End-to-End Network Slicing: Securing Sensitive Data Across the Network

Xhulio Limani*, Arno Troch*, Michiel Van Kenhove‡, Alexandra Papageorgiou§
Filip De Turck‡, Erik Pohle†, Leonard Schild†, Martin Zbudila†, Aysajan Abidin†, Bruno Volckaert‡
Miguel Camelo Botero*, Johann M. Marquez-Barja*, and Nina Slamnik-Kriještorac*
*University of Antwerp - imec, IDLab - Faculty of Applied Engineering, Belgium
‡IDLab, Department of Information Technology, Ghent University - imec, Ghent, Belgium
§Centre for IT and IP Law, †KU Leuven, COSIC

***Abstract*—Internet of Things (IoT) devices are increasingly being deployed in critical applications, such as eHealth systems, enabled by advancements in 5G technology, which offer more than 100 Mbps of throughput, less than 5 ms of latency, and 99,999% of reliability. However, to overcome computing limitations and security measures, IoT devices rely on cloud-based solutions to outsource data processing. This dependency introduces significant security concerns, as sensitive data must be transmitted over the network and processed in external environments, increasing the risk of interception, unauthorized access, and data breaches. To mitigate these security risks, within the scope of the MOZAIK project, we deploy Network Slicing to ensure end-to-end inter-slice and intra-slice isolation across all network domains i.e., 5G Core (5GC), Transport Network (TN), and Radio Access Network (RAN). We deploy a synergy across the entire network infrastructure i.e., 5GC, TN, and RAN, to isolate the IoT data flows from the moment the data is generated until it reaches the cloud, safeguarding sensitive data during transmission. The results of our real-life experiments demonstrate that our proof of concept provides robust isolation between slices, effectively addressing the security concerns of IoT devices and enhancing the reliability and security of IoT applications. Additionally, we also include aspects of secure data storage and secure data processing, covered in the MOZAIK project.**

***Index Terms*—5G, Network Slicing, O-RAN, MPC, isolation.**

I. Introduction

The deployment of 5G networks has introduced critical enablers for the adoption of Internet of Things (IoT) devices across various applications e.g., healthcare, smart homes, and industrial automation. 5G networks address the network requirements needed by different types of applications i.e., Ultra-Reliable Low-Latency Communication (URLLC), enhanced Mobile Broadband (eMBB), and massive Machine-Type Communications (mMTC), with features such as high reliability (99,999%), low latency (5 ms), support for dense device networks (1 million of devices per cell), and high throughput (more than 100 Mbps) [1]. Hence, IoT devices with limited computational power can operate in use cases previously constrained by network limitations, as IoT devices (5G-compatible) can transfer large amounts of data to the cloud with ultra-low latency and high reliability. For instance, in eHealth applications wearable devices and connected medical systems rely on URLLC (e.g., less than 1 ms for assisted surgery [2]) to provide real-time patient monitoring and emergency response capabilities. Similarly, eMBB supports the transmission of high-resolution imaging and data streams for diagnostics. In smart home applications, IoT devices such as cameras, sensors, and connected devices depend on mMTC to handle the simultaneous connectivity of numerous devices without degradation in performance.

However, IoT devices with limited computational resources rely on the network for security, outsourcing data processing and data storage to the cloud. Additionally, ensuring data privacy and compliance with regulations, such as General Data Protection Regulation (DGPR), is particularly complex in shared networks where sensitive data traverses multiple endpoints. This dependency increases the risk of vulnerabilities during data transmission, such as interception or unauthorized access. To address these challenges, Network Slicing in 5G networks plays a crucial role in enabling secure and efficient communication. For example, in eHealth applications, slices dedicated to Vital-Sign Monitoring applications can ensure low latency and high reliability, while smart home applications must utilize separate slices to support security systems like cameras or entertainment services like video streaming. Isolated slices guarantee customized network requirements, without interfering with each other.

As part of our work on the MOZAIK project, this paper provides an overview of the vulnerabilities present in Network Slicing, the requirements for implementing secure Network Slicing, and the methods to evaluate the isolation between slices. We ensure isolation across all parts of the network infrastructure, i.e., 5G Core (5GC), Transport Network (TN) and Radio Access Network (RAN). Furthermore, we establish synergy between each network domain to achieve end-to-end isolation throughout the entire slice. To validate our approach and evaluate the MOZAIK Proof-of-Concept (PoC), we conducted real-life experiments. Additionally, for completeness, we introduce the other components of the MOZAIK architecture, such as secure data collection, secure data storage, and secure data processing.

II. Background

A. The Mozaik Project

MOZAIK is a research and innovation project focused on developing a secure, end-to-end data-sharing platform that meets the growing security and privacy requirements in the IoT domain. The goal of the project is to deploy scalable, privacy-friendly solutions at every stage of data management, from data collection to regulatory constraints.

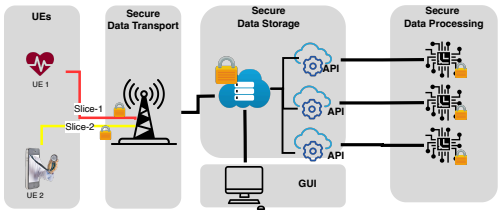

Fig. 1: MOZAIK Architecture.

To address these challenges, MOZAIK integrates advanced security mechanisms and data protection technologies, involv-

ing the entire data lifecycle, starting from data generation (e.g., sensors). The architecture of MOZAIK, shown in Figure 1, is designed to eliminate single points of failure and supports local data processing, reducing latency and limiting the exposure of sensitive information. MOZAIK deploys a platform that involves mechanisms of secure data transport, secure data collection, secure data storage, and secure data processing to increase the security of the entire chain:

*1) **Secure data transport**:* Data transport serves as the backbone of the MOZAIK ecosystem, underpinning the seamless flow of information between collection and storage nodes. Any breach during this stage compromises the integrity of the entire chain, undermining trust in the ability of the system to protect sensitive data. Within this context, Network Slicing emerges as a critical enabler of secure data transport.

*2) **Secure data collection**:* In certain cases, the data transmission between the IoT device and the cloud could be intercepted by an adversarial, e.g., when the network (or a slice) is not properly isolated. One of the most effective methods to safeguard data confidentiality and integrity is to encrypt the data at the point of origin, to guarantee high confidentiality and privacy.

*3) **Secure data storage**:* When data is stored in the cloud, precautions must be taken to ensure the data is encrypted, properly access-controlled, and protected against unauthorized tampering or exposure. Storing and processing data in the cloud exposes the data to threats such as data breaches, unauthorized access, and hacking. These challenges are further amplified when the application and its data are hosted on third-party, multi-tenant Cloud Service Providers (CSP). Data have to remain private, inaccessible to unauthorized users, and shielded from malicious tenants or even the CSP itself.

*4) **Secure data processing**:* The computation of data in cloud-based environments introduces its own set of vulnerabilities. Traditional methods of processing sensitive data on third-party clouds often expose it to potential misuse or unauthorized access. This is where emerging technologies, i.e., Multy-Party Computation (MPC) and Fully Homomorphic Encryption (FHE), come into play. MPC enables secure collaborative computation by allowing multiple parties to jointly compute a function over their inputs without revealing those inputs to each other or the processing entity. In contrast, FHE enables computations directly on encrypted data without requiring interaction among parties. For more details about MPC and FHE, we refer the reader to the project deliverable D4.2 (Full Implementation of The Proof of Concept) [3].

B. Network Slicing

Network Slicing is a key technology in 5G networks that enables the support of diverse types of applications i.e., URLLC, eMBB, and mMTC, on a shared physical network infrastructure [1]. This approach aims to fulfill the Service Level Agreements (SLAs) of different types of applications, by providing tailored network capabilities to guarantee different requirements in terms of performance, reliability, and security. To enable Network Slicing, multiple virtual and isolated networks (slices) need to be deployed within the same network infrastructure i.e., 5G Core, Transport Network, and Radio Access Network, leveraging technologies such as i) Software Defined Networking (SDN) [1], and ii) Network Function Virtualization (NFV) [1]. The architecture of the 5G Standalone (5G SA) relies on a Service-Based Architecture (SBA) and micro-services virtualization. Deployment of Network Functions (NFs) is based on containers and Virtual Machines (VMs). Such a virtualized environment enables dynamic and flexible resource allocation e.g., CPU, memory, and radio resources [4]. Hence, slices can be deployed by allocating network resources (e.g., computing, radio, and network links) coming from different network domains, to a logical network. Each slice must be deployed to operate as an independent network, ensuring that resources allocated to a slice remain unaffected by concurrent slices. For example, if an IoT device in the mMTC slice is compromised, the critical URLLC slice transmitting real-time patient data remains unaffected

III. State of the Art

Third Generation Partnership Project (3GPP) has standardized Network Slicing security from Release 15 to Release 17, with further studies ongoing in Release 18. Release 15 addressed security management, including User Equipment (UE) and NF authorization, as well as confidentiality and integrity protection of slice identifiers. Release 16 introduced Network Slice Specific Authentication and Authorization (NSSAA), which ensures that only authenticated and authorized entities can access certain slices. Release 17 focused on Application Function (AF) authorization with confidentiality protection of network slice identifiers, ensuring secure and authorized application interactions with the network. However, NSSAA and AF have been designed for the interaction between the network and applications, hence, NSSAA and AF could be deployed outside the operator domain. Moreover, the isolation of the underlying infrastructure layers remains still challenging.

Academic research has also primarily focused on specific parts of the network and does not offer a unified approach that covers the RAN, TN, and 5GC segments simultaneously. Yu et al. [5] proposed an isolation-aware slice mapping algorithm for RAN using a three-layer architecture and Wavelength Division Multiplexing (WDM) metro-aggregation networks. This approach ensures traffic isolation for URLLC, eMBB, and mMTC services, minimizing active processing nodes and wavelength channels under latency, bandwidth, and isolation constraints. Simulation-based experiments show a reduced resource usage by optimizing the placement of RAN functions, i.e., Central Unit (CU) and Distributed Unit (DU). Korrai et al. [6] exploit Orthogonal Frequency Division Multiple Access (OFDMA) to reallocate Resource Blocks (RBs) on-the-fly between eMBB and URLLC. Their optimisation model keeps eMBB users above their minimum rate even under heavy URLLC load, whereas a faster heuristic prioritises URLLC, decreasing latency by reducing queues. Simulation-based experiments show that both methods multiplex traffic without sacrificing isolation. However, these studies focus solely on the RAN domain. In our work, we consider all the network domains i.e., 5GC, TN, and RAN, to enable a full end-to-end Network Slicing configuration.

Escolar et al. [7] extend OpenvSwitch (OVS) with the Network Self-Protection (NSP) scheme within the RIGOUROUS framework, enabling isolation between slices to reduce the attack surface across multi-tenant, multi-domain networks. The results show only a negligible latency increase over the standard OVS. Cunha et al. [8] detail the EU 5Growth project, which strengthens slice orchestration and secure isolation. Combining OpenFlow with P4, Cunha et al. monitor and rate limit per-slice, adding ONOS-level P4 metering so OVS and P4-switch (bmv2) datapaths work together. Their results confirm seamless, correct traffic policing on both switch types throughout validation. For the 5GC, Qian et al. [9] address UPF resource isolation in private 5G. They combine

throughput, loss, and delay into a single isolation index and use a Particle Swarm Optimization (PSO) algorithm to size compute, and storage for the UPF at minimum cost. The results from simulation-based experiments show the relationship between resource requirements and cost, highlighting the "walkable" region where the isolation requirements are met. However, the paper does not provide specific details about the simulation environment used, such as the number of UPFs, types of servers, or traffic loads. Similarly, Esmaeily et al. [10] present a solution to improve the security and performance of isolated slices in 5G networks. The authors investigated the use of VPN solutions (such as WireGuard, IPSec, and OpenVPN) to provide isolation between the slices. Evaluation results show that WireGuard offers better isolation and performance (higher throughput for eMBB and lower latency for URLLC) than the other VPNs tested. These works address infrastructure security but fail to provide end-to-end isolation from the RAN to the 5GC.

In a nutshell, current research is fragmented, with most studies focusing on specific network domains or layers. To the best of our knowledge, none works provide a unified solution addressing the deployment of secure network infrastructure across RAN, TN, and 5GC simultaneously. Furthermore, most of these studies still rely heavily on simulations and emulations, leaving a gap in practical validation. While Network Slicing is a key technology for deploying 5G networks, there are no contributions that present results using functional prototypes at a Technology Readiness Level (TRL) of 2 or 3, where experimental proof-of-concept systems are tested.

IV. Vulnerabilities and Requirements

A. Data transport

Although standards already exist for building 5G networks, there are still no well-established guidelines on how to deploy Network Slicing to guarantee data confidentiality and integrity.

From a practical approach, a slice is built by a portion of network resources, whether physical or virtual, such as memory, vRAM, CPUs, vCPUs, storage, and radio resources. These resources come from various domains within the network infrastructure i.e., 5GC, TN, and the RAN. To create end-to-end slices, it is necessary to allocate to the slices specific network resources, depending on SLAs, coming from each of these domains. The deployment of end-to-end slices across multiple domains introduces the concept of inter-slice, where resources from different network domains are interconnected. Additionally, the granularity of resource allocation can be enhanced by creating intra-slices, which are sub-partitions of a slice. For example, in a smart hospital, applications like vital data monitoring and real-time robotic-assisted operations both demand URLLC requirement. However, their Quality of Service (QoS) differ: vital data monitoring requires latency no greater than 100 ms, while robotic-assisted operations demand latency of 20 ms [2]. In this scenario, two sub-slices would exist within the same URLLC slice, each tailored to the specific QoS needs of the respective application.

Hence, it is essential to ensure that the resources of one slice are isolated from those of other slices or sub-slices, to guarantee i) confidentiality, ii) integrity, and iii) availability.

User data, which flows from the UE to the internet through the RAN, TN, and 5GC, forms an appealing surface that must be protected. Multiple slices share the same underlying infrastructure: any congestion, cyberattack, or management issue occurring in one slice must not propagate to the others.

However, there are no ultimate specifications on how network operators should develop and enable isolation, at the infrastructure level, for Network Slicing. Deploying isolated slices requires defining partitioning mechanisms that establish distinct resource quotas from different domains among the various slices. These mechanisms depend on the domain in which they operate: for example, in the RAN, Physical Resource Blocks (PRBs) scheduling algorithms are used; in the transport network, encapsulation and bandwidth management solutions are employed; whereas for computing resources, cloud orchestration engines are utilized to properly create and assign VMs and containers. To help guide a secure deployment of Network Slicing, the National Security Agency (NationalSA) and Cybersecurity and Infrastructure Security Agency (CISA) have issued detailed recommendations regarding isolation and segregation across slices. These recommendations are elaborated in Table I. Consequently, isolation can be assessed along different dimensions: i) performance, to maintain the Key Performance Indicators (KPIs) of a slice regardless of the conditions of other slices; ii) management, to enable the operator to manage each slice as if it were a separate network; and iii) security and privacy, to prevent threats or attacks on one slice from affecting the data and information of the others.

B. Data collection

A secure data collection solution must provide high performance while minimizing overhead in key processes such as reading sensor data, performing necessary data conversions, encrypting the data, and transmitting it. At the same time, it must ensure the integrity and availability of the data. To mitigate potential data misuse in the event of unauthorized data access by a bad actor, it is essential to involve a secure data collection system that safeguards the data immediately at its source, directly following its generation or measurement.

C. Data storage

When storing data in the cloud, especially with multi-tenant third-party providers, it is vital to protect against unauthorized access and data tampering. A robust access control mechanism prevents data theft or modification at the application level, while securing the underlying infrastructure ensures data remains safe even if the system is compromised. Because breaches can still occur, compromised data should not reveal any sensitive information or insights about data subjects. Lastly, to handle potentially life-critical health data, the storage solution must support high-frequency data ingestion from numerous devices.

D. Data processing

Common practices for processing sensitive data in the cloud generally rely on complete trust in the cloud provider. This creates two main vulnerabilities: i) Data misuse or leakage, where sensitive input or intermediate data, especially personal health information, may be leaked, collected, or sold, and ii) Lack of correctness guarantees, where a single cloud server offers no assurance of accurate computations. A malicious provider could alter inputs or outputs. In a smart hospital setting, this may lead to misleading or incorrect diagnoses with serious real-world consequences.

To develop a truly end-to-end secure data-sharing platform, private data must remain confidential throughout processing so that only the data owner can access the inputs or diagnostic results.

TABLE I: Recommendations for secure Network Slicing.

Recommendation	Description
Logical Isolation	Each network slice should have logically separated resources (e.g., virtualized network functions), preventing unauthorized cross-slice interaction.
Performance Isolation	Guarantee distinct performance metrics (e.g., bandwidth, latency) for each slice so that spikes or resource contention in one slice do not degrade others.
Physical Resource Isolation	For high confidentiality, integrity, and availability (CIA) requirements, allocate dedicated physical resources (e.g., servers, hardware accelerators) to a single slice.
Separate Management Systems & Admins	Use unique management tools and assign distinct administrative privileges per slice, minimizing the impact of configuration changes on other slices.
Data Plane Segregation	Data plane activities in one slice must not influence or be influenced by another slice's data plane.
Control Plane Independence	Control plane actions (creation, update, deletion) in one slice should not affect other slices, preventing service interruptions or security leaks.

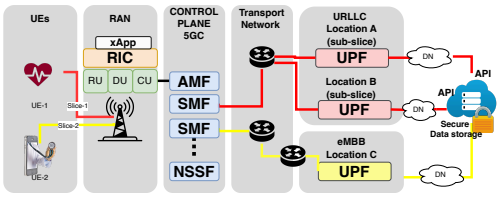

Fig. 2: PoC Architecture.

E. Challenges with GDPR compliance

Since the deployment of end-to-end scalable and secure data-sharing platforms can involve the processing of personal data, such platforms must adhere to the strict requirements of the GDPR. While end-to-end encryption is a robust privacy mechanism, it is classified as pseudonymization under GDPR, meaning that encrypted data is still considered personal data. Furthermore, in the context of secure data processing, there is ambiguity regarding whether the data being processed with technologies such as MPC or FHE is considered personal. Although MPC and FHE prevent any single entity from accessing the entire dataset, encrypted key shares could potentially compromise privacy if mismanaged. To mitigate these risks, entities identified as data controllers or processors must implement the necessary technical and organizational safeguards.

V. MOZAIK: Real-life Proof of Concept

In this section, we focus on the deployment of secure data transport by enabling seamless Network Slicing, creating a synergy across all the network domains i.e., 5GC, TN, RAN. To validate the feasibility of addressing the vulnerabilities and requirements described in Section IV within the MOZAIK architecture, we deployed a 5G SA network as a PoC (Figure 2). For completeness, in this section, we also provide an overview of the solutions implemented for the other components of the MOZAIK project, including data collection, data storage, data processing, and regulatory compliance.

A. Data transport

Within the MOZAIK architecture, shown in Figure 1, we enable Network Slicing guaranteeing isolation by i) deploying a decentralized 5G SA architecture, and ii) creating a synergy between all the network domains i.e., 5GC, TN, and RAN. The combination of NFV and SDN used in 5G networks realize a SBA where data plane and control plane are separated [1]. The SBA used by the Core Network (CN) establishes a cloud-based model where each NF can be placed in a different location and/or in a different type of machine e.g., bare metal, VM, container, and pod.

*1) **Core Network**:* We deployed the primary 5GC Network Functions (Access and Mobility Management Function (AMF), Session Management Function (SMF), Network Slice Selection Function (NSSF), and User Plane Function (UPF)) in separate VMs to i) ensure isolation among the NFs, and ii) to enable isolation between the resources of the VMs. Each VM has assigned its own pool of resources e.g., vCPUs, VRAM, storage, and network interfaces, so that each NF operates independently. We separate the control plane from the data plane for each slice, logically and physically: as shown in 2, each slice type has its own SMF (which belongs to the control plane) for the inter-slice isolation, while each sub-slice has a dedicated UPF (for data-plane operations). These UPFs are placed in different data centers to achieve intra-slice isolation through physical resource isolation and data plane segregation, in line with the recommendations from the Non-Standalone (NSA) and CISA. By adopting such a decentralized architecture for the 5GC, in the MOZAIK project we guarantee inter-slice and intra-slice isolation.

From a technical point of view, in the 5GC, each slice is identified by a pair of values: the Slice/Service Type (SST) and the Session Description (SD). The SST defines the slice type, such as eMBB, URLLC, or mMTC, while the SD provides an additional classification that can represent sub-slices or other specialized service characteristics. Furthermore, the SST and SD values are tagged with the Data Network Name (DNN), which UEs uses to request access to a specific slice or sub-slice.

*2) **Transport Network**:* In order to achieve robust isolation both across different slices (inter-slice) and within the same slice (intra-slice), the connections between SMF and UPF is assigned to a different TN path and interface. In that way the control plane of each slice is kept physically separated. By doing so, the control commands for one slice or sub-slice are insulated from any impact by other slices, preventing attacks and vulnerabilities in one slice from propagating to another. Similarly, for data-plane traffic, each UPF corresponding to a specific slice or sub-slice is connected to the DU of the RAN via a dedicated TN interface and physical connection. Although assigning each slice or sub-slice to a dedicated transport path and physical connection between the SMF and UPF (and likewise between UPF and DU) ensures proper isolation, it inevitably introduces scalability and cost challenges. For those slices or sub-slices with stringent reliability and security requirements, full physical separation remains a viable strategy. For less critical slices, logical separation through Virtual Local Area Network (VLAN), Virtual Routing and Forwardings (VRFs), or Virtual Private Network (VPN) tunnels may suffice, achieving isolation at a lower cost.

*3) **Radio Access Network**:* At the RAN domain, we deploy Network Slicing by leveraging the Open Radio Access Network (O-RAN) paradigm and deploying an xApp. The xApp dynamically allocates a specific amount of radio resources

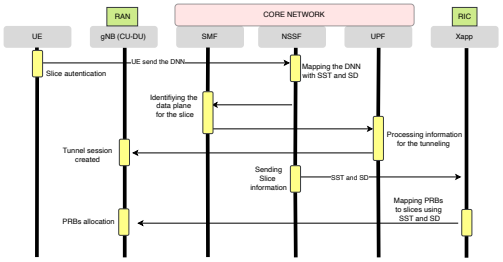

Fig. 3: Diagram chart of the messages exchange to enable synergy.

i.e., RBs, to each slice, identified by the pair of SST and SD values. We first use the SST value to define a maximum resource quota for each slice type (e.g., eMBB, URLLC, or mMTC) allocating a certain amount of radio resources to a slice in order to prevent inter-slice interferences. At the same time, this quota forms a shared "pool" from which the sub-slices within that slice type can draw. Next, to enable intra-slice isolation within the RAN domain, we assign a maximum allocation of PRB for each sub-slice using the SD values, to ensure that each sub-slice receives a dedicated share of the resource pool. This hierarchical allocation (slice-level first, then sub-slice-level) enables flexibility in distributing resources and prevents one sub-slice from monopolizing the entire pool of radio resources available.

4) ***Synergy****:* Slice selection and management begin the 5GC, where the NSSF assigns each UE to the appropriate slice by matching the SST and SD values, as show in Figure 3. Once the slice is determined, the SMF establishes dedicated data-plane tunnels between the CU and the UPF, where strong isolation using physical separation at the TN level is used. At the RAN level, the xApp applies the same SST-SD pair to allocate radio resources for each slice or sub-slice. This ensures that the slicing policies defined at the 5GC are fully reflected in the RAN domain, enforcing slice-aware scheduling and preventing cross-slice interference.

Through this slice identification in the 5GC, in the transport layer, and RAN, we deploy a full end-to-end Network Slicing mechanism with inter-slice and intra-slice isolation across every segment of the network.

B. Data collection

To ensure data confidentiality and integrity, the collected data from the IoT devices is encrypted using a state-of-the-art lightweight encryption scheme as soon as possible after generation at the source. Every sample generated by the IoT device is encrypted using AES-GCM with a key size of 128-bit as specified in NIST SP 800-38Dor RFC 5288. More details about the deployment of Data collection phase are described in the project deliverable D4.2 (Full Implementation of The Proof of Concept) [3].

C. Data storage

The data storage unit utilized within MOZAIK is Obelisk [3], specifically Obelisk High Frequency Streaming (HFS). Obelisk HFS is a cloud-based time series storage platform, implemented by means of an event-based and asynchronous microservices-based architecture, that focuses on high-frequency IoT data. A MOZAIK specific software layer with a focus on privacy-preserving properties is built around Obelisk-HFS, which we call MOZAIK-Obelisk. This includes an appropriate encryption key management system where encrypted key shares are stored, an overarching API for unified, mediated and secure access to the data, and state-of-the-art isolation techniques. For more technical details, we invite the reader to consult the MOZAIK-Obelisk documentation [3].

D. Data processing

Within the MOZAIK project, we offer solutions for private data analytics based on both MPC and FHE.

For MPC-based processing, the workflow involves three steps: distributed decryption of input data, privacy-preserving analytics, and distributed encryption of the results. Data remains encrypted from its collection at IoT devices to its decryption on the doctor's local device, ensuring end-to-end confidentiality. Secure key distribution is critical to prevent processing parties from reconstructing the full decryption key, safeguarding data privacy.

In FHE-based processing, the workflow is simpler. Data encrypted at the source with the CKKS FHE scheme [3] remains encrypted throughout analysis, ensuring confidentiality. The analysis, performed on FHE-encrypted data, uses the OPENFHE [3] library to evaluate a neural network.

Both approaches ensure secure, efficient, and privacy-preserving data analytics, tailored to meet the challenges of IoT environments and real-time processing requirements.

E. GDPR compliance in MOZAIK

We conducted a Data Protection Impact Assessment (DPIA), required under GDPR Article 35, to map how data flows through the system, identify risks, and highlight sensitive areas like patient data in smart hospitals. Furthermore, to assess and address privacy risks, within the MOZAIK project, we use the LINDDUN framework [3], a tool for identifying and solving privacy issues in its architecture.

VI. Validation and Results

In this section, we assess our PoC to determine whether our Network Slicing deployment ensures i) integrity, ii) availability, and iii) confidentiality (as defined in Section IV), both inter-slice and intra-slice. Our evaluation covers: i) network performance (throughput, packet loss), ii) isolation, and iii) security and privacy.

We recreated a smart hospital scenario on a real 5G testbed [11], considering three slices: i) a eMBB slice (70 Mbps) for Clinical Data Access, enabling staff to use tablets and stream/download large files; ii) a URLLC sub-slice for Robotic-Assisted Procedures requiring 40 Mbps ; and iii) another URLLC sub-slice for Vital-Sign Sensors (e.g., ECGs) demanding 20 Mbps.

Under normal network conditions, the network has enough resources to handle conventional traffic loads, without showing how the network behaves when multiple applications compete for limited resources coming from the same pool (same network infrastructure). Hence, to stress the network, we used three UEs sending 80 Mbps User Datagram Protocol (UDP) traffic each (via Iperf[1]): Clinical Data (eMBB), Robotic-Assisted Procedures (URLLC), and Vital-Sign Sensors data (URLLC). Unlike TCP, with UDP is possible to overload the network to expose the network limits.

To establish a baseline, we first disabled all slicing configurations, then we introduced sequentially the different traffic flows to observe the network's behaviour. As shown

[1]Iperf:https://iperf.fr

in Figure 4, the Clinical Data service starts at second 0 and runs until second 20, maintaining 80 Mbps with 0% packet loss. At second 20, we add the Robotic-Assisted traffic flow, increasing the demand of network resources needed to satisfy the QoS for both services. As a result, at second 21 the throughput for both services falls below 80 Mbps, while packet loss rises to 20%. At second 40, the addition of Vital-Sign Data further degrades performance, with 40% packet loss across all services—indicating the network's inability to maintain QoS under load. At second 60, we dynamically enable our Network Slicing configuration to allocate quotas of resources (dynamically configurable) from the different network domains i.e., the 5GC, TN, and RAN. In fact, at second 60 the network behavior shown in the upper side of Figure 4 changes completely. The eMBB slice guarantees 70 Mbps to the Vital-Sign Data service, while the URLLC slice guarantees 40 Mbps to the Robotic-Assisted Procedures and 16 Mbps to the Vital-Sign Data Service. However, the Vital-Sign Data service continues to request 80 Mbps, but since its URLLC subslice has not enough network resources, the Vital-Sign Data service is able to achieve only 16 Mbps, leading to 80% of packet loss. Similarly, the Robotic-Assisted Procedures flow is not able to achieve more than 40 Mbps, having 50% of its packets loss. The eMBB slice experiences only 10% packet loss, maintaining higher throughput due to adequate resource allocation. These results confirm that our work effectively prevents one slice from using network resources allocated to other slices, ensuring i) integrity and ii) availability intra-slice (URLLC-URLLC), and inter-slice (URLLC-eMBB).

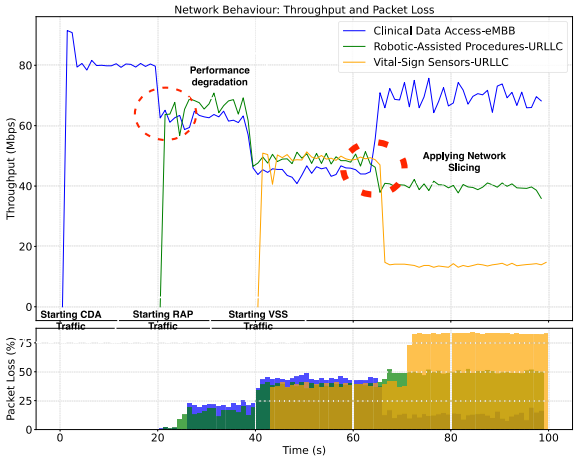

Fig. 4: Real-life experiment results.

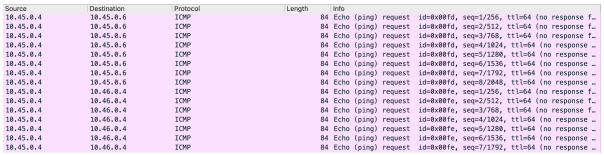

Fig. 5: Packets tracer.

For what concerns the confidentiality intra-slice and inter-slice, we performed a different experiment. Since all the UEs are connected to the same network infrastructure e.g., same RAN, the UEs should be able to ping each other. Hence we first ping the UEs belonging to the same slice i.e., URLLC, but different sub-slices. Than we ping the UEs belonging to different slices i.e., URLLC and eMBB. In Figure 5 we show the results of our experiment, using a packet tracer. The first packets are between two UEs within the same slice i.e., URLLC, as the IP of the soruce and the destination describes (both 10.45.X.X). The *Info* column idicates that there are obly ping requested sent, without responses. The same applies to the second part of Figure 5, where the ping is sebd between two UEs belonging to different slices URLLC and eMBB. Also in that case, the tracer shows only ping requests, with no responses from the destination.

VII. Conclusion

Within the scope of the MOZAIK project, this work presents a network configuration that establishes a synergy across all network domains—5G Core, Transport Network, and Radio Access Network—to enable end-to-end Network Slicing with both inter-slice and intra-slice isolation. This synergy is achieved by aligning slice selection, control plane, and data plane mechanisms, ensuring that network resources are dynamically and securely partitioned according to application-specific requirements. To validate our approach, we implemented a decentralized 5G Standalone architecture as a real-world Proof-of-Concept (PoC). This PoC integrates physical and logical isolation mechanisms across the entire infrastructure, thereby ensuring the confidentiality, integrity, and availability of sensitive data throughout its collection, transport, and processing lifecycle.

Acknowledgement

This work has been performed in the framework of the Flemish Government through FWO SBO project MOZAIK S003321N, and the European project SNS JU TrialsNet (Grant Agreement No. 101095871).

References

[1] 3GPP, "3GPP TS 28.530 V16.6.0 (2023-03) Technical Specification: Management and orchestration; Concepts, use cases and requirements (Release 16)," tech. rep.

[2] (3GPP), "Technical specification group services and system aspects; study on communication services for critical medical applications (release 17)," 2021.

[3] MOZAIK Project Team, "D4.2: Full PoC final report," 2024. https://www.esat.kuleuven.be/cosic/projects/mozaik/wp-content/uploads/sites/2/2024/06/D4.2_full_PoC_final.pdf.

[4] G. F. Pittalà, L. Rinieri, A. Al Sadi, G. Davoli, A. Melis, M. Prandini, and W. Cerroni, "Leveraging data plane programmability to enhance service orchestration at the edge: A focus on industrial security," *Computer Networks*, 2024.

[5] Y. H. et al., "Isolation-aware 5G RAN slice mapping over wdm metro-aggregation networks," *Journal of Lightwave Technology*, vol. 38, no. 6, pp. 1125–1135, 2020. url: https://ieeexplore.ieee.org/document/9365081.

[6] P. K. et al., "A RAN Resource Slicing Mechanism for Multiplexing of eMBB and URLLC Services in OFDMA Based 5G Wireless Networks," *IEEE Access*, vol. 8, pp. 45674–45686, 2020. url: https://ieeexplore.ieee.org/document/9027575.

[7] A. M. Escolar, J. B. Bernabe, J. M. A. Calero, Q. Wang, and A. Skarmeta, "Network slicing as 6g security mechanism to mitigate cyber-attacks: the rigourous approach," in *Proceedings of the 2024 IEEE 10th International Conference on Network Softwarization (NetSoft)*, pp. 387–392, IEEE, 2024.

[8] C. et al., "5 growth: Secure and reliable network slicing for verticals," in *2021 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*, pp. 347–352, 2021. doi: 10.1109/EuCNC/6GSummit51104.2021.9482536.

[9] S. Qian, F. Chen, H. Ning, T. Lin, and Y. Wang, "Computing and Storage Resources Allocation of UPF Based on Isolation in Private 5G Networks," *Vehicular Technology Conference (VTC)*.

[10] A. Esmaeily and K. Kralevska, "Orchestrating Isolated Network Slices in 5G Networks," *Electronics*, vol. 13, no. 1548, 2024.

[11] X. Limani, V. Charpentier, A. Troch, M. Camelo, J. M. Marquez-Barja, and N. Slamnik-Kriještorac, "Network slicing as the ultimate enabler of enhanced service quality in vehicular-to-everything (v2x) world," in *2024 IEEE 100th Vehicular Technology Conference (VTC2024-Fall)*, pp. 1–5, 2024.