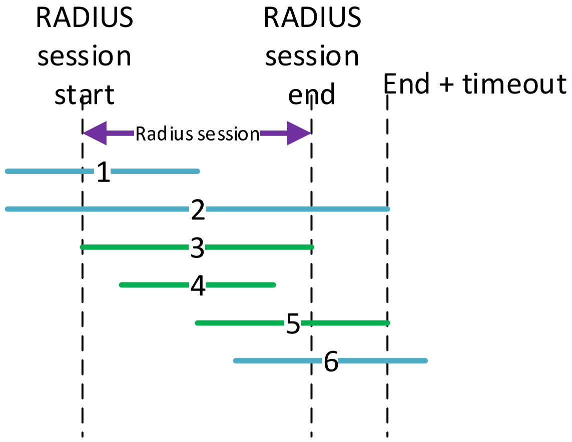

**Figure 20: RADIUS sessions and CGNAT sessions**

**Table 5: Radius and CGNAT session search results**

| CGNAT session number | Matches Radius session? | Reason |
|---|---|---|
| 1 | No | CGNAT session starts outside RADIUS session |
| 2 | No | As above |
| 3 | Yes | CGNAT session completely within RADIUS session |
| 4 | Yes | As above |
| 5 | Yes | CGNAT session starts within RADIUS session and finishes within the timeout period after the RADIUS session ends |
| 6 | No | CGNAT session finishes outside the RADIUS session end and after the timeout period has expired |

Figure 20 demonstrates how the CGNAT session start time can be used to validate which records are associated with an AAA session and therefore which CGNAT events can be said to be associated with those AAA records.

# Annex A: Change History

| Status of Technical Report ETSI TR 103 829<br>LEA support services;<br>IP address retention and traceability | | |
|---|---|---|
| **TC LI approval date** | **Version** | **Remarks** |
| July 2022 | 1.1.1 | First publication of the TR after approval by Remote Consensus after ETSI TC LI#60 (28-30 June 2022, Paris) |

# History

| **Document history** | | |
|---|---|---|
| V1.1.1 | August 2022 | Publication |
| | | |
| | | |
| | | |
| | | |