# 5G SA Roaming Performance: Empirical Evaluation of LBO, HR, and Distributed UPF Architectures

Oliver Zeidler
*Technical University of Munich*
Munich, Germany
oliver.zeidler@tum.de

Mert Günes
*Technical University of Munich*
Munich, Germany
mert.guenes@tum.de

Daniel Fraunholz
*Independent Researcher*
Munich, Germany
daniel.fraunholz@gmail.com

Wolfgang Kellerer
*Technical University of Munich*
Munich, Germany
wolfgang.kellerer@tum.de

***Abstract*—Mobile roaming ensures seamless international connectivity and primarily relies on two modes in 5G: Local Breakout (LBO) and Home Routed (HR) roaming. While network operators prefer HR for its centralized security and traffic control, tunneling data back to the home network incurs significant latency penalties. Conversely, LBO routes traffic locally, drastically reducing latency but diminishing the home operator's visibility. To address the lack of empirical comparison between these modes, we conduct a comprehensive end-to-end performance evaluation. Using an Open5GS-based testbed, we investigate five modes: a non-roaming baseline, standard LBO, standard HR, an IPsec-secured HR variant, and our proposed hybrid architecture, HR-remote-UPF. Evaluated across multiple visited networks at diverse server locations, our results quantify the operational trade-offs of each mode across a variety of setups. We show the notable overheads of HR over LBO and of IPsec in high-load scenarios. Ultimately, we demonstrate how HR-remote-UPF bridges this gap, achieving user plane latencies near those of LBO while providing home-operator control approaching that of conventional HR.**

***Index Terms*—5G, Roaming, Network Security, Performance Measurements**

## I. Introduction

The transition from 4G Long-Term Evolution (LTE) to 5G New Radio (NR) has introduced a paradigm shift in mobile networking, moving away from rigid hardware-based structures toward a cloud-native Service Based Architecture (SBA). While 5G promises unprecedented data rates and ultra-low latency, these benefits must extend beyond the subscriber's home network. As global mobility increases, the efficiency and security of international roaming have become critical focal points for Mobile Network Operators (MNOs).

In the 5G roaming landscape, two primary roaming modes dictate how data is routed: Local Break-Out (LBO) and Home Routed (HR) roaming. Currently, HR is the ubiquitous standard used across global deployments [1], [2]. In HR, all subscriber traffic is tunneled back from the visited Public Land Mobile Network (PLMN) (vPLMN) to the home PLMN (hPLMN). This allows the home operator to maintain full visibility and apply consistent charging and security policies. However, the resulting path inflation can dominate the end-to-end latency experienced by roaming users, undermining 5G's low-latency capabilities even when the radio access is not the bottleneck.

Conversely, LBO allows the User Equipment (UE) to access the internet or local services directly through the vPLMN's User Plane Function (UPF). While this drastically reduces latency and backhaul costs, it complicates the home operator's ability to monitor traffic and enforce security protocols. Consequently, despite the clear performance advantages of LBO, administrative and security concerns have hindered its widespread adoption.

To inform future roaming deployments, it is vital to quantify LBO's performance advantages while considering the associated security and operational trade-offs. Furthermore, this stark dichotomy between performance and control motivates the exploration of a hybrid mode. By combining the short user plane path of LBO with greater home-operator control than conventional LBO, we implement and evaluate a fusion approach, denoted *hr-remote-upf*. This was inspired by the widely discussed concept of Edge-UPFs, e.g. in [3]–[5]. We call it remote-UPF since the UPF is not located at the network edge, but in another country.

Despite theoretical discussions of these modes, there is a distinct lack of empirical data comparing their performance across diverse network conditions. This work closes this gap by performing realistic end-to-end tests using the open-source systems Open5GS [6] and PacketRusher [7]. To ensure a rigorous and comprehensive evaluation, our analysis incorporates a *non-roaming* baseline to establish absolute performance metrics. We evaluate this baseline against HR, LBO, and *hr-remote-upf* modes. Furthermore, we implement an HR roaming mode enhanced with Internet Protocol Security (IPsec) to accurately quantify the performance overhead introduced by robust security tunnels.

In detail, the contributions of this work are:

(C1) We provide a comprehensive overview of 5G roaming modes, focusing on LBO and HR.
(C2) We introduce *hr-remote-upf* as a fusion of LBO and HR, reducing the user plane path length of conventional HR while retaining greater home-operator control than LBO.
(C3) We provide an Open5GS-based, dockerized roaming testbed with configurable delays, experiment automation, and multiple UE-based user plane tests, available as open-source software at https://github.com/tum-lkn/LKN_5G_Roaming_Testbed.
(C4) We conduct a quantitative performance evaluation of five roaming modes (a *non-roaming* baseline, *lbo*, standard *hr*, *hr-ipsec*, and *hr-remote-upf*) across three vPLMNs and four Data Network (DN) server locations.

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

(C5) We discuss the performance and security trade-offs between LBO and HR 5G Standalone (SA) roaming, incorporating the empirical results to evaluate the viability of the *hr-remote-upf* mode.

The remainder of this work is structured as follows: First, related work is discussed in Section II. Section III provides an overview of 5G SA roaming, its security, and the two roaming modes LBO and HR, before introducing the proposed *hr-remote-upf* mode. Section IV focuses on the testbed setup and methodology of the performance tests, while Section V discusses the results of the performance evaluation. In Section VI, we discuss the security-performance trade-off between LBO, HR, and *hr-remote-upf* in more detail. Finally, Section VII discusses limitations and Section VIII concludes this work.

## II. Related Work

Several studies have conducted real-world measurements of roaming scenarios, though they largely focus on legacy or Non-Standalone (NSA) deployments. Mandalari et al. [1] and Fezeu et al. [8] conducted extensive UE-based measurements, but both investigations predated wide 5G SA roaming adoption. More recently, Trichias et al. [9] presented network-level measurements of neighboring 5G NSA networks, demonstrating that HR roaming significantly increases round-trip times compared to direct fiber interconnections. In contrast to these measurement campaigns, our work provides an empirical performance evaluation specifically targeting fully standardized 5G SA environments.

To evaluate core network behavior, various testbed approaches have been proposed. Keller et al. [10] focused on 4G/5G interworking, while Bhuyan et al. [11] demonstrated 5G roaming using Open5GS, though the latter only addressed control plane security in LBO. Similarly, Vázquez-Gallego et al. [12] investigated cross-border handovers with HR roaming, but omitted a standard-compliant Security Edge Protection Proxy (SEPP) implementation. The most similar methodological setup to ours is by Pawana et al. [13], who used Open5GS and PacketRusher to generate data for training Intrusion Detection Systems (IDSs). While they focused on deep learning for attack recognition, we utilize a similar standard-compliant toolchain to systematically compare the performance and structural trade-offs of HR and LBO configurations across all involved Network Functions (NFs).

The general security landscape of the 5G core and its critical interfaces (e.g., N9, N32) has been systematically analyzed in literature [14]–[16], but these works do not differentiate between HR and LBO roaming modes. Most closely aligned with our architectural focus, Motallebighomi et al. [2] discussed how the choice between LBO and HR fundamentally alters the network trust boundary regarding packet inspection and data handling. Our work builds directly on this architectural premise, quantifying the exact performance and infrastructure trade-offs required to maintain these distinct trust boundaries under standardized 5G SA configurations.

## III. Overview of 5G Roaming Modes

This section introduces the 5G roaming architecture, then details its security components and features. Furthermore, the two roaming modes LBO and HR, and our proposed fusion of the two, *hr-remote-upf*, are introduced.

### A. 5G Roaming

The advancement from the 4th generation Evolved Packet System (EPS) to the 5G SA system marks a fundamental change in telecommunications infrastructure. While previous mobile generations used roaming as a functional add-on of existing core network structures, 5G SA integrates roaming features deeply into its SBA. The interaction between vPLMN and hPLMN is not based on fixed point-to-point connections, but on flexible, Application Programming Interface (API)-based services, which run via Hyper Text Transfer Protocol version 2 (HTTP/2) and JavaScript Object Notation (JSON). The architecture of 5G systems is modular, with Control and User Plane Separation (CUPS) playing an important role. This separation enables routing control and data plane messages independently, which allows to differentiate between the two possible roaming modes, namely LBO and HR.

### B. Security in 5G SA Roaming

Both roaming modes rely on the SEPP, a new 5G NF located at the edge of the core network. The SEPP acts as a security gateway to protect the inter-PLMN communication via the N32 interface. This interface is divided into two logical parts: N32-c is used for the initial handshake, the negotiation of security parameters, and the exchange of error messages. N32-f is used to forward the actual NF-to-NF Service Based Interface (SBI) messages, after the negotiation of the security parameters via N32-c [17], [18].

On the control plane, the SEPP is responsible for all security-related activities. It implements mechanisms such as topology hiding to conceal internal network structures and message filtering to block malicious signaling traffic. Therefore, it is the main target of Denial-of-Service (DoS) attacks, impersonation, and fraud. Packets between SEPPs must be confidentiality- and integrity-protected, which is achieved through protocols such as Transport Layer Security (TLS), or the newer PRotocol for N32 INterconnect Security (PRINS) [19]. While TLS offers the best security and performance, PRINS enables Roaming Value Added Services (RVASs) and more flexible routing through the IP Exchange Network (IPX), enabling multiple business models [20]. Access to intermediary IPX-Providers (IPX-Ps) must be controlled and regulated to avoid privacy risks to critical user data needed for authentication. The SEPP improves the 5G SA roaming design by mitigating a number of existing roaming attacks, such as key theft and rerouting attacks, that were possible in Signaling System No.7 (SS7). Network node impersonation and source address spoofing, which were possible with Diameter, have also been mitigated [14].

The roaming user plane's UPF initially uses unencrypted GPRS Tunneling Protocol in the User plane (GTP-U) tunnels without mechanisms to filter malicious traffic, such as DoS

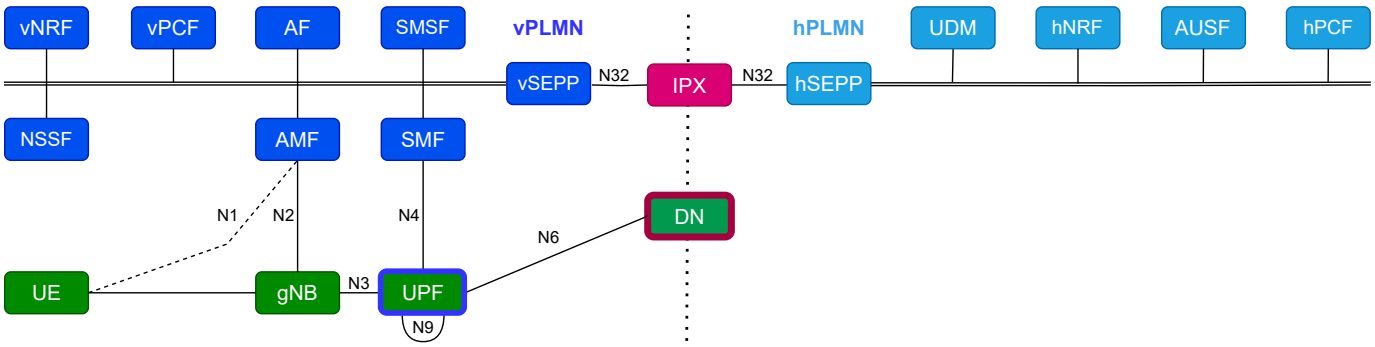



Fig. 1: Overview of LBO roaming functionality and involved NFs. The control plane of the vPLMN is shown in dark blue, the hPLMN in light blue, the user plane in green, and the IPX and DN in magenta.

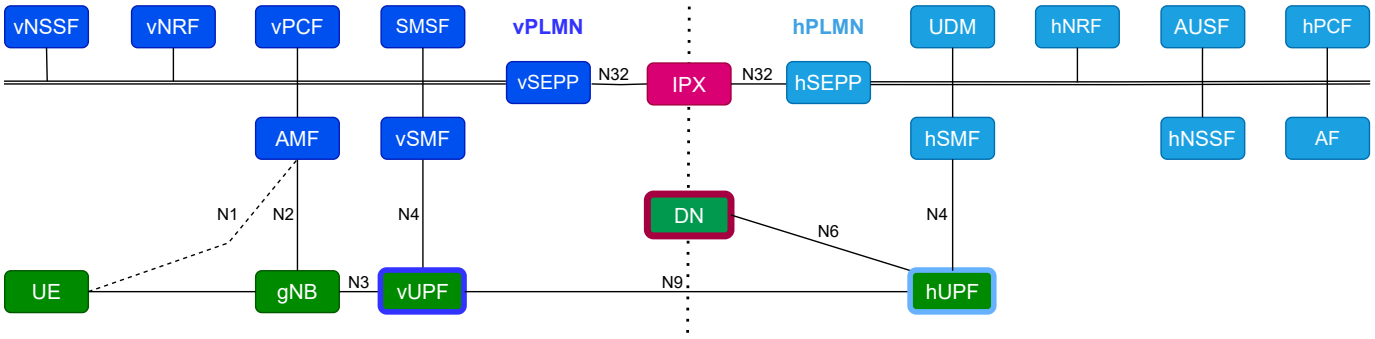



Fig. 2: Overview of HR roaming functionality and involved NFs. The control plane of the vPLMN is shown in dark blue, the hPLMN in light blue, the user plane in green, and the IPX and DN in magenta.

attacks. Because it is completely oblivious to security concerns, the 3rd Generation Partnership Project (3GPP) introduced the Inter PLMN User Plane Security (IPUPS) to address them. The IPUPS, which can be co-located with the UPF, is able to mitigate DoS attacks by accepting only packets from known tunnels that were established by the control plane NFs. To encrypt the GTP-U traffic, the 3GPP recommends using IPsec on the N9 interface between UPFs in roaming scenarios [17].

### C. Local Break-Out Roaming

LBO, the first discussed roaming mode, handles the roaming connection entirely in the control plane, while routing the user plane directly to the DN. Figure 1 shows all involved NFs and entities in the LBO roaming mode. In a normal user registration procedure, the UE connects via the base station (gNB) to the Access and Mobility Management Function (AMF) as the first control plane NF. As in a non-roaming scenario, the AMF is responsible for user registrations. The difference in a roaming scenario is that the UE's Mobile Country Code (MCC) and Mobile Network Code (MNC) do not match the ones of the vPLMN. This makes authenticating the user within the visited network impossible. Since only the hPLMN can authenticate the user, the visited Network Repository Function (NRF) (vNRF) must reach out to the home NRF (hNRF), which belongs to the home network and has the respective MCC and MNC. This connection relies on the SEPPs of both networks and on the secure use of the N32 interface, which passes through the IPX network. After the Authentication Server Function (AUSF), Policy Control Function (PCF), and Unified Data Management (UDM) in the home network have authenticated the roaming UE, the registration is finalized.

The next procedure is called Packet Data Unit (PDU) Session Establishment, which establishes the user's data session and enables internet connectivity. While the home PCF (hPCF) can still influence the policies applied to the user, the visited PCF (vPCF) can also act in accordance with predefined policies, e.g., those defined in a roaming agreement. In LBO roaming, the PDU session is routed directly to the attached DN from the visited UPF (vUPF). Here, the PDU Session Anchor (PSA) is the vUPF, leading to all roaming users receiving Internet Protocol (IP)-addresses from the visited MNO (vMNO)'s address range.

### D. Home Routed Roaming

The second discussed roaming mode, HR roaming, is defined as a connection between two PLMNs in which all user data is routed via its respective home network. Figure 2 displays the HR roaming mode and its distinct user plane setup. To make this user plane setup work, more control plane NFs need to be involved. The initial UE registration procedure works as in LBO, where the hPLMN authenticates the UE roaming in the vPLMN. The PDU session establishment procedure

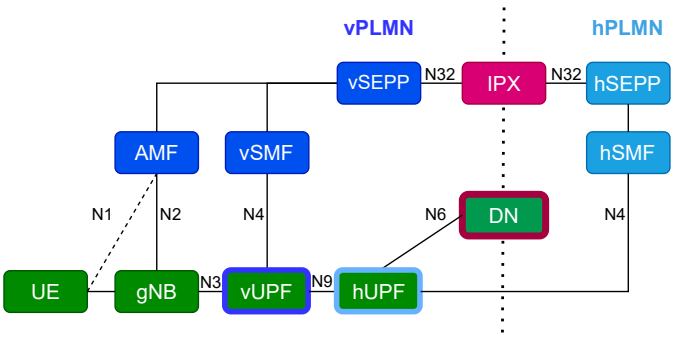



Fig. 3: Architecture overview of the proposed *hr-remote-upf* roaming mode. The key difference is the placement of the hUPF near the vPLMN to reduce user plane latency.

does differ between the roaming modes. For HR roaming, the session management is distributed among the visited Session Management Function (SMF) (vSMF) and its counterpart, the home SMF (hSMF). Therefore, the PDU session establishment involves vSMF-hSMF interactions, which are needed to set up a set of user plane tunnels between UE, gNB, vUPF, and home UPF (hUPF). To enable network slicing in the PDU session, the Network Slice Selection Functions (NSSFs) of both networks must exchange parameters via the roaming interface as well. Finally, user data flows via the N9 roaming interface between the vUPF and the hUPF, where the hUPF serves as the PSA. This means that all IP addresses assigned to roaming users stem from the home MNO (hMNO)'s address range.

### E. HR-remote-UPF (Proposed fusion of HR and LBO)

As mentioned earlier, our introduced mode *hr-remote-upf* is a combination of LBO and HR. It combines the HR roaming architecture with the Edge-UPF concept, where the UPF is located remotely in the visited country. Figure 3 shows the architecture of *hr-remote-upf*, which is based on the HR architecture. The only change is the location of the hUPF, which is moved across the international border into the vPLMN. Therefore, the user plane setup is comparable to LBO with an additional UPF, while the N4 interface between hSMF and hUPF gets extended beyond the roaming border. The expected impact is lower user plane latency and a higher control plane latency. Furthermore, the hMNO retains the ability to inspect and control the traffic through its hUPF, as in standard HR, although the extended N4 path introduces additional security requirements. As in HR, the PSA is the hUPF, so the roaming UEs get IP addresses from the home network.

## IV. Experimental Roaming Setup

This section overviews the hardware and software used and explores the investigated set of control plane procedures and user plane tests. Our fully software-based and Dockerized two-PC setup applies the 5G concept of Network Function Virtualization (NFV) while clearly distinguishing between vPLMN and hPLMN [21]. We focus on the 5G core network, while measuring the end-to-end delay from a UE's perspective.

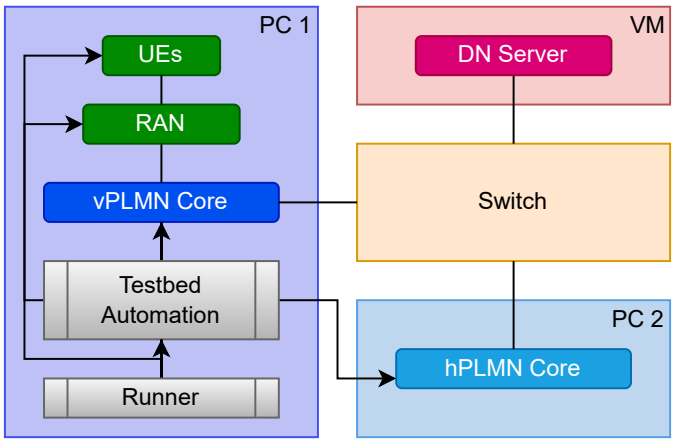



Fig. 4: Testbed architecture consisting of two PCs, one VM (red) and a switch (yellow). PC 1 hosts the vPLMN (dark blue), Radio Access Network (RAN) (green), and UEs (green), as well as testbed automation and runner scripts. PC 2 hosts the hPLMN (light blue), the VM hosts a DN server application, and the switch connects the three systems.

### A. Testbed

For this work, we have set up an open source roaming testbed consisting of two Open5GS systems [6] on separate, off-the-shelf PCs with PacketRusher [7] as a UE and RAN simulator. Both are widely used for open-source 5G experimentation and provide the functionality required for our roaming setup. We containerized all Open5GS and PacketRusher components using Docker, and use docker compose for start-up and shut-down.

The testbed supports a variety of configurable deployment modes that can be selected prior to each test run. Users can choose between HR and LBO roaming scenarios, each with optional TLS encryption for the communication between NFs. HR roaming can additionally be configured with IPsec on the N9 interface. Combined with a non-roaming baseline mode, these options yield multiple testable configurations, enabling systematic evaluation across a range of security and roaming scenarios.

Figure 4 displays our roaming testbed architecture with four major entities. To perform user plane tests, we implemented a simple DN server, which allows for ping, download and iperf testing. This DN server is hosted on a Virtual Machine (VM) and a switch is used to connect the two PCs and the VM. While PC 2 hosts only the hPLMN, PC 1 hosts the vPLMN, the RAN and UE simulations, and the scripts for testbed automation and large-scale test runs.

The `testbed_automation` script performs a variety of tasks, including starting the vPLMN and hPLMN cores simultaneously, where the latter is started via *ssh*. It can start PacketRusher a specified number of times back-to-back, differentiating between cold and warm start scenarios, and collecting more data for each time the two core networks are (re-)started. Each time PacketRusher is started, it performs multiple pre-defined user plane tests, like pinging, file downloading or iperf tests. For automated, large-scale

TABLE I: Overview of test run scenarios.

| Metric | Options |
|---|---|
| Mode | *non-roaming*, *lbo*, *hr*, *hr-ipsec*, *hr-remote-upf* |
| vPLMN | *France*, *USA*, *Australia* |
| DN Location | *Frankfurt*, *Paris*, *Boston*, *Sydney* |

test runs, we use `batch_runner`, which handles automated configuration file generation and repeated triggering of `testbed_automation`. We use 104 different configuration files, which will be motivated and explained in the following.

### B. Scenarios

The following settings allow us to differentiate between different constellations of roaming modes, vPLMNs, and DN locations. Table I provides an overview of all considered options. Note that we consider a fully TLS-secured control plane setup, as the performance impact of TLS on the 5G core was found to be small, and we aim to adopt the most secure, standard-compliant setup [22]. In detail, all SBI-communication in the vPLMN, the hPLMN, and on the N32-interface between the networks is encrypted with TLS.

*1) Roaming Modes:* We compare the five previously introduced roaming modes: *non-roaming*, which we use as a baseline, *lbo*, standard *hr*, *hr-ipsec*, and *hr-remote-upf*. We use italic, lowercase labels when referring to the corresponding testbed configurations.

*2) Visited Networks:* As the vPLMN's location significantly influences latency, we differentiate between three different vPLMNs, with the hPLMN being *Germany*: *France* for a shorter inner-european connection, *USA* for a longer transatlantic connection, and *Australia* for a worst-case connection around the globe. In our tests, we emulate the delay between the hPLMN and vPLMN by applying it between the respective NFs. For all roaming modes, except for *non-roaming*, we add a delay between both SEPPs. The delay is added as a one-way delay on both ends, resulting in a full delay with respect to Roundtrip Times (RTTs). Furthermore, in *hr* and *hr-ipsec*, we apply a delay between vUPF and hUPF. For *hr-remote-upf*, the second delay is added between hSMF and the remote hUPF, while no user plane delay is added for *non-roaming* and *lbo*.

*3) DN Server Locations:* Since the roaming modes can also affect the choice of DN server locations, we selected four different DN server locations for our tests: *Frankfurt*, *Paris*, *Boston*, and *Sydney*. Depending on the configured DN location, the specified one-way delay is added to outgoing packets on the vUPF, for *non-roaming*, *lbo*, and *hr-remote-upf*, or on the hUPF, for *hr* and *hr-ipsec*. The other side of the one-way delay is applied to the DN server, providing realistic delay conditions.

*4) Performance Metrics:* For our performance analysis, we look at the effect of the previously introduced sets of settings on the UE. Therefore, we always start a single UE, to avoid mixing procedures and times. Each time we start a UE, we measure certain control plane procedures and perform user plane tests to collect performance data.

The two most relevant control plane procedures are the UE registration and PDU session establishment procedures. We measure both individually, by capturing the time between the `initialUE registration`, the `PDUSessionEstablishmentRequest`, and the `PDUSessionSetupResponse` messages from the PacketRusher logs. These timestamps yield two Key Performance Indicators (KPIs): the UE registration and the PDU session setup times.

TABLE II: One-way DN latencies for home and local break out [23].

| From -> To | Frankfurt | Paris | Boston | Sydney |
|---|---|---|---|---|
| Frankfurt | 0 | 16±1 ms | 44±1 ms | 150±2 ms |
| Paris | 16±1 ms | 0 | 40±1 ms | 139±2 ms |
| Boston | 44±1 ms | 40±1 ms | 0 | 104±2 ms |
| Sydney | 150±2 ms | 139±2 ms | 104±2 ms | 0 |

Once the UE can access the internet, it performs three subsequent tests. First, it pings the IP address of the DN server on the VM. This is done for a specified number of times with a predefined payload size, and returns the minimum, average, and maximum ping latencies. Second, the UE downloads a chosen file from the DN server and tracks the following KPIs: Time To First Byte (TTFB), download speed, and total download duration. Third, the UE performs an iperf test with the DN server as its endpoint. The test duration, direction (uplink or downlink), and other factors can be set in advance. In terms of KPIs, it tracks bits per second (bps), jitter, and packet loss. Additionally, we measure the effective throughput as defined in Equation 1.

$$throughput_{eff} = target_bandwidth * (1 - \frac{packet_loss}{100}) \quad (1)$$

### C. Model Configuration

This section presents the specific parameter choices for all previously introduced settings. Table II displays the one-way latencies that are applied between UPF and DN server. With the hPLMN being *Germany* and the vPLMNs being *France*, *USA*, and *Australia*, the first line also displays the latencies applied between the PLMNs. If the vPLMN is *France*, 16 ms latency with 1 ms jitter are applied to every packet leaving the vPLMN towards the hPLMN. The same delay is applied on the way back. For *USA* as a vPLMN, the delay is 44±1 ms, while it is increased to 150±2 ms for *Australia*. The values were taken from Wondernetwork.com on 30.01.2026 [23] and, therefore, mark realistic values for these distances.

For the DN server latencies, all possible combinations between *Frankfurt*, *Paris*, *Boston*, and *Sydney* are considered. Here, *Frankfurt* corresponds to the hPLMN *Germany*, while *Paris*, *Boston*, and *Sydney* correspond to the vPLMNs *France*, *USA*, and *Australia*, respectively.

Table III lists the chosen metrics for the user plane tests. We limit the parameter choice to two ping payloads of 56 and 1400 B respectively, to avoid bloating the number of configurations to test. A smaller number of configurations allows for more test runs and, therefore, a better statistical significance. Each UE pings the DN server 20 times, with an

TABLE III: Chosen user plane test metrics.

| Metric | Options/Choice |
|---|---|
| Number of Pings | 20 |
| Ping Payload | 56 B, 1400 B |
| Ping Interval | 0.2 s |
| Download File Size | 10 MB |
| iperf Duration | 5 s |
| iperf Direction | Uplink |
| iperf Bandwidth | 25 Mbps, 100 Mbps |

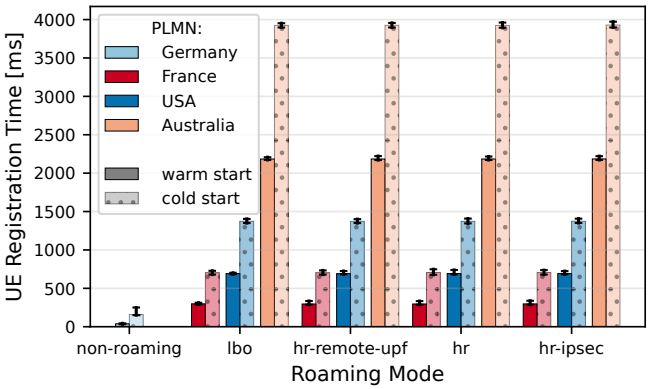



Fig. 5: UE Registration time for the five roaming modes according to the vPLMN location.

interval of 0.2 s, and a payload of either 56 or 1400 B. We chose these values to limit the time each UE spends pinging, after we found ping latencies to be quite stable in preliminary tests. Similarly, we chose a 10 MB file as the download target, rather than 50 or 100 MB, to reduce the time spent per UE. For the iperf test, we set the duration to 5 s with the same intent. As we already perform a download test, we limit the iperf test to an upload scenario.

## V. Performance Evaluation

In this section, we evaluate the results of our performance tests across the 104 configurations. Each configuration was executed in 100 test runs, with five sequential UE executions per run, yielding 100 warm and 400 cold start observations. The following is divided into two subsections focusing on the control and user plane results, respectively. The bar plots display the medians with error bars spanning the 2nd to 98th percentiles.

### A. Control Plane

Since the control plane results do not depend on the DN location or the ping payload, the number of valid data points increases. All control plane data show low variance, indicating stable behavior across repeated test runs. The figures show both cold and warm start values, where the first UE corresponds to the cold start, while UEs two to five correspond to the warm start.

*1) UE Registration:* Figure 5 shows the duration of the UE registration procedure for all five roaming modes according to the vPLMN location. It becomes apparent that *non-roaming* is faster than any roaming approach. Within the other roaming

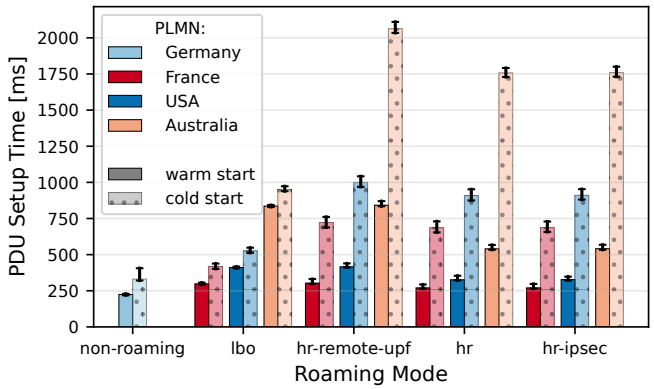



Fig. 6: UE PDU Session Establishment time for the five roaming modes according to the vPLMN location.

modes, there are no visible differences. The most notable differences are between the vPLMNs, which shows the effect of different delays induced by physical distance.

In the warm start, roaming adds a delay of around seven RTTs per UE. Two RTTs are needed for the authentication procedure, two for UDM registration and subscription, one for policy establishment, and a final two for finalizing PCF and UDM interactions. Those subprocedures happen for every registration [24].

The cold start overhead stems from the selection procedures for AUSF, UDM, and PCF. Each selection procedure involves the hNRF for service discovery and, thus, adds one RTT. Additionally, the N32 connection setup is triggered by the UE registration. This adds another three RTTs for the TLS handshake and HTTP/2 init. Finally, the cold start overhead adds up to 6 RTTs. Naturally, the UE registration latency scales with the delay between the networks. The RTTs for *France*, *USA*, and *Australia* are 32 ms, 88 ms, and 300 ms, respectively. Since a cold start adds around 13 RTTs, and a warm start adds 7 RTTs, this accumulates to a theoretical overhead of 3900 ms in a cold start and 2100 ms in a warm start for the example of *Australia*. This theoretical overhead matches the results shown in Figure 5.

*2) PDU Session Establishment:* Figure 6 displays the same modes and networks for the duration of the PDU session establishment procedure. This metric provides more differences between the roaming modes and the cold and warm start scenarios. The *non-roaming* and *lbo* modes exhibit a cold start overhead of 106-118 ms, regardless of the vPLMN. Compared to *non-roaming*, *lbo* adds an overhead of around two RTTs. This can be explained by the standardized PDU session establishment procedure itself, which involves the home UDM (hUDM) and hPCF. Since both functions were already involved in the UE registration procedure, no significant service discovery overhead is observed here. The vSMF has to discover both hUDM and hPCF, but when asking the vNRF it directly receives the answer from the vNRF's cache, therefore, avoiding inter-PLMN RTTs. However, in every cold and warm start, both home NFs must be called separately, adding a delay of two

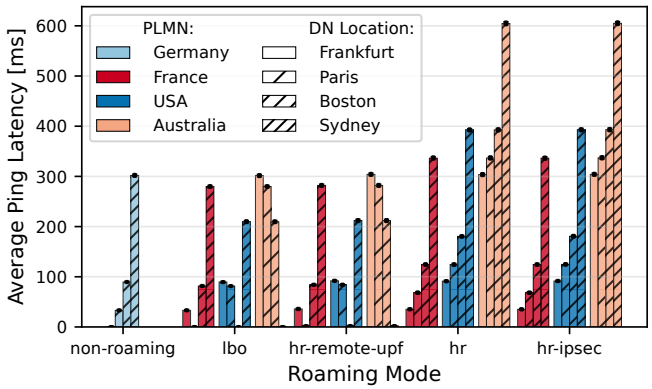



Fig. 7: Average ping latency for the five roaming modes according to the vPLMN and DN locations.

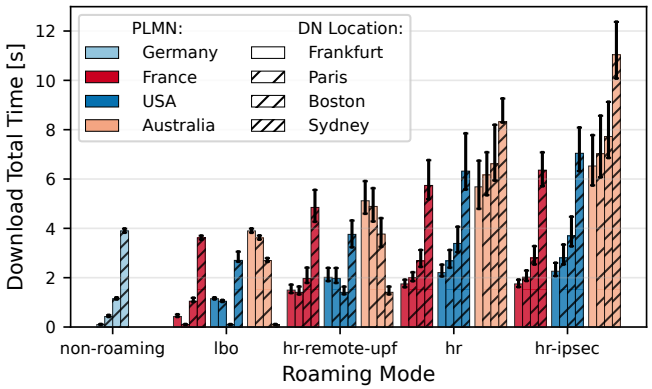



Fig. 8: Total download times of a 10 MB file for the five roaming modes according to the vPLMN and DN locations.

RTTs [24].

For the remaining modes, the cold start incurs a considerable overhead of 3 RTTs. One is needed for the hSMF service discovery by the vSMF and the other two are for the NSSF discovery and interaction. Since the exchanged data is cached in the vPLMN, the home NSSF plays no role in the warm start [24]. The modes *hr* and *hr-ipsec* have equal values since their notable difference, the IPsec encryption, only affects the user plane measurements. The absolute cold start overhead for *hr* and *hr-ipsec* is also equal to *hr-remote-upf*, since no handshake is required on the N4 interface.

*hr-remote-upf* has the highest peak in the cold start values since the PDU session establishment involves the hUPF, and in *hr-remote-upf* delay is added between hUPF and hSMF. Therefore, in this scenario, the inter-PLMN delay is added on two relevant interfaces, namely N32 and N4.

Surprisingly, while *lbo* is notably faster than all *hr* related roaming modes in the cold start, that changes in the warm start. Here, the values of *lbo* and *hr-remote-upf* are quite close and yield longer procedure durations than *hr* and *hr-ipsec*. This might seem counterintuitive; however, in HR, the hSMF handles the procedures with hUDM and hPCF within the home network. Thus, there is only one inter-PLMN RTT between vSMF and hSMF. In contrast, LBO only uses the vSMF, which has to communicate with both hUDM and hPCF via the roaming interface, adding two RTTs. Similarly, *hr-remote-upf* adds two RTTs, one between vSMF and hSMF and another one with the Packet Forwarding Control Protocol (PFCP) messages between hSMF and hUPF.

### B. User Plane

In the following user plane results, the four DN locations become relevant, while there is no notable difference between cold and warm start results. Thus, cold and warm start results are joined.

*1) Ping latencies:* Figure 7 shows ping latencies of all five roaming modes, all four countries (three vPLMNs and the hPLMN), and four DN server locations. It allows for drawing multiple conclusions about the different setup and mode combinations. First, there is no visible difference between *lbo* and *hr-remote-upf*, supporting the assumption that placing the hUPF in the visited country can achieve ping latencies comparable to LBO. Second, there is no visible overhead of IPsec, making *hr-ipsec* comparable to *hr*. The mean overhead of IPsec is 0.25 ms, indicating negligible IPsec overhead for ping latency in our setup. With latencies ranging up to 605 ms, the IPsec overhead for ping latencies is negligible. Third, ping latencies scale consistently with the configured delays, both for vPLMNs and for DN server locations. For example, if the roaming mode is *hr*, the vPLMN is *Australia*, and the DN location *Sydney*, the ping latency reaches its maximum of 605 ms. This is almost exactly one RTT for *Australia* and one RTT for *Sydney*. Fourth, Figure 7 includes both ping payloads of 56 B and 1400 B and all data between the 2nd and 98th percentiles. Thus, the ping payload has no influence on the latencies.

The TTFB metric from the 10 MB file download shows the exact same trends, but with a scaling factor of 2. Therefore, it is not shown here. The scaling factor is introduced because initiating the file download takes 1.5 RTTs, and the first download data arrives at the UE after two RTTs.

*2) File download:* Figure 8 displays the total time taken for the UE to download the 10 MB file from the performance server VM. In contrast to the highly stable ping results, the download times exhibit greater variance, which scales proportionally with increasing network delay. Furthermore, the download performance reveals a more pronounced differentiation among the *lbo*, *hr-remote-upf*, *hr*, and *hr-ipsec* modes.

Introducing an additional UPF in the *hr-remote-upf* configuration incurs a measurable performance penalty, resulting in an overhead of 0.92 s to 1.31 s compared to *lbo*. This impact remains significant even when the UPF is deployed within the visited country. When comparing *hr* to *hr-remote-upf*, standard *hr* introduces a further overhead ranging from 0.25 s to 4.27 s (extending up to 6.92 s under extreme routing conditions).

This trend is clearly illustrated when examining the vPLMN *Australia*. For a local user accessing a server in *Sydney*, *lbo* achieves a median download time of 0.10 s. Under the same

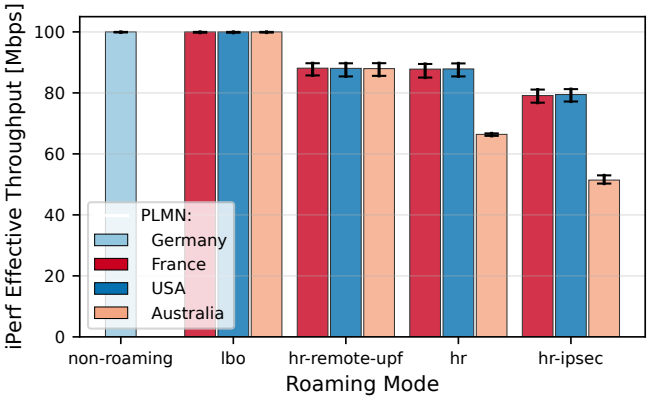



Fig. 9: Effective iperf throughput with a target bandwidth of 100 Mbps for the five roaming modes according to the vPLMN locations.

conditions, *hr-remote-upf* increases the download time to 1.41 s (a 1.31 s overhead), while standard *hr* escalates the download time to 8.33 s—representing an overhead of 8.23 s over *lbo* and 6.92 s over *hr-remote-upf*. Conversely, when evaluating a server located in the home country (*Frankfurt*), *hr-remote-upf* adds 1.22 s on top of *lbo* (5.12 s vs. 3.90 s), whereas standard *hr* adds 1.78 s (5.68 s vs. 3.90 s). Consequently, shifting the UPF to the visited country effectively reduces the severe routing overhead of standard *hr*, even though *hr-remote-upf* inherently introduces a baseline processing overhead.

Figure 8 also provides insights into the performance impact of securing the N9 interface via IPsec. The additional overhead of *hr-ipsec* over standard *hr* ranges from 0.0 s up to 2.71 s. The worst case occurs when routing to *Sydney* from *Germany* while roaming in *Australia*, where the median rises from 8.33 s to 11.04 s. This demonstrates that IPsec overhead scales heavily with the underlying network delay. Across all tested scenarios, the average IPsec overhead is 0.62 s, indicating that securing the roaming interface compromises download performance by up to 33 % in high-latency environments.

*3) iperf throughput:* Figure 9 displays the effective throughput of our iperf User Datagram Protocol (UDP) tests with a target bandwidth of 100 Mbps. Both *non-roaming* and *lbo* consistently achieve the full 100 Mbps target with no noticeable packet loss. In contrast, the HR-based architectures experience throughput degradation due to three distinct bottlenecks.

First, in our implementation, the secondary UPF introduces a baseline throughput reduction of approximately 12 %, with *hr-remote-upf* stabilizing at roughly 88 Mbps across all visited networks due to static processing limitations.

Second, severe propagation delay combined with an unmanaged UDP stream triggers heavy packet loss. In high-latency scenarios, extended transit times aggravate buffer constraints on intermediate routers. Without a transport-layer feedback loop to throttle the sender, buffers saturate, causing packet drops and out-of-order deliveries. This is highly visible when the vPLMN is *Australia* under standard *hr*, where throughput drops sharply to a median of 66.40 Mbps. Conversely, *hr-remote-upf* avoids this additional 22 % degradation (maintaining 87.96 Mbps) because the remote UPF keeps the user plane path short, minimizing transit bottlenecks.

Third, IPsec processing on the N9 interface introduces a severe compounding penalty. In lower-latency roaming environments (*France* and *USA*), *hr-ipsec* reduces throughput further to approximately 79 Mbps. However, when combined with the extreme latency of the Australian path, the lack of UDP flow control results in severe packet drops, reducing the median throughput to 51.41 Mbps.

In comparison, with a target bandwidth of 25 Mbps, there was no packet loss, and the effective throughput remained at 25 Mbps across all modes, vPLMNs, and DN locations.

*4) Cross-Protocol Configuration Impacts on Performance:* During our tests, we observed severe performance degradation and unexpected behaviors that were not explainable by 5G or security considerations. To help future researchers avoid the same problems, we want to highlight some notable findings.

The Maximum Transmission Unit (MTU), especially at the UPFs, plays an important role regarding GTP-U and IPsec encapsulation. If the MTU is exceeded, packets are fragmented, increasing the load on the UPFs. Once delays are added, packet fragmentation and out-of-order deliveries accumulate, reducing performance. We suggest reducing the MTU size from 1500 B to 1350 B for user plane performance tests.

While GTP-U relies on UDP, end-to-end applications often use Transmission Control Protocol (TCP). Depending on the TCP variant, such as Cubic and Reno, large delays can lead to a significant performance drop. To mitigate this transport-layer bottleneck, we utilized TCP Bottleneck Bandwidth and RTT (BBR). Unlike Reno or Cubic, BBR models the actual maximum bandwidth and minimum RTT of the path rather than treating isolated packet loss as a congestion signal. This allows BBR to maintain high throughput and stable download times even in high-latency, loss-prone roaming environments (such as the *Australia* scenarios). Consequently, we recommend utilizing TCP BBR alongside mechanisms such as TCP Fast Open and an optimized initial congestion window size for performance testing.

## VI. Discussion of the Security-Performance Trade-off

Previously, we performed a quantitative performance evaluation to compare the five roaming modes defined in this work. In this section, we aim to discuss the fundamental security-performance trade-offs inherent to these architectural choices, focusing specifically on our proposed mode, *hr-remote-upf*.

Regarding control plane performance, *hr-remote-upf* adds overhead to the PDU session establishment procedure, especially in the rare case of a cold start. With a warm start, its remote UPF location incurs higher overhead than *hr*, but achieves the same performance as *lbo*. For user plane metrics like ping latency and download TTFB, *hr-remote-upf* reaches the same speeds as *lbo*, clearly outperforming standard *hr*. However, high-load user plane evaluations reveal deeper architectural distinctions. In the total download time evaluation,

the secondary UPF in *hr-remote-upf* introduces a measurable baseline overhead of up to 1.31 s over *lbo*, while standard *hr*'s download times increase by up to 8.23 s over *lbo*. Similarly, upload stress tests using iperf demonstrate that *hr-remote-upf* incurs a constant 12 % throughput drop compared to *lbo*. This confirms that, in our implementation, processing by the additional UPF degrades performance under heavy traffic loads, even when geographic delays are minimal. Crucially, *hr-remote-upf* avoids the severe RTT-induced throughput degradation that plummets standard *hr* performance by an additional 22 % in high-latency scenarios. Furthermore, high-load scenarios amplify security overhead: while *hr-ipsec*'s latency impact appears minor in low-load ping tests, it increases download times by up to 33% and reduces the reported UDP throughput from 66.40 Mbps to 51.41 Mbps in the most demanding scenario. Additionally, performance tests with TCP and larger MTU sizes highlighted the vulnerability of HR-based roaming modes to network misconfigurations.

From a security perspective, the choice between HR and LBO fundamentally shifts the risk distribution between the vPLMN and the hPLMN. Generally, standard HR limits the home network's exposure to the visited network, whereas LBO expands the attack surface by routing user plane traffic directly through the vPLMN. For instance, if a vMNO is malicious or compromised, HR is inherently more secure than LBO. While the ultimate security of the hPLMN always depends on the configurations deployed by the hMNO, ensuring the confidentiality and integrity of a roaming UE's data is significantly more challenging for the hMNO under an *lbo* architecture than under *hr*.

The proposed *hr-remote-upf* mode operates similarly to HR roaming, but introduces specific structural security considerations regarding the N4 interface and PFCP. Although the 3GPP specifies that these connections should be protected via IPsec when physical security cannot be guaranteed [17], deployment interpretations vary, leaving potential room for vulnerabilities. Because *hr-remote-upf* increases the physical distance and involves more intermediate parties in forwarding and processing data across the N4 interface, the network's attack surface expands. This makes the deployment more susceptible to Man-in-the-Middle (MitM) attacks and other known PFCP threats [25]–[27]. Consequently, adopting *hr-remote-upf* necessitates stricter security controls, such as deploying IDSs to detect and mitigate PFCP-specific attacks [28], [29].

## VII. Limitations

While this work demonstrates the feasibility and performance characteristics of the proposed roaming architectures, several limitations must be acknowledged.

First, our quantitative evaluation was conducted using open-source software stacks deployed on resource-constrained mini PCs rather than carrier-grade hardware or large-scale server farms. Additionally, we utilized a simplified roaming topology that omits the complex, multi-hop IPX network setups typical of real-world deployments. Thus, the absolute results are testbed-specific, while the relative performance trends provide a controlled comparison of the five configurations and may not fully capture production-level processing, routing, and scaling effects. Validation with commercial UEs and an SDR-based RAN remains future work, as our emulated setup intentionally isolates core-network and roaming-path effects. This is particularly relevant to our security layer: while IPsec introduces notable computational and encapsulation overhead in high-traffic scenarios, these bottlenecks can be significantly mitigated in production environments through dedicated hardware acceleration, such as cryptographic offloading, and optimized architectural designs. Furthermore, due to space constraints, our security analysis is limited to a high-level discussion of architectural risks rather than an empirical assessment of vulnerabilities or cryptographic mechanisms.

Beyond the experimental setup, the practical deployment of *hr-remote-upf* introduces substantial operational and administrative challenges for an hMNO. Establishing and securing a physical or virtual site in a foreign country involves complex regulatory and legal hurdles, including compliance with local lawful interception mandates, telecom licensing, and strict data sovereignty laws. Financially and logistically, hosting infrastructure abroad requires either deploying hMNO personnel to remote locations or outsourcing maintenance to third-party providers, which introduces significant trust and supply-chain security concerns.

Finally, the success of *hr-remote-upf* depends on the routing capabilities of the visited network; the vPLMN must be capable of routing user plane data locally to the remote hUPF instead of utilizing the standard IPX network path, which would otherwise nullify the latency benefits of our architecture. Given these complexities, *hr-remote-upf* is not a universal replacement for existing roaming modes. Instead, we propose it as a targeted solution, ideally restricted to a limited number of strategic roaming partnerships that account for the majority of an operator's roaming traffic.

## VIII. Conclusion

The rollout of 5G SA networks promises ultra-low latency and high bandwidth, yet realizing these benefits for roaming subscribers remains a significant architectural challenge. In this work, we addressed the stark trade-off between the performance advantages of LBO and the administrative control and security benefits of HR roaming. To move beyond theoretical discussions, we deployed an Open5GS-based end-to-end testbed to empirically evaluate five distinct roaming configurations: *non-roaming*, *lbo*, *hr*, an IPsec-secured HR variant, *hr-ipsec*, and our proposed hybrid architecture, *hr-remote-upf*.

Our multi-layered performance evaluation clearly quantified the heavy performance degradation inherent to standard HR roaming under high-stress conditions. While LBO consistently achieved baseline performance, standard *hr* exhibited severe routing latency, increasing 10 MB download times by up to 8.23 s compared to *lbo* and reducing UDP throughput to 66.40 Mbps in high-latency scenarios. This performance gap was drastically widened by the integration of IPsec. Under heavy stress, *hr-ipsec* added a severe compounding penalty

due to bottlenecks and transit packet loss, escalating download times by up to 2.71 s and plummeting UDP throughput to a median of 51.41 Mbps.

To bridge this performance-security gap, we introduced and evaluated the *hr-remote-upf* configuration. Our results demonstrate that this hybrid approach successfully reconciles these competing requirements. By strategically distributing the UPF to the visited country, *hr-remote-upf* kept ping latencies close to *lbo*. Furthermore, under high-load testing, it effectively decoupled throughput from wide-area path latency, maintaining a stable median throughput of 88 Mbps and isolating the user plane from RTT-driven packet drops. This allows the architecture to preserve home-operator-controlled charging, policy enforcement, and traffic-steering capabilities associated with HR, subject to additional safeguards for the extended N4/PFCP path, without sacrificing 5G's low-latency promises. However, deploying and managing remote UPF locations entails operational costs that MNOs must carefully weigh.

Future work will focus on evaluating the effectiveness of attacks implemented within our roaming testbed, alongside investigating the deployment and performance impacts of Post-Quantum IPsec to safeguard against emerging cryptographic threats. Furthermore, future iterations will explore IPUPS mechanisms to evaluate their efficiency and robustness in securing user plane traffic within this architecture.

## Acknowledgment

The authors would like to thank the roaming team of Deutsche Telekom for the fruitful discussions.

Gemini 3.1 Pro was used for language editing of selected paragraphs. The authors reviewed and take responsibility for all content in the manuscript.

The authors acknowledge the financial support by the Federal Ministry of Research, Technology and Space of Germany (BMFTR) in the programme of "Souverän. Digital. Vernetzt." joint project 6G-life, project identification number 16KIS2414.

## References

[1] A. M. Mandalari, A. Lutu, A. Custura, A. S. Khatouni, Alay, M. Bagnulo, V. Bajpai, A. Brunstrom, J. Ott, M. Trevisan, M. Mellia, and G. Fairhurst, "Measuring Roaming in Europe: Infrastructure and Implications on Users' QoE," *IEEE Transactions on Mobile Computing*, vol. 21, no. 10, pp. 3687–3699, Oct. 2022.

[2] M. Motallebighomi, J. Veara, E. Bitsikas, and A. Ranganathan, "eSIMplicity or eSIMplification? Privacy and Security Risks in the eSIM Ecosystem," in *Proceedings of the 34th USENIX Security Symposium*. USENIX, Aug. 2025.

[3] Y. Li, X. Ma, M. Xu, A. Zhou, Q. Sun, N. Zhang, and S. Wang, "Joint placement of upf and edge server for 6g network," *IEEE Internet of Things Journal*, vol. 8, no. 22, pp. 16 370–16 378, 2021.

[4] E. Goshi, H. Harkous, S. Ahvar, R. Pries, F. Mehmeti, and W. Kellerer, "Joint UPF and Edge Applications Placement and Routing in 5G & Beyond," in *2024 IEEE MASS*, 2024, pp. 446–452.

[5] E. Goshi, E. N. Şen, H. Harkous, S. Ahvar, R. Pries, F. Mehmeti, and W. Kellerer, "Profitable Energy-Efficient UPF and Application Placement and Routing in Dynamic 5G Edge," *IEEE Transactions on Network Science and Engineering*, vol. 13, pp. 5698–5720, 2026.

[6] S. Lee, "Open5GS: Open Source implementation for 5G Core and EPC," Feb. 2026. [Online]. Available: https://github.com/open5gs/open5gs

[7] V. D'Emmanuele and A. Raguideau, "PacketRusher: High performance 5G UE/gNB Simulator and CP/UP load tester," Nov. 2023.

[8] R. A. K. Fezeu, C. Fiandrino, E. Ramadan, J. Carpenter, D. Chen, Y. Tan, F. Qian, J. Widmer, and Z.-L. Zhang, "Roaming across the European Union in the 5G Era: Performance, Challenges, and Opportunities," in *IEEE INFOCOM 2024 - IEEE Conference on Computer Communications*, May 2024, pp. 2378–2387.

[9] K. Trichias, S. Col, I. Masmanidis, A. Berisha, F. Setaki, P. Demestichas, S. Papavassiliou, and N. Mitrou, "5g for connected and automated mobility - Network level evaluation on real neighboring 5G networks: The Greece - Turkey cross border corridor," *Computer Communications*, vol. 232, p. 108047, 2025.

[10] R. Keller, D. Castellanos, A. Sander, A. Robison, and A. Abtin, "Roaming in the 5G System: The 5GS Roaming Architecture," *Ericsson Technology Review*, vol. 2021, no. 6, pp. 2–11, Jun. 2021.

[11] A. Bhuyan and M. Gheorghevici, "Real-world Demonstration of Secure 5G Roaming for Military Communications," in *MILCOM 2024 - 2024 IEEE Military Communications Conference (MILCOM)*, 2024, pp. 639–640.

[12] F. Vázquez-Gallego, J. Nasreddine, M. Codina, B. Cordero, E. Carmona-Cejudo, M. Trullenque, D. Camps-Mur, Y. Murillo, P. Seguret, J. Polo, and J. L. Luque, "Performance Evaluation of 5G Standalone Seamless Home Routed Roaming for Connected Mobility in Cross-Border Scenarios," in *2024 IEEE 99th Vehicular Technology Conference (VTC2024-Spring)*, Jun. 2024, pp. 1–7.

[13] I. W. A. J. Pawana, V. Abella, J. K. Lastre, Y. Ko, and I. You, "Enhancing Roaming Security in Cloud-Native 5G Core Network through Deep Learning-Based Intrusion Detection System," *Computer Modeling in Engineering & Sciences*, vol. 145, no. 2, pp. 2733–2760, 2025.

[14] Q. Tang, O. Ermis, C. D. Nguyen, A. D. Oliveira, and A. Hirtzig, "A Systematic Analysis of 5G Networks With a Focus on 5G Core Security," *IEEE Access*, vol. 10, pp. 18 298–18 319, 2022.

[15] G. M. Køien, "On Threats to the 5G Service Based Architecture," *Wireless Personal Communications*, vol. 119, no. 1, pp. 97–116, Jul. 2021.

[16] M. Mahyoub, A. AbdulGhaffar, E. Alalade, E. Ndubisi, and A. Matrawy, "Security Analysis of Critical 5G Interfaces," *IEEE Communications Surveys & Tutorials*, 2024.

[17] 3GPP, "3GPP TS 33.501 V18.6.0 - Security architecture and procedures for 5G System," Jul. 2024.

[18] GSMA, "GSMA NG.113-v9.0 Roaming Guidelines," Feb. 2024.

[19] 3GPP, "3GPP TS 29.573 V19.2.0 - Public Land Mobile Network (PLMN) Interconnection," Mar. 2025.

[20] O. Zeidler, D. Fraunholz, J. Sturm, H. König, and W. Kellerer, "Performance Evaluation of 5G Roaming Security Based on PRINS," in *2025 IEEE CNS*, Sep. 2025, pp. 1–10.

[21] O. Zeidler, M. Günes, S. Anirudh M., and W. Kellerer, "DEMO: 5G SA Roaming Testbed for Post Quantum IPsec Security Evaluations," in *ACM WiSec'26*, Jun. 2026.

[22] O. Zeidler, J. Sturm, D. Fraunholz, and W. Kellerer, "Performance Evaluation of Transport Layer Security in the 5G Core Control Plane," in *ACM WiSec'24*, May 2024, pp. 78–88.

[23] WonderNetwork, "Global Ping Statistics," 2026. [Online]. Available: https://wondernetwork.com/pings/

[24] 3GPP, "3GPP TS 23.502 V16.19.0 - Procedures for the 5G System," Apr. 2024.

[25] G. Amponis, P. Radoglou-Grammatikis, T. Lagkas, W. Mallouli, A. Cavalli, D. Klonidis, E. Markakis, and P. Sarigiannidis, "Threatening the 5G core via PFCP DoS attacks: The case of blocking UAV communications," *EURASIP Journal on Wireless Communications and Networking*, vol. 2022, no. 1, p. 124, Dec. 2022.

[26] Y. Cao, Y. Chen, and W. Zhou, "Detection of PFCP Protocol Based on Fuzz Method," in *2022 2nd International Conference on Computer, Control and Robotics (ICCCR)*, Mar. 2022, pp. 207–211.

[27] N. Makondo, E. Baloyi, H. I. Kobo, and T. E. Mathonsi, "A Review of PFCP Cyber Attacks in 5G Standalone for Robotic Telesurgery Services," in *NextComp 2024*, Oct. 2024, pp. 1–7.

[28] G. Amponis, P. Radoglou-Grammatikis, G. Nakas, S. Goudos, V. Argyriou, T. Lagkas, and P. Sarigiannidis, "5G Core PFCP Intrusion Detection Dataset," in *2023 12th International Conference on Modern Circuits and Systems Technologies (MOCAST)*, Jun. 2023, pp. 1–4.

[29] R. Pell, M. Shojafar, and S. Moschoyiannis, "LSTM-Based Anomaly Detection of PFCP Signaling Attacks in 5G Networks," *IEEE Consumer Electronics Magazine*, vol. 14, no. 1, pp. 56–64, Jan. 2025.