Secure Hardware Accelerated Data Analytics for 6G Networks: the PRIVATEER approach

Ilias Papalamprou*, Aimilios Leftheriotis†, Apostolos Garos‡, Georgios Gardikis‡,
Maria Christopoulou§, George Xilouris§, Lampros Argyriou¶, Antonia Karamatskou¶, Nikolaos Papadakis¶,
Emmanouil Kalotychos‖, Nikolaos Chatzivasileiadis‖, Dimosthenis Masouros*, Dimitrios Soudris*
*National Technical University of Athens, Greece †University of Patras, Greece
‡R&D Department, Space Hellas S.A., Greece §NCSR "Demokritos", Institute of Informatics and Telecommunications, Greece
¶Infili Technologies S.A., Greece ‖Ubitech Ltd., Digital Security & Trusted Computing Group, Greece
*{ipapalambrou, dmasouros, dsoudris}@microlab.ntua.gr, †aleftheriotis@ac.upatras.gr
‡{agaros, ggar}@space.gr, §{maria.christopoulou, xilouris}@iit.demokritos.gr,
¶{largyriou, akaramatskou, npapadakis}@infili.com, ‖{mkalotychos, nchatzivasileiadis}@ubitech.eu

***Abstract*—Next generation 6G networks are designed to meet the requirements of modern applications, including the need for higher bandwidth and ultra-low latency services. While these networks show significant potential to fulfill these evolving connectivity needs, they also bring new challenges, particularly in the area of security. Meanwhile, ensuring the privacy is paramount in 6G network development, demanding robust solutions following "privacy-by-design" principles. To address these challenges, PRIVATEER project strengthens existing security mechanisms, introducing privacy-centric enablers tailored for 6G networks. This work, evaluates key enablers within PRIVATEER, focusing on the development and acceleration of AI-driven anomaly detection models, as well as attestation mechanisms for both hardware accelerators and containerized applications.**

***Index Terms*—6G, Security, Privacy, AI, Hardware Acceleration**

I. Introduction

The development of beyond-5G (B5G) and upcoming 6G networks marks a pivotal shift in mobile communications, promising to significantly reduce latency, enhance bandwidth, and improve overall performance [1]. The planned evolution towards 6G aims to enable transformative technologies, e.g. augmented reality, massive Internet of Things (IoT), and AI-driven applications, by offering faster, more reliable and ultra-responsive network experiences [2]. The shift to 6G not only plans to enhance connectivity and speed but also needs to address emerging security and privacy challenges arising from immersive use cases.

PRIVATEER project[1] addresses the evolving security and privacy challenges of 5G and envisioned 6G networks through four key areas [3]: i) Decentralized Security Analytics, ii) Infrastructure and Service Attestation to ensure integrity across a multi-actor environment (including Service Providers, Mobile Network Operators, Infrastructure Providers, and End Users), iii) Privacy-Aware Orchestration, and iv) Privacy-Friendly Cyber Threat Intelligence (CTI) Sharing. Supporting these technologies, a Distributed Ledger (DL) provides a transparent, accountable framework to protect sensitive data and store immutable records, such as trustworthiness evidence.

Given the complex nature of 6G networks, the adoption of AI is one of the most promising solutions to support advanced network management as well as sophisticated security mechanisms [4]. However, due to the high computational demands of these applications, general-purpose processors (CPUs) alone are insufficient for deployment. This challenge underscores the need for a heterogeneous computing continuum - integrating various specialized hardware accelerators to meet the low-latency and energy-efficiency requirements of 6G services. Among these options, Field Programmable Gate Arrays (FPGAs) distinguish as a flexible option, providing low-latency, energy-efficient performance well-suited to 6G's needs [5]. However, the intricate structure of 6G networks increases the risk of complex security threats, that existing security mechanisms may be inadequate to handle [2]. For instance, unique threats such as malware updates, trojan detection, Distributed-Denial-of-Service (DDoS) attacks etc. require tailored approaches for the heterogeneous hardware devices to safeguard the network's integrity and protect sensitive data.

This paper focuses on specific aspects of PRIVATEER that have been developed in the first half of the project and are crucial for the secure operation of 6G services, particularly in the domain of data analytic functions and their secure deployment at the edge. More specifically, the contributions of PRIVATEER project presented in this paper are the following:

- **Anomaly detection AI models** for detecting Distributed Denial-of-Service (DDoS) attacks in the network.
- **Hardware Acceleration of AI models** aiming for both low latency and energy efficient execution of data-analytics applications in the edge.
- **Secure configuration of hardware accelerators** for ensuring their secure deployment in the edge.
- **Attestation for virtualized infrastructures** for establishing trustworthiness for the associated containerized applications.

This work is funded by the Smart Networks and Services Joint Undertaking (SNS JU) under the EU Horizon Europe programme PRIVATEER under Grant Agreement No. 101096110.

[1]https://www.privateer-project.eu

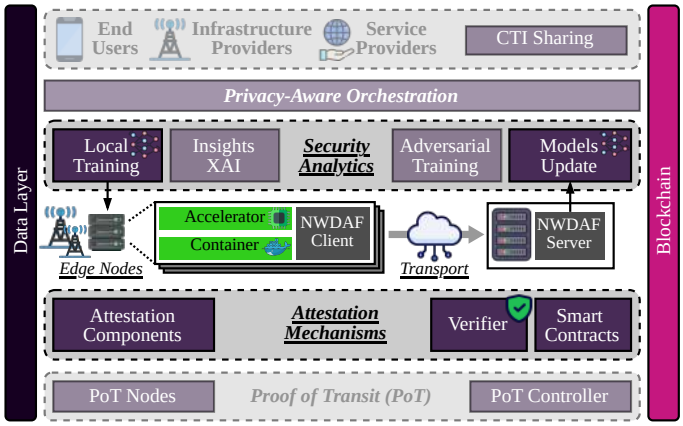

Fig. 1: PRIVATEER's high-level architecture.

The rest of the paper is organized as follows: Section II presents an overview of PRIVATEER's architecture. Section III dives into the details of specific security enablers that have been developed. Section IV showcases the evaluation results. Finally, Section V summarizes the progress up-to-date, as well as mentioning some future directions for the project.

II. PRIVATEER's Vision & Overview

PRIVATEER Horizon 2020 research and innovation programme, initiated on February 2023, with an expected duration of 3 years. The abstract architecture of PRIVATEER's framework, as analyzed in [3] and illustrated in Fig. 1, aims to develop innovative security enablers *following a privacy-by-design approach*, in alignment with EU standards, including GDPR. PRIVATEER introduces innovative methodologies to ensure privacy preservation across all stakeholders, encompassing End Users, Infrastructure Providers/Neutral Hosts, and Service Providers. The proposed security enablers include the following:

i) Decentralized Robust Security Analytics: The existing security analytics are enhanced using AI, for detection and classification of network threats from traffic data and logs. Furthermore, to improve privacy, it adopts decentralized federated learning, supported by adversarial training and anonymization techniques [6], that ensure trustworthiness in terms of privacy, fairness, robustness, and explainability.

ii) Privacy-aware Slicing and Orchestration: PRIVATEER integrates privacy awareness and user intent, while securing the service path. The framework prioritizes creating trusted network topologies by incorporating Proof-of-Transit (PoT) services and implementing privacy-focused orchestration mechanisms based on individual user needs.

iii) Distributed Attestation: Various mechanisms for distributed privacy-preserving attestation, identification and threat sharing are introduced. This includes tools for distributed verification through digital trusted wallets and verifiable credentials, as well as components necessary for remote attestation of the 6G services and the heterogeneous hardware infrastructure.

iv) Cyber Threat Intelligence (CTI) Sharing: CTI sharing enables privacy-preserving secure exchange of threat intelligence.

In the final stage of the project, PRIVATEER's enablers will be evaluated across five use cases within two distinct environments: Intelligent Transport Systems (ITS) and Smart Cities, demonstrating the framework's effectiveness in representative, real-world scenarios.

This paper provides a deep-dive analysis of PRIVATEER's hardware-accelerated data analytics engine and attestation mechanisms, which enable the deployment of secure monitoring across heterogeneous distributed edge computing environments. The security enablers developed are open-source and available through PRIVATEER's GitHub repository[2].

III. PRIVATEER's Security Enablers

A. AI-Driven Hardware Accelerated Security Analytics

With the rapid increase in connected devices within 6G networks, the ability to efficiently detect attacks is essential for maintaining reliable and secure network performance. Detecting anomalies in network traffic is particularly important for identifying threats like DDoS attacks, which attempt to overwhelm a targeted server, service, or network by flooding it with malicious internet traffic. Particularly in the 6G ecosystem, DDoS attacks pose significant security risks for various critical technologies, such as Network Function Virtualization (NFV) and Software-Defined Wide-Area Networks (SD-WAN) [7]. To address these risks, PRIVATEER has developed AI-based security analytics, leveraging data from the Network Data Analytics Function (NWDAF), and incorporate a Federated Learning (FL) training scheme to enhance privacy preservation. NWDAFs are distributed across the network continium and are responsible for continuously monitoring network traffic, providing insights into network data production and consumption. On top of that, to support real-time threat detection, often required within a few milliseconds [8], PRIVATEER leverages FPGA hardware accelerators to provide efficient and high-performing threat detection across the network.

1) LSTM Autoencoder for Anomaly Detection: PRIVATEER leverages a Long Short-Term Memory (LSTM) Autoencoder Neural Network architecture to identify anomalies on streaming data coming from the NWDAF. LSTM networks, a popular variant of Recurrent Neural Networks (RNNs), excel in managing long-term dependencies, making them particularly effective for time-series analysis in network monitoring [9].

The LSTM Autoencoder (LSTM-AE) combines the strengths of Autoencoders, which learn efficient representations of input data, with the temporal modeling abilities of LSTMs. This hybrid architecture is well-suited for detecting anomalies in sequential data [10], as it learns typical patterns and identifies deviations based on reconstruction error. In this setup, the encoder transforms input sequences into a latent representation, while the decoder attempts to reconstruct the sequence from this representation. For anomalous data, the reconstruction error is usually significantly higher, allowing the model to effectively signal irregular patterns.

[2]https://github.com/privateer-project

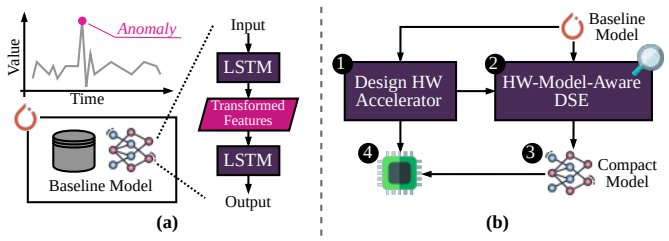

Fig. 2: **(a)** LSTM-AE development. **(b)** FPGA acceleration.

To fine-tune the LSTM autoencoder for optimal performance, PRIVATEER employs a systematic hyperparameter-tuning approach using Grid Search. This method involves defining a comprehensive search space and systematically evaluating the model's performance across all possible combinations of hyperparameter values. The process ensures that we identify the set of hyperparameters that yields the best results according to predefined metrics. The search space for our model includes: the window size, the model's architecture, layer normalization, dropout rate, loss function and learning rate. This Grid Search identified a regularized overcomplete LSTM-AE as the optimal architecture, which we refer as *Baseline* model.

2) Acceleration of Data Analytics on FPGAs: PRIVATEER leverages FPGAs to facilitate the deployment of energy-efficient, low-latency AI-based data analytics applications at the edge. FPGAs are well-suited for such tasks due to their parallel-processing capabilities, which allow for significant acceleration of computationally intensive operations.

The process of hardware acceleration of PRIVATEER's AI models, illustrated in Fig. 1b, begins with the design and optimization of the hardware accelerator itself. In Step ❶, PRIVATEER's *Baseline* model serves as the foundation, where design choices for key model components (e.g., LSTM layers) are translated into architectural specifications suited for FPGA deployment. The FPGA design includes specific considerations to accommodate the structure and requirements of the LSTM model, ensuring that the hardware design can support the model's inference tasks efficiently.

In Step ❷, a Hardware-Model-Aware Design Space Exploration (DSE) is conducted. This stage integrates both hardware design and model characteristics, guiding the DSE process to identify an optimal configuration that balances the performance requirements of the AI model with the capabilities and limitations of the target FPGA platform. By combining hardware-awareness and model-awareness, this DSE process not only maximizes model performance for its intended anomaly detection task but also adjusts the model's structure to fully capitalize on the specialized hardware design developed in Step ❶, achieving a highly efficient and compact design.

The DSE process in Step ❷ results in a refined version of the *Baseline* model—referred to as the *Compact* model ❸. The DSE ensures that the *Compact* model achieves the best possible performance trade-offs, balancing minimal accuracy degradation with significant gains in energy efficiency and processing speed.

Finally, Step ❹ involves the implementation of the *Compact* model on the FPGA platform. It involves the integration of the identified *Compact* model ❸ with the hardware design developed in Step ❶, to realise a complete FPGA-based inference accelerator.

B. Secure Services in Heterogeneous Edge Systems

As B5G technologies progress, new functionalities introduce fresh challenges, highlighting the need for adaptive security measures to address evolving threats. In 6G, robust security is critical not only for services but also for the infrastructure that supports them, including virtualized services, Virtual Network Functions (VNFs), and a heterogeneous edge infrastructure composed of CPU-based systems and diverse hardware accelerators. These complex systems are susceptible to various attacks, such as reverse engineering attempts, malicious code injection, hardware trojans, as well as physical attacks in multi-tenant scenarios [11]. To address these risks, PRIVATEER employs remote attestation protocols in conjunction with dedicated hardware (e.g. encryption modules, Trusted Execution Environments (TEEs)), to enable secure service deployment across the edge, covering both hardware-accelerated applications on FPGAs and containerized applications. This enables a system to prove its trustworthiness to a remote verifier by responding to a challenge with specific security evidence, which the verifier then validates against some reference values (Fig. 3a). Hence, following Zero Trust principles -where no entity is inherently trusted and all must continuously demonstrate their trustworthiness- we ensure secure service deployment across the edge, covering both hardware-accelerated applications on FPGAs and containerized applications. Additionally, for collecting any security evidence for service deployment, a decentralized secure storage mechanism is employed, i.e., the Blockchain.

1) Secure configuration of FPGAs: For edge nodes featuring FPGA hardware accelerators, a custom hardware/software solution is presented to ensure secure configuration. As shown in Fig. 3b, four key entities are involved:

Application Provider: Refers to the individual that provides the FPGA accelerated application, namely the bitstream.

Edge Node: The computing platform with the FPGA accelerator. Apart from hosting the hardware accelerated application, the edge node contains supplementary hardware (*encryption kernel*) and software (*attestation service*) modules responsible for verifying the integrity and authenticity of all individual components. The *attestation service* serves as the backbone of the system, handling the collection of any security evidence using hardware and software components, as well as exchanging those data with an external attestation server.

Attestation Server: Acts as an external authority whose objective is to verify the integrity of all individual components and deploy them securely in the target edge node. It coordinates with all involved parties, namely the application provider, the edge node, and PRIVATEER's Blockchain.

Blockchain: Serves as a decentralized storage system for any security evidence collected during the attestation process (details in Section III-B3).

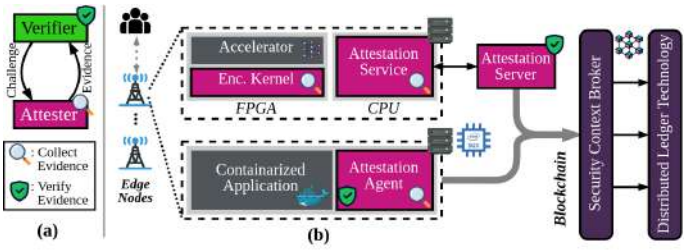

Fig. 3: Security methodologies for edge nodes: **(a)** Overview of remote attestation principle. **(b)** Secure configuraiton mechanisms for heterogeneous edge nodes

When a hardware accelerated application needs to be deployed to the edge, the procedure is the following: Initially an offline phase is required, in which the application provider encrypts the application's bitstream and transfers it to the edge node. Meanwhile, the attestation server collects the reference values used for verification (e.g., cryptographic checksum, digital signatures), of both the attestation service installed in the edge node and the application's bitstream. After this preparation, the online phase initiates. The verification process is twofold:

- Firstly, the integrity of the *attestation service* preinstalled in the edge node is ensured.
- Subsequently, the remote attestation protocol proceeds by verifying the authenticity of the application bitstream.

Only after both steps are successfully completed, the application's bitstream is decrypted and loaded onto the FPGA. Furthermore, after each attestation process is finalized, the results are forwarded to the Blockchain, to be stored in a decentralized location.

2) Secure Containerized Applications: For addressing the vulnerabilities of the containerized applications, TEEs are leveraged to ensure secure and reliable operation of such applications. TEEs are secure areas within a device's processor that provide isolated execution of code and data, shielding sensitive information from the rest of the system. To achieve robust security for sensitive applications, PRIVATEER leverages Intel Software Guard Extensions (Intel SGX) technology [12], as its TEE solution. Intel SGX is a set of instructions designed to enhance the security of application code and data by creating "SGX Enclaves", providing protection from attacks launched by untrusted applications.

Additionally, beyond providing advanced protection through Intel SGX, PRIVATEER leverages the TEE-enabled infrastructure for extracting (verifiable) evidence; thus performing attestation tasks, ensuring the integrity and authenticity of the code running within an enclave. More specifically, a core component is proposed, i.e., the *Attestation Agent*, that is responsible for monitoring and verifying:

- The static properties of the containerized application's integrity, focusing on secure launching of a container.
- The runtime properties of the underlying (virtualised) infrastructure configuration integrity.

The output of the *Attestation Agent* is a proof of correctness, which does not reveal the exact evidence, providing privacy-preservation. Furthermore, the extracted attestation evidence, similar with those collected from the edge nodes with FPGAs, is stored in PRIVATEER's Blockchain, guaranteeing its integrity and accessibility.

3) Blockchain: Serves as the decentralized storage for preserving security evidence acquired from edge nodes, and is comprised of two components: The *Security Context Broker (SCB)*, a custom service that acquires the attestation appraisals from the heterogeneous edge nodes and the *Distributed Ledger Technology (DLT)*, where those evidence are stored. Regarding the *DLT*, Hyperledger Besu [13] is adopted, due to its compatibility with both private and public Blockchain networks. Strong privacy features and a high transaction volume processing capacity are two things that Besu offers, which is particularly appealing to PRIVATEER considering that trustworthy evidence may include a lot of data. Furthermore, PRIVATEER's interoperability with Ethereum is quite helpful since it makes it possible to leverage the variety of resources and apps that the Ethereum ecosystem has to offer.

Alongside auditability, encryption, integrity protection, authentication and access control mechanisms are employed to restrict access to only authorized entities and to ensure accountability for all actions taken. *Smart Contracts (SC)*, which are self-executing agreements with predefined rules encoded on the Blockchain, provide advanced access control capabilities. These contracts enforce access policies, specifying which parties are authorized to access particular data or execute specific actions on the Blockchain, thereby enhancing security and accountability. They offer distinctive methods to facilitate trustworthiness evidence, such as attestation, CTI, and proof of transit, while maintaining data integrity and confidentiality across networks.

IV. Evaluation of Security Enablers

A. AI-Driven Hardware Accelerated Security Analytics

1) Development of AI models: To evaluate the accuracy of PRIVATEER's AI models, we utilize the NCSRD-DS-5GDDoS dataset [14]. This dataset is composed of multivariate time series-data, which includes multiple variables recorded in regular time intervals. Each series in this dataset consists of various metrics related to 5G radio and core networks, capturing the performance and behaviour of network components during normal operation and sporadic DDoS attacks. This format enables us to observe the dynamics of network traffic and security incidents over time, making it suitable for temporal analysis and anomaly detection.

Centralized Model Training: We first evaluate PRIVATEER's autoencoder architecture using a centralized training approach. While PRIVATEER leverages a FL approach for training, this analysis allows to determine the maximum achievable model accuracy. As mentioned in Section III-A2, to determine the optimal LSTM-AE configuration for anomaly detection, we performed a DSE focusing on different LSTM-AE parameters, using Grid Search. The DSE identified a regularized overcomplete autoencoder, containing two LSTM layers in the

TABLE I: Accuracy Metrics for different AI models

Model	Accuracy	Precision	Recall	F1 Score
Baseline	0.9982	1.0000	0.9822	0.9910
Baseline-FL	0.9939	0.9640	0.9763	0.9701
Compact	0.9936	1.0000	0.9416	0.9699

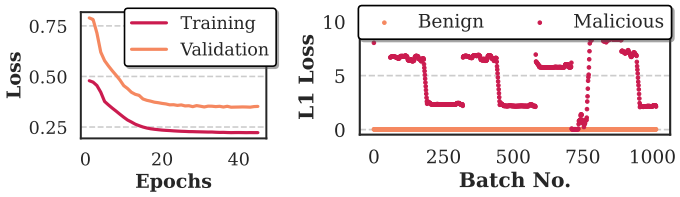

(a) Training (b) Benign and Malicious L1 Loss

Fig. 4: Experimental Results for Security Analytics

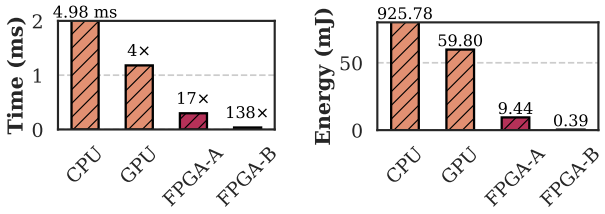

(a) Latency (b) Energy

Fig. 5: Experimental Results for FPGA acceleration

TABLE II: LSTM-AE FPGA Resource Utilization†

Name	Device	LUT (%)	FF (%)	BRAM (%)	DSP (%)
FPGA-A	Alveo U280	16.66	9.50	14.41	3.71
FPGA-B	MPSoC ZCU104	19.04	9.19	5.13	18.98

† Includes LSTM-AE kernel and CPU/FPGA communication

encoder, two LSTM layers in the decoder, and processing input sequences of 120 timesteps. We refer to this model as *Baseline* model and its configuration was chosen as it provided a good balance between reconstruction accuracy and computational efficiency. Fig 4a depicts details on the training process.

Table I demonstrates the performance metrics of the *Baseline* model and its high effectiveness in anomaly detection. Recall, or the true positive rate, stands at 0.9822, signifying that the model successfully captures 98.22% of all actual anomalies. The F1 Score, which balances precision and recall, is at a robust 0.9910, further underscoring the model's reliability. Additionally, Fig 4b depicts in detail the evaluation of the *Baseline* model on the testing set, illustrating the loss for each input batch and the subsequent classification of the input.

Federated Learning: To assess the effectivenes of FL, we evaluated the *Baseline* model in a FL training scenario. The FL scheme was evaluated through an experiment involving 60 federated rounds, each consisting of four local epochs—representing the number of epochs before each client's model weights are centrally aggregated and updated, The results presented in Table I as *Baseline-FL*, indicate that model performance remains largely consistent with that of the centrally trained model Notably, Recall, which is the most important accuracy metric in an anomaly detection application, experienced a minimal degradation of less than 0.6%, underscoring the robustness of the FL approach.

2) Acceleration of Data Analytics on FPGAs: As described in Section III-A2 and depicted in Step ❷ of Fig. 2b PRIVATEER employs a Hardware-Model-Aware DSE. The resulting model, which we refer as *Compact* model, is a regularized overcomplete LSTM-AE that processes inputs of 8 timesteps, containing one LSTM layer in the encoder and another LSTM layer coupled with a linear layer in the decoder.

To evaluate the effectiveness of PRIVATEER's FPGA-accelerated implementation for anomaly detection, we conducted a comparative evaluation focusing on inference latency and energy efficiency across different hardware platforms. Specifically, we evaluated 4 setups: the *Baseline* model deployed on an Intel Xeon Gold 6530 @2.10 GHz (`CPU`) and on a Nvidia V100 (`GPU`), and the *Compact* model deployed on two diverse FPGA-enabled edge nodes, i.e., `FPGA-A`, which consists of an Alveo U280 FPGA, and `FPGA-B`, which is based on the Multiprocessor System-on-a-Chip (MPSoC) ZCU104.

Fig. 5 illustrates the average inference latency and the average energy per inference consumption comparison of each implementation, across 1000 inferences. The `FPGA-B` implementation demonstrated an average latency of 0.036 ms, achieving substantial speedups compared to the rest of the setups, meeting even the most stringent low-latency restrictions. The large difference in latency between `FPGA-A` and `FPGA-B` can be explained by the increased overheads in data transfers that occur in the ALVEO U280. Moreover, `FPGA-B` reduced energy consumption dramatically, requiring 2374×, 153.33×, and 24× less energy per inference than the `CPU`, `GPU`, and `FPGA-A` implementations respectively. Those results indicate that our FPGA-based accelerated data analytics implementation, achieves both goals of energy-efficiency and low-latency.

Resource utilization on the FPGA is also a key metric in our evaluation. As summarized in Table II, the *Compact* model utilizes less than 20% of all resources across all resource types of both `FPGA-A` and `FPGA-B`. The similarity in resource usage can be explaining by the higher communication logic overhead required by the Alveo U280. This efficient use of FPGA resources provides the option of horizontal scaling to improve achieve high-throughput targets, by hosting multiple accelerator instances on the same FPGA device. Furthermore, low resource implementations ensure that additional security mechanisms, such as data encryption modules, could be deployed without exceeding the FPGA's resource limitations. Thus, the implementation remains aligned with PRIVATEER's objectives of not only delivering performance but also enabling secure, resource-constrained edge deployments.

Finally, we assessed the trade-off between model compactness and detection performance. Table I compares the *Baseline* and *Compact* models, demonstrating that the *Compact* model exhibits minimal degradation in key performance metrics—less than 4% across accuracy, precision, recall, and F1 score.

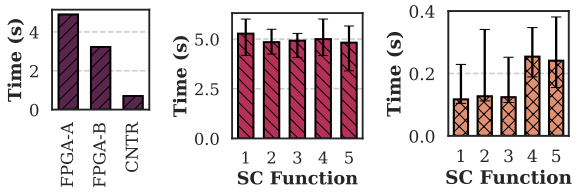

(a) Attestation (b) Store (c) Retrieve

Fig. 6: Experimental results for secure service deployment

This minor reduction in detection accuracy indicates that the *Compact* model remains highly effective for anomaly detection tasks, while significantly benefiting from the reduced latency and energy consumption afforded by FPGA deployment.

B. Edge Service Deployment in Heterogeneous Environments

1) Secure Service Deployment: The methodologies for the integrity verification of deployed services (e.g. data analytics), are evaluated with their execution time. For the FPGA-enabled edge nodes, as in Section IV-A2 we have the two distinct FPGAs, `FPGA-A` and `FPGA-B`. For evaluating the attestation latency of containerized applications, we follow a similar approach to [15], which applies a zero-knowledge methodology. Denoted as `CNTR`, it targets a Raspberry Pi Model 4B equipped with a hardware-based Trusted Platform Module (TPM), that is responsible for securely calculating attestation evidence. Fig. 6a presents the time required for attestation across various edge nodes. We observe that in all platforms, the attestation process takes less than 5sec., having therefore minimal overhead to the existing infrastructure. For evaluating the performance of applications over Intel SGX enclaves, the *Baseline* LSTM-AE model is deployed on SGX enclaves using Gramine [16]. Gramine is a lightweight library OS, that enables running applications on SGX enclaves, without code modifications. The measured latency for the LSTM-AE model is $\sim$18ms., having therefore $\times 3.6$ slowdown over the CPU baseline without an enclave. This slowdown when employing SGX, is due to the memory encryption engines, as well as the designated Enclave Page Cache (EPC) with limited size.

2) Blockchain Interaction: Apart from the execution time of collecting and verifying the security evidence, we also evaluate the time required for the Blockchain functions, namely the SCs. In Fig. 6b and 6c the time required for each of the SC to store and retrieve evidence is depicted. Several functions were implemented to address the unique requirements of various components that interact with DLT, for faster and more effective execution when retrieving data from the Blockchain. In both operations, the amount of data and the complexity of the inquiry determines the response time of each function. Regarding store operations, we observe that small deviation in execution time exist, since the data structures in the DLT are equivalent (Fig. 6b). On contrary, retrieve functions return a list of multiple pieces of evidence based on specific attributed, hence in Fig. 6c large variations in execution times are noticed.

V. Conclusion & Future Steps

PRIVATEER project aims to enhance the security enablers of 5G networks by following a privacy-by-design approach. In this paper, we evaluate several of PRIVATEER's methodologies developed in the first phase of the project, with a focus on developing hardware-accelerated AI models for anomaly detection in 6G networks. Additionally, we outline the infrastructure necessary for secure service deployment across heterogeneous edge environments, that incorporate both CPUs and custom hardware accelerators. In the project's second phase, we plan to expand detection capabilities to address a broader range of attacks in 6G networks. This phase will also integrate security evidence collected from deployed services into a comprehensive privacy mechanism. Service deployment and management will be overseen by a privacy-aware orchestrator, ensuring robust privacy and security management throughout.

References

[1] A. Dogra, R. K. Jha, and S. Jain, "A survey on beyond 5g network with the advent of 6g: Architecture and emerging technologies," *IEEE access*, vol. 9, pp. 67512–67547, 2020.

[2] V.-L. Nguyen, P.-C. Lin, B.-C. Cheng, R.-H. Hwang, and Y.-D. Lin, "Security and privacy for 6g: A survey on prospective technologies and challenges," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 4, pp. 2384–2428, 2021.

[3] D. Masouros, D. Soudris, G. Gardikis, V. Katsarou, M. Christopoulou, G. Xilouris, H. Ramón, A. Pastor, F. Scaglione, C. Petrollini, *et al.*, "Towards privacy-first security enablers for 6g networks: the privateer approach," in *International Conference on Embedded Computer Systems*, pp. 379–391, Springer, 2023.

[4] W. Jiang, B. Han, M. A. Habibi, and H. D. Schotten, "The road towards 6g: A comprehensive survey," *IEEE Open Journal of the Communications Society*, vol. 2, pp. 334–366, 2021.

[5] V. Ziegler, H. Viswanathan, H. Flinck, M. Hoffmann, V. Räisänen, and K. Hätönen, "6g architecture to connect the worlds," *IEEE Access*, vol. 8, pp. 173508–173520, 2020.

[6] M. Cunha, G. Duarte, R. Andrade, R. Mendes, and J. P. Vilela, "Privkit: A toolkit of privacy-preserving mechanisms for heterogeneous data types," in *Proceedings of the Fourteenth ACM Conference on Data and Application Security and Privacy*, pp. 319–324, 2024.

[7] S. B. Prathiba, G. Raja, S. Anbalagan, K. Arikumar, S. Gurumoorthy, and K. Dev, "A hybrid deep sensor anomaly detection for autonomous vehicles in 6g-v2x environment," *IEEE Transactions on Network Science and Engineering*, vol. 10, no. 3, pp. 1246–1255, 2022.

[8] M. A. Ferrag, O. Friha, B. Kantarci, N. Tihanyi, L. Cordeiro, M. Debbah, D. Hamouda, M. Al-Hawawreh, and K.-K. R. Choo, "Edge learning for 6g-enabled internet of things: A comprehensive survey of vulnerabilities, datasets, and defenses," *IEEE Communications Surveys & Tutorials*, 2023.

[9] K. Greff, R. K. Srivastava, J. Koutník, B. R. Steunebrink, and J. Schmidhuber, "Lstm: A search space odyssey," *IEEE transactions on neural networks and learning systems*, vol. 28, no. 10, pp. 2222–2232, 2016.

[10] M. Said Elsayed, N.-A. Le-Khac, S. Dev, and A. D. Jurcut, "Network anomaly detection using lstm based autoencoder," in *Proceedings of the 16th ACM Symposium on QoS and Security for Wireless and Mobile Networks*, pp. 37–45, 2020.

[11] W. Hu, C.-H. Chang, A. Sengupta, S. Bhunia, R. Kastner, and H. Li, "An overview of hardware security and trust: Threats, countermeasures, and design tools," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 40, no. 6, pp. 1010–1038, 2020.

[12] V. Costan, "Intel sgx explained," *IACR Cryptol, EPrint Arch*, 2016.

[13] "Welcome — Besu documentation."

[14] N. C. of Scientific Research "Demokritos" and S. H. (Greece), "NCSRD-DS-5GDDoS: 5G Radio and Core metrics containing sporadic DDoS attacks," Oct. 2024. https://doi.org/10.5281/zenodo.13900057.

[15] H. B. Debes and T. Giannetsos, "Zekro: Zero-knowledge proof of integrity conformance," in *Proceedings of the 17th International Conference on Availability, Reliability and Security*, pp. 1–10, 2022.

[16] "Gramine - a Library OS for Unmodified Applications."