Evaluation of Security-Induced Latency on 5G RAN Interfaces and User Plane Communication

Sotiris Michaelides
michaelides@spice.rwth-aachen.de
RWTH Aachen University
Aachen, Germany

Jakub Lapawa
jakub.lapawa@rwth-aachen.de
RWTH Aachen University
Aachen, Germany

Daniel Eguiguren Chavez
daniel.eguiguren@rwth-aachen.de
RWTH Aachen University
Aachen, Germany

Martin Henze*
henze@spice.rwth-aachen.de
RWTH Aachen University
Aachen, Germany

Abstract

5G promises enhanced performance—not only in bandwidth and capacity, but also latency and security. Its ultra-reliable low-latency configuration targets round-trip times below 1 ms, while optional security controls extend protection across all interfaces, making 5G attractive for mission-critical applications. A key enabler of low latency is the disaggregation of network components, including the RAN, allowing user-plane functions to be deployed nearer to end users. However, this split introduces additional interfaces, whose protection increases latency overhead. In this paper, guided by discussions with a network operator and a 5G manufacturer, we evaluate the latency overhead of enabling optional 5G security controls across internal RAN interfaces and the 5G user plane. To this end, we deploy the first testbed implementing a disaggregated RAN with standardized optional security mechanisms. Our results show that disaggregated RAN deployments retain a latency advantage over monolithic designs, even with security enabled. However, achieving sub-1 ms round-trip times remains challenging, as cryptographic overhead alone can already exceed this target.

CCS Concepts

• **Security and privacy** → **Security protocols**; **Mobile and wireless security**; • **Networks** → *Mobile networks.*

Keywords

5G, RAN, URLLC, Security, Latency, IPsec, TLS

ACM Reference Format:
Sotiris Michaelides, Jakub Lapawa, Daniel Eguiguren Chavez, and Martin Henze. 2026. Evaluation of Security-Induced Latency on 5G RAN Interfaces and User Plane Communication. In *Proceedings of the 19th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '26), June 30-July 03, 2026, Saarbrücken, Germany.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3765613.3797453

*Also with Fraunhofer FKIE.

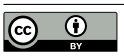

WiSec '26, Saarbrücken, Germany

ACM ISBN 979-8-4007-2201-1/2026/06
https://doi.org/10.1145/3765613.3797453

1 Introduction

The fifth generation of mobile networks (5G) introduces a major architectural transition, shifting from hardware-centric, monolithic deployments to software-defined, cloud-native designs [2]. This evolution enables 3GPP, for the first time, to extend its focus beyond traditional consumer connectivity and to address industrial and mission-critical sectors that demand ultra-low latency [13].

In such sectors, the modular and cloud-native architecture of 5G—enabled by the disaggregation of the 5G Core (5GC) and Radio Access Network (RAN) into individual network functions—allows latency-critical User Plane (UP) components (e.g., the Control Unit-User Plane (CU-UP) and User Plane Function (UPF)) to be deployed closer to end devices, thereby reducing latency [13, 17, 18] and enabling 5G to target Round-Trip Times (RTT) as low as 1 ms [9].

Beyond low latency and high reliability, these critical sectors also demand strong security measures as their importance in societal, economic, and public safety-critical functions makes them a frequent target for powerful attacks [12, 19]. Consequently, as 5G is being deployed in these environments—expanding the attack surface through new interfaces and components—strong security controls must be implemented to mitigate resulting threats.

Fortunately, 5G introduces a comprehensive suite of security controls across its architecture, covering nearly every component and interface. Although each component is required to support these mechanisms, the activation of most remains optional, granting network operators flexibility [11, 13, 14, 20]. As enabling these controls inevitably introduces cryptographic operations that add latency, it remains open whether 5G can still achieve its promised sub-1 ms RTT targets while maintaining robust security.

The importance of this issue is highlighted by the extensive prior work examining security controls within the 5GC, N3, and Uu interfaces [11, 14, 20]. While these studies provide insights into individual controls, they focus on isolated interfaces without considering the broader implications of RAN disaggregation, which is critical for low-latency communication [13, 18]. Thus, there is a need to evaluate security controls across the entire 5G architecture and their overall impact on system performance.

In this paper, we present the first comprehensive evaluation of optional security controls across the 5G UP, with a focus on disaggregated RAN deployments, crucial for achieving ultra-low latency [13, 18]. Our work is guided by discussions with a Mobile Network Operator (MNO) and a 5G vendor, providing practical

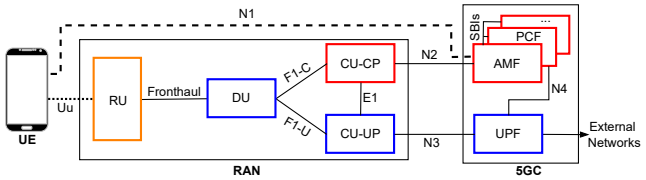

Figure 1: The modularity of 5G enables time-critical components to be placed closer to end devices on the edge cloud infrastructure to reduce latency. Control-plane components are hosted on a remote and centralized cloud to simplify network control. The Radio Unit (RU) is always deployed on-site

insights into the benefits of disaggregated RAN deployments and the use of optional security controls in real networks, ensuring that our findings have direct practical relevance. More specifically, our key contributions are:

(1) Practical insights from an MNO and a 5G vendor on security and performance in disaggregated RAN deployments, offering guidance for secure low-latency deployments (§3).
(2) The first *open-source* 5G testbed supporting all optional internal RAN security controls (§4); available at [1].
(3) Evaluation of the impact of security on latency across RAN interfaces for both control and user planes (§5), identifying IPsec as the most latency-efficient option, with configurations adding less than 60 µs of overhead.
(4) End-to-end UP latency analysis from User Equipment (UE) to UPF in monolithic and disaggregated setups (§6), showing that security alone already pushes RTTs beyond 1 ms.

2 5G Networks and Security Controls

From a high-level view, 5G follows the same basic structure as earlier mobile networks, consisting of the UE, the 5GC, and the RAN (cf. Fig. 1). It also maintains a split between the Control Plane (CP), which handles signaling such as session management and mobility, and the User Plane (UP), which handles user-data forwarding [14]. This section outlines the 5G architecture (§2.1), the main interfaces and their associated security controls (§2.2), and the optional security mechanisms examined in related work (§2.3).

2.1 Architectural Components

The UE is the end device that accesses network services provided by the 5GC using authentication credentials [3]. The 5GC consists of interconnected Network Functions (NFs) responsible for control and user-plane operations: CP functions such as the Authentication and Mobility Function (AMF) and Session Mangement Function (SMF) manage authentication, mobility, and session control, while the UPF on the UP handles routing of user data [3]. The RAN mediates between the UE and the 5GC by providing the wireless connectivity, and can be deployed either monolithically or disaggregated [6].

In monolithic RAN deployments, the Baseband Unit (BBU) implements the full protocol stack except radio transmission, handled by the Radio Unit (RU). In a disaggregated deployment, the BBU is split into Control Unit (CU) and Distributed Unit (DU) [6]. The DU handles real-time lower-layer functions, including scheduling, signal processing, modulation, coding, and beamforming, while the CU manages higher-layer functions such as resource control, security, intra-RAN mobility, and DU management. The CU is split into Control Unit-Control Plane (CU-CP) and Control Unit-User Plane (CU-UP), reflecting control- and user-plane separation, allowing independent scaling, optimization, and placement of each plane. Placing UP functions and the DU on the edge cloud (blue in Fig. 1) near the UE is crucial for low-latency communication (cf. §3), while keeping CP centralized (red in Fig. 1) enables better control.

2.2 5G Interfaces and Security

Following the control–user plane separation, we examine the 5G interfaces and their associated security controls.

Control Plane: CP data is exchanged between network nodes to maintain and support 5G operations. This includes signaling between the UE and AMF over the N1 interface for authentication and mobility management, as well as signaling over the Uu interface between the UE and the RAN for requesting and allocating radio resources. These messages are mandatorily integrity-protected using the NIA algorithms based on SNOW, AES, or ZUC [5, 11], and may optionally be encrypted using the corresponding NEA schemes. Within the RAN, control signaling occurs over the F1-C and E1 interfaces. These interfaces support operations such as intra-RAN handovers, the exchange of CP data to and from the UEs, radio resource management, radio resource control forwarding, and UP control commands. Control signaling between the RAN and the 5GC takes place over the N2 interface, which carries handover messages, mobility management information, and session-related control data between the CU-CP and the AMF. All three interfaces are optionally protected using IPsec or DTLS [5, §9.8.2, §9.8.3, §9.2]. Within the 5GC, CP data is exchanged between NFs for tasks such as session management and retrieving authentication credentials. Data over the Service Based Interfaces (SBIs) can be optionally secured with TLS or IPsec [5, §13.1.0], and with IPsec over the N4 interface [5, §9.9]. While most security controls are optional *"because they can be replaced with physical security"* [14], disaggregated RAN deployments—where physical security is infeasible—often rely on these controls (cf. §3). Given that CP security is critical for the network operation, we evaluate how security controls within the RAN affect crucial CP operations of the 5G system (§5).

User Plane: The UP is responsible for carrying the user data from the UE to the UPF. Between the UE and the RAN (CU-UP), this data is optionally secured using the NIA/NEA schemes [5, §5.2]. While these controls also cover data over the F1-U interface between the CU-UP and DU, security on this interface is crucial, as F1-U not only transports UP data but also carries control traffic used by the DU to regulate downlink data and prevent buffer overflows [7, §5.3]. IPsec is defined as an optional security control for this interface. The same applies to the N3 interface, which carries UP data between the CU-UP and the UPF [5, §9.3]. Security over UP interfaces adds latency and may affect 5G's low-latency capabilities. In this paper, we first assess the impact of F1-U security in isolation (§5), as it is the only UP interface not studied in prior work (cf. §2.3), before evaluating all UP security controls collectively (§6).

The special case of the Fronthaul (FH) Interface: The FH carries CP and UP data and allows the DU to manage RUs (e.g., synchronization). Defined by the Common Public Radio Interface

Table 1: While prior work examined the impact of various controls on isolated interfaces, our work is the first to evaluate the impact of security on (1) RAN interfaces and (2) UP interfaces collectively to assess end-to-end latency overhead.

-	Interf.	Security	[8]	[11]	[14]	[20]	This
Control Plane	N1&Uu	NEA$^{\beta}$	-	-	-	-	-
		NIA$^{\alpha}$	-	-	-	-	-
	SBIs	TLS$^{\beta}$	-	-	✓	✓	-
		IPsec$^{\beta}$	-	-	✓	-	-
	N2	DTLS$^{\beta}$	-	-	-	-	-
		IPsec$^{\beta}$	-	-	-	-	-
	N4	IPsec$^{\beta}$	-	-	-	-	-
	E1	DTLS$^{\beta}$	-	-	-	-	✓
		IPsec$^{\beta}$	-	-	-	-	✓
	F1-C	IPsec$^{\beta}$	-	-	-	-	✓
		DTLS$^{\beta}$	-	-	-	-	✓
User Plane	F1-U	IPsec$^{\beta}$	-	-	-	-	✓+
	Uu	NEA$^{\beta}$	-	-	-	-	+
		NIA$^{\beta}$	-	✓	-	-	+
	N3	IPsec$^{\beta}$	-	-	✓	-	+
Both	FH	MACsec $^{\gamma}$	✓	-	-	-	-
		IPsec $^{\gamma}$	-	-	-	-	-

$^{\alpha}$Mandatory Usage $^{\beta}$Optional Usage $^{\gamma}$Recommended Support
+ Combined Evaluation ✓Isolated Evaluation

(CPRI) consortium rather than 3GPP, proprietary CPRI links up to LTE transmitted analog signals, with security only feasible at the hardware level. In 5G, enhanced CPRI (eCPRI) gradually replaces CPRI with a packet-based interface that enables partial distribution of DU functions to the RU; however, implementing security controls (e.g., MACsec) is only recommended, not mandated, by the specification. As the FH lacks standardized security controls and has diverse vendor-specific implementations, we exclude it from our evaluation, though we reference related work.

2.3 Related Work

The studies summarized in Tab. 1 highlight the current state of research on optional security controls in 5G systems and motivate our work. Heijligenberg et al. [11] were the first to examine the impact of optional security controls, evaluating NIA schemes on the UP in terms of bandwidth and latency and showing a noticeable yet comparable impact across all schemes. Zeidler et al. [20] analyzed TLS as the primary security mechanism over the SBIs, reporting minimal performance degradation in an operational 5G system but significant delays during initial security establishment. Similarly, in prior work [14], we investigated IPsec within the 5GC and across the N3 interface, demonstrating that IPsec can outperform TLS due to its ability to pre-establish security associations, while introducing only minimal latency on N3. Lastly, Dik et al. [8] examined MACsec (although not standardized) on the FH, showing a negligible impact on latency. Overall, these studies reveal the absence of an architecture-wide evaluation of the UP and a lack of consideration for disaggregated RAN deployments—vital for achieving low latencies—thereby underscoring the gap filled by our contributions.

3 Insights on Disaggregated RANs

To guide our research on disaggregated RAN architectures and optional security controls in real-world 5G deployments, we consulted a major 5G equipment manufacturer and a European MNO (who requested to not be named). These discussions highlight the practical benefits of such deployments (§3.1) and the use of optional security controls (§3.2), underscoring the relevance of this work.

3.1 Practical Benefits of Disaggregated RANs

During our discussions with the MNO, we learned that most current RANs deployments are monolithic due to legacy infrastructure. However, they emphasized that new RAN deployments are disaggregated and virtualized to improve cost efficiency, availability, and scalability. As they explained, it is more convenient to dynamically deploy instances of the DU or CU rather than the entire stack, significantly reducing costs. These deployments also allow the MNO to quickly identify and isolate faulty instances and redeploy replacements, minimizing downtime—a top operational priority [14]. Finally, the MNO, which is currently in the process of deploying a low-latency network slice, pointed out that internal measurements indicate that without physically reducing the distance between UP components and end-users, achieving low-latencies is infeasible.

3.2 Security in Real-World 5G Networks

The equipment manufacturer we contacted is also responsible for configuring networks according to the requirements of clients, which include both MNOs and private companies. The manufacturer reported an increasing demand for disaggregated and virtualized RANs, which aligns with findings from Dell'Oro (and Nokia) [15], projecting that by 2028, 20% of RANs will be disaggregated and virtualized. When asked about security, the manufacturer noted that optional security controls are typically not used in traditional 5G deployments due to bandwidth considerations and potential performance bottlenecks especially on the UP, which bears more traffic. In contrast, optional security mechanisms are almost always employed in disaggregated RAN deployments, as these components are often geographically distributed across cloud infrastructure where physical security is not feasible (cf. §2.2).

Our discussions highlight the benefits of disaggregated RAN deployments, their role in low latency, and the use of optional security controls in real-world 5G networks, demonstrating the relevance and real-world applicability of evaluating security-induced latency on 5G RAN interfaces and UP communication.

4 Deploying a Disaggregated-RAN 5G Testbed

To lay the foundation for such evaluations, we first identify optional security controls applied to the RAN interfaces (§4.1), before deploying them by extending an existing state-of-the-art testbed (§4.2).

4.1 Identified Optional Security Controls

To identify optional security controls, we examine the relevant 3GPP specifications and determine the mandatory-to-support configurations, as these ensure a degree of compatibility between devices from different manufacturers. For IPsec, we leverage the same controls as those we discovered in our previous work on evaluating tunneling protocols over the N3 and the SBIs [14]. These comprise

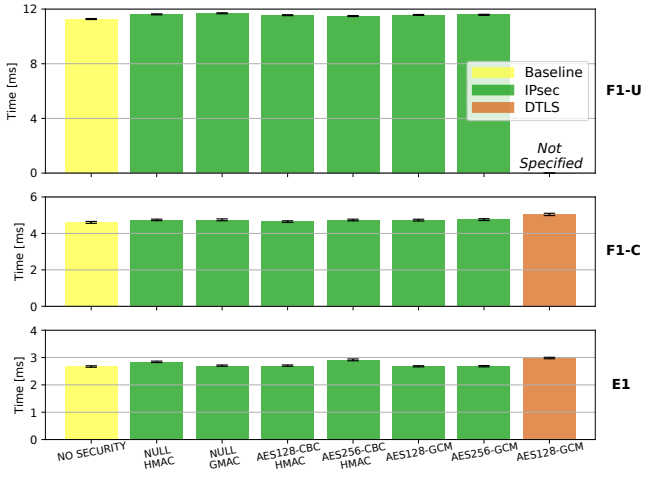

Figure 2: Securing internal-RAN interfaces using IPsec adds only minimal latency overhead regardless of the configuration, making it suitable to support low-latencies.

six different algorithm combinations, which can be used with either certificates or pre-shared keys. For DTLS, we extract three configurations from [4, §6.2]. All three configurations employ the same Authenticated Encryption with Associated Data (AEAD) algorithm for encryption and integrity protection (AES-GCM with 128-bit keys) but differ in their key exchange and authentication mechanisms. Interestingly, while standard TLS configurations mandate support for version 1.3, DTLS only requires version 1.2.

During our experiments, the RAN components establish their security (i.e., authentication and key exchange) once, during the discovery phase at deployment. Thus, our measurements focus on the overhead of encryption and integrity protection, evaluating six IPsec ESP configurations and one DTLS configuration (cf. Fig. 2). This mirrors real-world 5G systems, where the network runs continuously and security associations are thus rarely re-established.

4.2 Testbed and Experimental Setup

To deploy and study the identified security protocols, we extend our previously presented open-source testbed [14], which implements IPsec over the N3 interface using UERANSIM and Open5GS. As UERANSIM lacks RAN disaggregation support and a full 5G protocol stack, we replace it with Open Air Interface (OAI) [10], a widely adopted in academia and industry. This upgrade not only supports a fully disaggregated RAN, but also allows seamless integration with software-defined radios, enabling both simulated and over-the-air deployments. Each network function—including the 5GC, disaggregated RAN functions, and the UE—runs as an independent Docker container, interconnected via dedicated subnetworks that emulate 3GPP interfaces. The F1-C, F1-U, and E1 links are represented as separate Docker bridges, allowing their precise monitoring.

For consistency with the original setup, we deploy IPsec over the F1-C, F1-U, and E1 interfaces using strongSwan, establishing IKEv2/ESP tunnels automatically at container startup. DTLS is implemented over the F1-C and E1 interfaces using the lightweight socat utility, which provides bidirectional DTLS tunnels between RAN components. Our setup also allows dynamic switching between security configurations and authentication methods via environment variables, making it easily adaptable. The complete testbed runs on a physical host with Ubuntu 22.04.5 LTS on an AMD Ryzen 7 Pro 6850U (8 cores, 16 threads) with an integrated Rembrandt GPU, 32 GB DDR5 RAM, and Linux kernel 6.8.0-60-generic. The CPU supports AES-NI for AES acceleration, which we leverage to reflect realistic performance in modern 5G deployments.

5 Impact of Security on Internal-RAN Interfaces

To assess the impact of security on RAN interfaces, we use our testbed to perform reproducible CP and UP tests, with and without security to quantify the overhead. We first outline the methodology we follow for each plane (§5.1) before presenting our results (§5.2).

5.1 Methodology

We use different methodologies for the two planes as we can directly generate traffic on the UP, whereas we cannot do so on the CP.

User Plane: To assess the latency impact on the UP over the F1-U interface, we generate user data. We employ the *ping* tool, which uses ICMP messages to measure the RTT between two hosts, following the methodology of related work [11, 14]. From within the UE container, we ping the UPF container, ensuring that the traffic remains confined within the same physical host to minimize external factors that may affect latency. We employ 1024-byte payloads to amplify the effect of cryptographic operations and repeat the experiment 20,000 times for statistical confidence.

Control Plane: In contrast to the UP, we cannot directly generate traffic on the CP interface, as it is triggered by control procedures such as authentication and handovers. Therefore, we use *UE registration*, one of the most common CP operations [20], to trigger CP activity. To this end, we deploy a UE that automatically sends a registration request. While we do not monitor the registration data itself, we capture the CP sub-procedures triggered during authentication to establish connectivity between the UE and the UPF. Over the F1-C, we monitor the UE Context Setup procedure, which establishes the UE context and provides the DU with the information needed to manage the UE. We measure the time from the UE Context Setup request sent by the CU-CP until the corresponding response, which indicates that the DU is ready to handle the UE. For the E1 interface, we monitor the E1 Bearer Context Setup procedure, responsible for allocating resources and establishing GTP tunnels between the UE and the UPF over CU-UP. We measure the time between the Bearer Context Setup request from the CU-CP to the CU-UP and the corresponding response, confirming the successful establishment of GTP-U tunnels for the F1-U and N3 interfaces. To ensure statistical confidence, we repeat the experiment 1,000 times.

5.2 Results

Our results (cf. Fig. 2) are shown as 99% confidence intervals of the average latency across all repetitions, providing statistical confidence in the true mean. Across all interfaces, the fastest IPsec configuration adds only minimal overhead on the order of tenths of microseconds, whereas DTLS consistently incurs higher latency.

User Plane: Over the F1-U interface, where only IPsec is specified, all IPsec configurations combined add on average ~120 µs

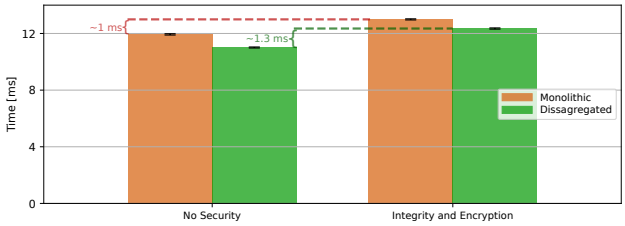

Figure 3: While the monolithic deployment benefits from lower security overhead due to fewer controls needing to be enabled, the disaggregated approach still outperforms it due to the independent and efficient handling of UP data

RTT (~60 µs one-way), with minimal variation across cipher suites. Differences between cipher suites are indistinguishable, due to overlapping confidence intervals. These results closely align with related work on the UP (over N3), where the same IPsec algorithms were shown to add on average 55 µs, also with overlapping confidence intervals [14]. This indicates that IPsec is a strong candidate for secure, low-latency UP data transmission across the 5G architecture, as the added latency is well below the 1 ms threshold.

Control Plane: On the CP, specifically the F1-C interface, all IPsec configurations add on average ~120 µs to UE context setup, compared to ~430 µs for DTLS. The fastest IPsec configuration (AES256-HMAC) adds only ~40 µs, making it nearly indistinguishable from unprotected traffic. Over the E1 interface, IPsec configurations introduce on average ~80 µs of overhead, whereas DTLS adds ~300 µs. However, IPsec with AES128-GCM performs best, adding only ~10 µs. From a practical perspective, both protocols can be used for procedures that are not latency-critical—such as UE authentication—without noticeable impact. For latency-critical operations, such as intra-RAN handovers, IPsec remains the preferred option due to its consistently lower overhead.

6 User Plane-Wide Evaluation

Having deployed and evaluated security within the RAN, we now assess its overall impact on latency when enabled across the entire 5G UP, from the UE to the UPF. We first describe our methodology (§6.1), present results with 99% confidence intervals (§6.2), and report tests conducted in a realistic 5G deployment (§6.3).

6.1 Methodology

To evaluate the impact across the 5G UP, we follow a similar approach as for the F1-U measurements. We perform 20,000 ping measurements between the UE and the UPF using 1024-byte packets. We assess latency overhead by comparing RTTs of a baseline without security to a deployment with security enabled on all UP links. We consider both a monolithic RAN, securing Uu and N3, and a disaggregated RAN, additionally securing F1-U.

To secure the interfaces, we have two options: either enable only integrity protection or enable both integrity and encryption. Encryption must always be integrity-protected for some interfaces (e.g., N3 [14]) and, by itself, provides no protection against tampering attacks [16]. However, we observed that enabling only integrity protection provides no measurable advantage over enabling both. Therefore, we choose to enable both encryption and integrity protection on all interfaces for enhanced security. Over N3 and F1-U, we use IPsec with AES-GCM128, which provides both integrity and encryption, with performance comparable to schemes that only provide integrity protection. The same applies to the Uu interface in which we measure no difference when enabling NIA and NEA simultaneously compared to using only NIA. There, we utilize NIA&NEA option 2 based on AES, as it has been shown to outperform the other two schemes in runtime and throughput [11].

6.2 Results

By examining the 99% confidence intervals in Fig. 3, we observe that disaggregated RAN deployments achieve better baseline performance than monolithic ones (11 ms vs. 12 ms), due to the independent handling of UP data—even without being physically closer to the end device. This confirms that disaggregated architectures are inherently well-suited for low-latency communication. Even with security enabled, disaggregated RAN maintains its performance advantage despite the additional interfaces to protect. Cryptographic overhead exceeds 1 ms in both cases—approximately 1 ms for the monolithic setup and 1.3 ms for the disaggregated deployment. Given minimal CPU load and comparable memory usage, we attribute the measured overhead primarily to cryptographic operations. While disaggregated RAN remains faster overall, the high cost of cryptographic operations shows that achieving sub-1 ms latency is extremely challenging for both RAN deployments, highlighting the combined effect of multiple optional security controls as a key limiting factor for ultra-low-latency applications.

6.3 Validation in a Realistic 5G Deployment

To ensure reproducibility in an isolated setup, our main results are obtained using a simulated radio and UE, where both disaggregated and monolithic RAN deployments are equally close to the UE, minimizing the benefits of disaggregation. To validate these results in realistic 5G deployments, we repeat experiments using a UE with a MediaTek Dimensity 700 baseband, a sysmocom ISIM-SJA5 SIM card, and a USRP B210, considering two scenarios differing in UP placement. In the disaggregated deployment, the UPF and CU-UP are colocated with the DU on one host, placing UP functions closer to the UE, while the CP resides on a separate host connected via a network switch. In the centralized deployment, the DU remains local, while the UP components are colocated with the CP on a second host, introducing "distance" from the UE, resembling a monolithic setup. The disaggregated scenario achieves 2.5 ms and 2.1 ms lower latency without and with security enabled, respectively, highlighting the benefits of disaggregation for low-latency communication and validating the trends observed in our main experiments.

7 IPsec Configuration Benchmarking

While we focus on the latency of security controls, throughput is of equal concern to MNOs, especially on the UP (cf. §3.2). Related work has studied NIA schemes, showing linear scaling with input size and identifying NIA2 as most efficient in throughput and runtime [11]. We complement these studies by benchmarking the impact of IPsec configurations using the Python Cryptography library.

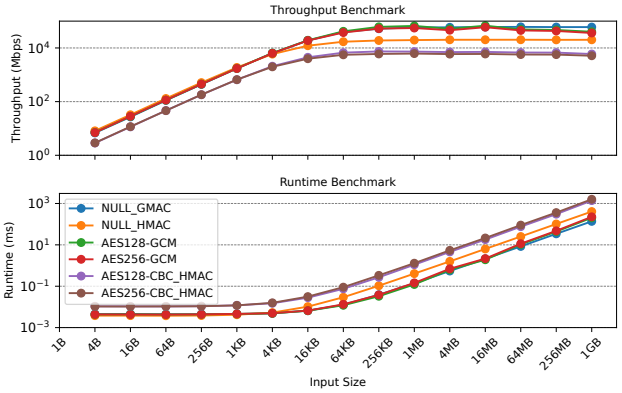

Figure 4: Our benchmarking of IPsec configurations confirms the linear scalability of all schemes and highlights the superior performance of GMAC and GCM schemes.

As shown in Fig. 4, all ESP schemes scale linearly in throughput before capping at an upper bound. GMAC achieves the highest throughput, reaching up to 7.4 GB s^{-1} for large data. The GCM schemes follow closely, averaging 6 GB s^{-1} to 6.5 GB s^{-1}, whereas CBC schemes remain the slowest, never exceeding 1 GB s^{-1}. These differences are reflected in runtime measurements: encrypting 1 GB of data requires 0.14 s with GMAC AES-128, 0.21 s to 0.23 s with AES-GCM, and 1.37 s to 1.61 s with AES + HMAC, highlighting a 6 × to 10 × runtime advantage for GCM-based schemes. The performance differences align with the underlying cryptographic constructions: GMAC is lightweight, AES-GCM combines encryption and GMAC in a parallelizable manner, while AES-HMAC is sequential and constrained by data dependencies. Consequently, operators can expect lower resource consumption and higher throughput when deploying GMAC or AES-GCM schemes.

8 Discussion

Finally, we discuss the broader implications of our results, including their transferability to real-world 5G systems, the underlying the root cause of high latencies, and potential avenues for improvement.

Transferability of Results: Even in our small-scale 5G network, latency already exceeds 1 ms, and overheads are likely to be higher in larger, real-world deployments with increased load. Our measurements show overheads in the order of tenths of microseconds—extremely fast—so further gains from more powerful machines would be minimal. Additional security controls, such as those over the FH or other non-parallelizable schemes (e.g., NIA1/NEA1 or AES-CBC_HMAC), would further increase latency. Lastly, in our experiments, we ping the UPF, but in real 5G networks, where communication might occur between two UEs, cryptographic actions would double, effectively also doubling the overhead. As such, our conclusion that achieving sub-1 ms RTTs is challenging is broadly transferable to real-world 5G deployments.

NIA/NEA as the Limiting Factor: Although IPsec achieves very low latencies using fast, parallelizable schemes such as AES-GCM, the overall end-to-end latency remains high. We attribute this high latency to the 5G NIA/NEA algorithms, which must be executed sequentially in a MAC-then-encrypt fashion [11], inherently introducing additional latency. While NEA2 is based on AES-CTR, which is parallelizable, NIA2 relies on CMAC mode, which is sequential, introducing data dependencies, further increasing latency.

Reducing Latency: Latency can be reduced by minimizing cryptographic operations on UP data. Approaches include security-aware traffic separation (e.g., avoiding re-encryption of UP data over F1-U) or true end-to-end security between UE and UPF. Lastly, using AEAD algorithms (e.g., AES-GCM or ChaCha20-Poly1305) in NIA/NEA schemes can further lower latency.

9 Conclusion

With this work, we underscore the importance of RAN disaggregation to meet low latencies in ultra-low-latency scenarios such as industrial networks. Guided by discussion with a real-world MNO and a RAN manufacturer, using a disaggregated RAN 5G setup, we show that the security impact on its interfaces can be kept to a minimum with IPsec, and—despite additional interfaces—it still outperforms its monolithic counterpart in an end-to-end UP evaluation. However, our results also reveal that achieving sub-millisecond RTTs remains difficult, as the cumulative cryptographic overhead already exceeds this threshold, even in a small-scale network.

Acknowledgments

Funded by the German Federal Ministry of Research, Technology and Space (BMFTR) under funding reference number 16KIS2409K (6GEM+) and the German Federal Office for Information Security (BSI) under funding reference number 01MO24003B (CSII). The authors are responsible for the content of this publication.

References

[1] 2025. https://github.com/RWTH-SPICe/Internal-Ran-Security.
[2] 3GPP. 2023. *TS 38.300 V18.7.0; NR; NR and NG-RAN Overall Description; Stage-2.*
[3] 3GPP. 2024. TS 23.501 V18.7.0: System Architecture for the 5G System (5GS).
[4] 3GPP. 2024. TS 33.210 V18.1.0: Network Domain Security (NDS); IP Network Layer Security.
[5] 3GPP. 2024. TS 33.501 V18.7.0: Security Architecture and Procedures for 5G System.
[6] 3GPP. 2024. TS 38.401 V18.7.0: 5G NG-RAN Architecture description.
[7] 3GPP. 2024. TS 38.470 V17.6.0: F1 general aspects and principles.
[8] Daniel Dik and Michael Stübert Berger. [n. d.]. Open-RAN Fronthaul Transport Security Architecture and Implementation. *IEEE Access.*
[9] Aoyu Gong et al. 2025. Towards URLLC with Open-Source 5G Software. In *ACM OpenRIT6G '25 Workshop.*
[10] OAI group. 2025. OpenAirInterface 5G Radio Access Network Project. https://openairinterface.org/oai-5g-ran-project/
[11] Thijs Heijligenberg et al. 2023. BigMac: Performance Overhead of User Plane Integrity Protection in 5G Networks. In *ACM WiSec.*
[12] Martin Henze et al. 2013. Maintaining User Control While Storing and Processing Sensor Data in the Cloud. *IJGHPC.*
[13] Sotiris Michaelides et al. 2025. Secure Integration of 5G in Industrial Networks: State of the Art, Challenges and Opportunities. *FGCS.*
[14] Sotiris Michaelides et al. 2025. Assessing the Latency of Network Layer Security in 5G Networks. In *ACM WiSec.*
[15] Nokia. 2024. Cloud RAN: Unlocking the power of 5G for a smarter, connected future. https://www.nokia.com/asset/i/214269/
[16] David Rupprecht et al. 2019. Breaking LTE on Layer Two. In *IEEE SP.*
[17] Asher Sajid et al. 2025. Towards SMPC-Enabled O-RAN: A Survey with Deployment-Oriented Insights. *Computer Networks.*
[18] Samsung. 2019. Virtualized Radio Access Network. https://www.samsung.com/global/business/networks/insights/white-papers/virtualized-radio-access-network/
[19] Ioannis Stellios et al. 2018. A Survey of IoT-Enabled Cyberattacks: Assessing Attack Paths to Critical Infrastructures and Services. *IEEE Commun. Surv. Tutor.*
[20] Oliver Zeidler et al. 2024. Performance Evaluation of Transport Layer Security in the 5G Core Control Plane. In *ACM WiSec.*