Towards AI-Driven Security in the Edge—Cloud Continuum: A Unified Architectural Perspective

Pedro R. Tomas[1][0000−0001−7938−4972], João Fernandes[1][0009−0008−9673−7416], Davide Miola[2][0009−0001−9996−0030], Grigorios Chrysos[3][0000−0001−8398−8365], Rajat Kandoi[4][0009−0008−9429−6581], Amir Javadpour[5][0000−0002−4932−1660], Riccardo Sisto[2][0000−0002−3142−2383], Matthias Dippold[6][0009−0001−9708−9886], Despina Kopanaki[3][0000−0003−3933−1619], Sotiris Ioannidis[3][0000−0001−9340−2241], Sofia Maragkou[6][0000−0001−6823−4223], Luís Cordeiro[1][0000−0001−5471−7064], and Tarik Taleb[7][0000−0003−1119−1239]

[1] OneSource, Consultoria Informática, Lda., Coimbra, Portugal {pedro.tomas, joao.fernandes, cordeiro}@onesource.pt
[2] Politecnico di Torino, Turin, Italy {davide.miola, riccardo.sisto}@polito.it
[3] Technical University of Crete, Greece {gxrysos, dkopanaki, sioannidis}@tuc.gr
[4] Ericsson Research, Finland rajat.kandoi@ericsson.com
[5] MOSA!C Lab / ICTFICIAL Oy, Finland a.javadpour87@gmail.com
[6] TU Wien, Austria {e1609401@student., sofia.maragkou@}tuwien.ac.at
[7] Ruhr University Bochum, Germany tarik.taleb@rub.de

Abstract. The evolution towards Beyond 5G (B5G)/6G systems is accelerating the emergence of distributed Edge–Cloud environments, where computation and intelligence span heterogeneous and dynamic infrastructures. While this enables latency-sensitive and data-intensive services, it also expands the attack surface, rendering traditional perimeter-based security insufficient.

In this context, Artificial Intelligence (AI)-driven security is emerging as a key approach for enabling adaptive monitoring, intelligent threat detection, and automated response. This paper presents an integration-oriented perspective on AI-driven security in the Edge–Cloud continuum. It identifies the main security requirements and design dimensions, and analyses representative building blocks, including extended Berkeley Packet Filter (eBPF)-based monitoring, hardware-accelerated intrusion detection, federated intelligence, and privacy-preserving mechanisms.

Based on these elements, the paper outlines a unified architectural framework that integrates telemetry collection, AI-driven detection, distributed learning, and trusted orchestration into an end-to-end security pipeline. The approach is further supported by insights from the ELASTIC and 6G-PATH projects, highlighting its applicability in realistic deployment scenarios. Finally, the paper discusses key challenges related to scalability, trust, and robustness in next-generation Edge–Cloud systems.

Keywords: Edge–Cloud Continuum · AI-Driven Security · eBPF · AI-Based Intrusion Detection · Federated Learning · 6G

1 Introduction

The evolution towards 6G networks is driving the emergence of a highly distributed computing paradigm, where intelligence, storage, and processing capabilities are seamlessly deployed across the Edge–Cloud continuum. This transformation enables a wide range of latency-sensitive and data-intensive applications, but also significantly expands the attack surface, introducing new security challenges that cannot be addressed by traditional, perimeter-based approaches [17]. In such environments, security must be intrinsically embedded across all system layers, from resource-constrained edge devices to cloud-native infrastructures.

AI is increasingly recognised as a key enabler for addressing these challenges, providing advanced capabilities for anomaly detection, threat prediction, and automated response. However, the adoption of AI in security introduces a set of new complexities. Security mechanisms must operate under strict latency and resource constraints, ensure data privacy across distributed domains, and remain robust against adversarial manipulation [3]. At the same time, they must remain deployable, scalable, and trustworthy in operational Edge–Cloud environments.

1.1 Research Gap and Paper Scope

Although substantial progress has been made in areas such as telemetry collection, AI-based intrusion detection, federated intelligence, and privacy-preserving analytics, these directions are often studied in isolation. As a result, fewer works address how such mechanisms can be combined into coherent, end-to-end security pipelines for practical Edge–Cloud environments.

This paper addresses this gap by analysing and positioning a set of complementary AI-driven security mechanisms within the Edge–Cloud continuum. In particular, we consider kernel-level telemetry collection through eBPF-based monitoring, AI-based intrusion detection techniques, and distributed intelligence approaches based on federated learning and privacy-preserving mechanisms. These components are examined not as standalone solutions, but as building blocks of a broader security pipeline. Accordingly, the scope of this paper is not to provide an exhaustive survey of all security solutions for 6G systems, but rather to focus on a representative set of complementary mechanisms that support an integration-oriented architectural perspective.

1.2 Main Objective and Contributions

Building on this analysis, we outline a unified architectural perspective that organises these mechanisms into a layered framework, highlighting their interactions and identifying key challenges related to interoperability, scalability, and trust. While full integration remains an open issue, this work provides a structured foundation for understanding how such technologies can be combined in future systems. More specifically, the main objective of this paper is to clarify how monitoring, intelligent threat detection, distributed intelligence, and

privacy-preserving security functions can be positioned and connected within a common Edge–Cloud security architecture.

The main contributions of this paper are as follows. First, we identify key requirements and design dimensions for AI-driven security in the Edge–Cloud continuum (Section 2). Second, we analyse a set of complementary building blocks, including eBPF-based monitoring, AI-driven intrusion detection, distributed intelligence via federated learning, and privacy-preserving mechanisms (Section 3). Third, we outline a unified architectural perspective integrating these elements into an end-to-end security pipeline (Section 4). Finally, we discuss deployment perspectives and open challenges based on insights from the ELASTIC and 6G-PATH projects (Section 5). Section 6 highlights open challenges and research directions, and Section 7 concludes the paper.

2 AI-Driven Security in the Edge–Cloud continuum

Building on the challenges outlined in Section 1, this section outlines the key requirements and design dimensions that underpin AI-driven security in the Edge–Cloud continuum.

2.1 Security Requirements

The transition to AI-driven security in the Edge–Cloud continuum requires a fundamental shift from perimeter-based protection to distributed, trust-centric architectures. In such environments, security mechanisms must satisfy several key requirements.

First, security solutions must be trustworthy and robust against adversarial manipulation. As AI becomes both a defensive tool and a potential attack target, systems must ensure resilience against poisoning, evasion, and model manipulation attacks [10]. Second, scalability and decentralisation are essential. The massive scale and distribution of connected devices in 6G environments require security intelligence to be distributed across edge and cloud domains, avoiding central bottlenecks while preserving data locality [24]. Third, explainability and verifiability are critical for operational deployment. AI-driven decisions must be interpretable and auditable, particularly in mission-critical scenarios where opaque behaviour may lead to unreliable or unsafe outcomes [8]. Fourth, security mechanisms must be resource-aware and low-latency. Edge devices operate under strict computational and energy constraints, requiring lightweight models and efficient execution while maintaining real-time responsiveness [18]. Finally, end-to-end trust and identity management must be ensured across the entire service chain. This includes consistent policy enforcement, secure identity propagation, and integrity guarantees across heterogeneous domains [7].

Together, these requirements define the operational constraints for AI-driven security systems and motivate the need for integrated, distributed security architectures in the Edge–Cloud continuum.

2.2 Design Dimensions

Based on these requirements, AI-driven security in the Edge–Cloud continuum can be characterised along key design dimensions that capture the trade-offs of distributed intelligent systems.

A first dimension is distributed intelligence, where decision-making is decentralised across edge and cloud nodes to enable scalability and localised threat detection [4]. Closely related is edge-centric inference and training, where AI models are deployed and updated near data sources, often through federated learning, improving privacy and reducing communication overhead [11]. Privacy preservation forms another critical dimension, supported by mechanisms such as federated learning and trusted execution environments, enabling secure data processing in distributed environments [2]. Low-latency operation is also essential, requiring efficient execution through hardware acceleration and in-kernel mechanisms such as eBPF to support real-time security enforcement [12]. In addition, resource awareness ensures that security functions are adapted to constrained edge devices through lightweight models and optimised deployment strategies [9]. Finally, high autonomy enables dynamic adaptation through advanced orchestration mechanisms, allowing security functions to scale and evolve with changing system conditions [13].

These dimensions define the design space and guide the development of integrated AI-driven security architectures for the Edge–Cloud continuum.

3 AI-Driven Security Mechanisms in the Edge–Cloud Continuum

Building upon the security requirements and design dimensions presented in Section 2, this section introduces the AI-driven security mechanisms that realise these principles in the Edge–Cloud continuum. These mechanisms constitute the core building blocks of the proposed framework, enabling distributed intelligence, privacy-preserving operation, low-latency enforcement, and autonomous security management.

3.1 Monitoring and Telemetry Collection

The foundational step in any AI-driven intrusion detection pipeline for cloud-native environments is the capture and injection of network data and derived metrics into downstream processing stages. Traffic traversing edge and cloud nodes must be intercepted and forwarded to the intrusion detection pipeline, making the choice of capture mechanism critical in terms of fidelity, latency, and scalability.

Two principal strategies exist. Hardware-assisted traffic replication relies on specialised equipment such as managed switches or SmartNICs to mirror traffic without CPU overhead, but assumes the availability of compatible infrastructure. In contrast, software-based capture intercepts traffic directly on the host using mechanisms such as network taps and Berkeley Packet Filter (BPF), offering greater flexibility for heterogeneous Edge–Cloud environments.

However, network interface-level capture becomes ineffective in cloud-native infrastructures where Container Network Interface (CNI) plugins or Service Mesh frameworks enforce pervasive encryption. In such settings, traffic is encrypted within the application namespace before reaching the host network interface, rendering it opaque to traditional capture mechanisms. A possible mitigation is the deployment of capture agents within each Pod to intercept traffic prior to encryption. However, this approach introduces operational complexity, potential interference with workloads, and scalability limitations in dynamic environments.

To address these limitations, we adopt an eBPF-based traffic capture mechanism that operates at Layer-4 within the host kernel, providing uniform visibility across all co-located workloads without per-Pod instrumentation. The key advantage of this approach lies in its position relative to the encryption boundary imposed by Service Mesh and CNI components (Figure 1). For outbound traffic, interception occurs before encryption, while for inbound traffic it occurs after decryption, ensuring visibility of cleartext data in both directions.

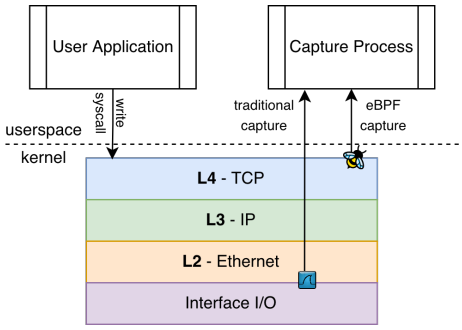

Fig. 1: Comparison of traditional traffic capture versus eBPF-based capture.

Furthermore, because eBPF programs execute within the shared host kernel rather than individual container namespaces, the system achieves comprehensive monitoring through a single, non-intrusive instrumentation point. This design is well suited to dynamic cloud-native environments, balancing deployment transparency, encryption-agnostic visibility, and scalability.

3.2 Hardware AI-based Threat Detection Framework

AI-driven intrusion detection enables the identification of both known and previously unseen attack patterns through data-driven analysis, complementing traditional rule-based approaches.

AI-IDS (Fig. 2) is a hardware-accelerated intrusion detection framework that combines Field-Programmable Gate Array (FPGA)-based processing with Machine Learning (ML)-driven analytics to address both known and emerging threats. The system integrates low-latency hardware monitoring with adaptive detection mechanisms, enabling real-time traffic analysis and high-precision intrusion detection.

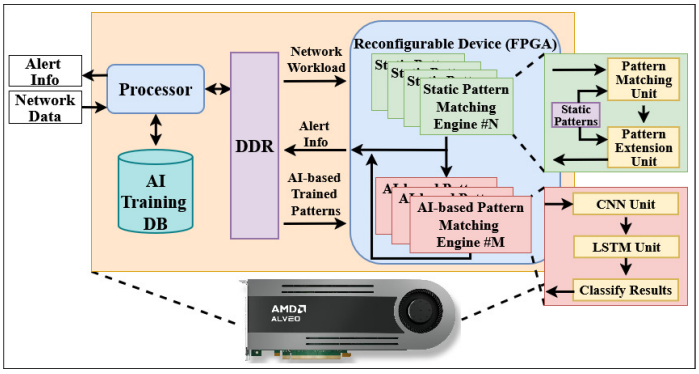

Fig. 2: Hardware-based AI-IDS architecture.

The framework adopts a hybrid hardware–software architecture, where a host-side interface preprocesses and forwards eBPF-derived traffic to the FPGA subsystem. The hardware design integrates two main components: (i) a pattern-matching engine for signature-based detection, and (ii) an AI-based module for adaptive pattern refinement.

The pattern-matching engine performs real-time inspection using predefined signatures, enabling deterministic, line-rate detection of known attacks with minimal latency [16].

The AI-driven module complements this capability through semi-supervised learning (e.g., Convolutional Neural Networkss (CNNs) and Long Short-Term Memory (LSTM)) to identify anomalous traffic patterns [5]. Suspicious instances are collected and used to periodically retrain the models, generating new signatures corresponding to previously unseen threats. These are translated into hardware-compatible representations and deployed to the FPGA, forming a closed-loop adaptation mechanism.

The FPGA-based implementation leverages parallel processing to enable high-throughput, low-latency inspection of multiple data streams, making it suitable for Edge–Cloud and Internet of Things (IoT) environments. Overall, this approach supports key design dimensions identified in Section 2.2, including low-latency operation, resource awareness, and adaptive intelligence.

3.3 Distributed and Federated Intelligence

The rapid growth of connected devices and distributed infrastructures has led to massive volumes of data that must be processed across the Edge–Cloud continuum [1]. Traditional centralised machine learning approaches, where data is aggregated and processed at a single location, are increasingly inadequate in such environments due to bandwidth constraints, latency requirements, and privacy concerns.

Federated Learning (FL) has emerged as a key paradigm to address these challenges by enabling collaborative model training across distributed nodes without requiring the exchange of raw data. Instead, each participant trains a local model and shares only model updates, which are aggregated to produce a

global model. This approach reduces communication overhead while preserving data locality and privacy, making it particularly suitable for edge-native security applications. Different FL configurations can be adopted depending on the system architecture and data distribution, including hierarchical approaches that introduce intermediate aggregation layers at the edge to improve scalability and reduce latency.

DeepGuardian: A FL-based approach DeepGuardian (DG) is a Network Intrusion Detection Prevention System (NIDPS) with a strong focus on data security and privacy. Relying on the core foundations of the Holistic Security and Privacy Framework (HSPF) [20], DG extends its core principles by combining FL, Continual Learning (CL), and unsupervised ML, with supervised ML to perform network traffic attacks classification, Privacy Enhancing Technologys (PETs) to ensure data privacy and security, an intuitive dashboard and policy enforcement mechanisms, aiming to provide scalable, privacy-preserving, and End-to-End (E2E) adaptive security for cloud-native environments.

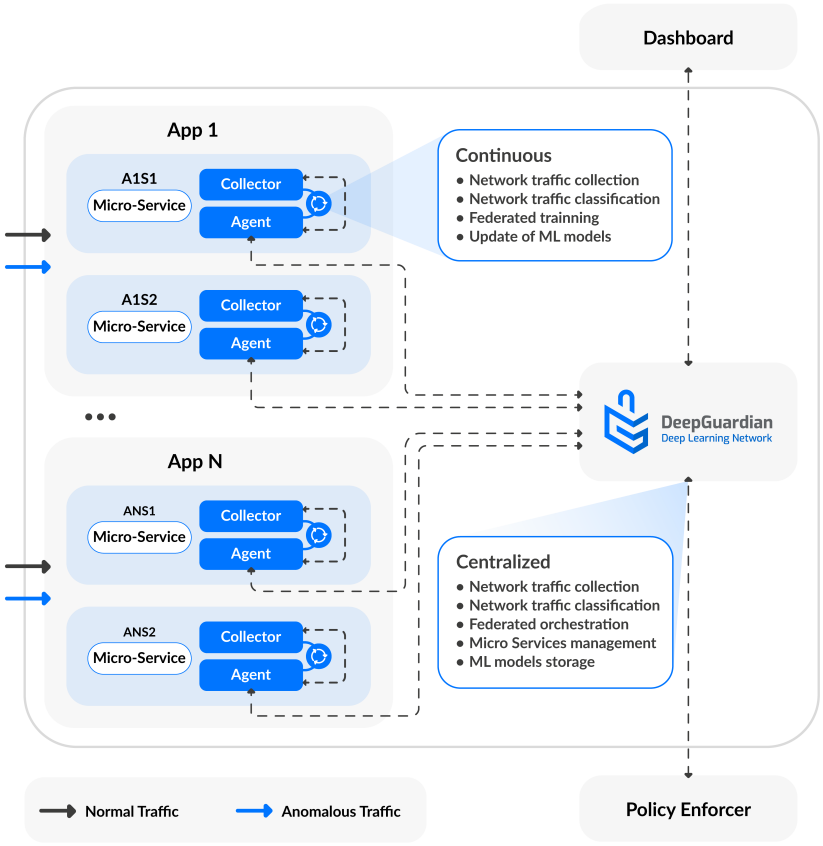

Fig. 3: DG Architecture

Figure 3 presents the DG architecture, which is composed of the following key architecture components: (i) Aggregator, (ii) Agent, (iii) Collector, (iv) Policy Enforcer, (v) Dashboard. The Aggregator acts as the main framework unit, coordinating the Federated Training process and handling the management and distribution of the various ML models used for anomaly detection across federated agents. The Collector component gathers inbound and outbound network traffic traversing the primary network interface. This data is aggregated into flows and stored until classification occurs. Subsequently, the Agent, where the federated logic resides, is responsible for performing inference on the stored data and detecting potential malicious flows. Additionally, the Agent periodically executes federated training rounds to keep the ML model updated and effective in identifying anomalies within the continuously evolving network communication landscape. The Dashboard provides system observability over all the applications being protected, while also enabling the Human-in-the-loop functionality by allowing the human operator to override the AI's classification on a flow basis. Lastly, the Policy Enforcement [21] mechanism enables the enforcement of network traffic policies in cloud-native environments, supporting different policy enforcement mechanisms (i.e., OPA, Istio, Calico, Cilium, Kyverno).

3.4 Privacy-Preserving Security Mechanisms

In distributed AI-driven security environments, sensitive data and model updates are exchanged across multiple domains, raising significant privacy and trust concerns. To address these challenges, a range of PETs can be integrated into the security pipeline.

Secure Multi-Party Computation (MPC) enables collaborative computation across multiple participants without exposing their private inputs, supporting secure aggregation and distributed analytics [25]. Differential Privacy (DP) provides formal guarantees against information leakage by introducing controlled noise into model updates, making it particularly suitable for federated learning scenarios [15]. Homomorphic Encryption (HE) allows computations to be performed directly on encrypted data, ensuring confidentiality throughout the processing pipeline [14, 6]. Zero-Knowledge Proof (ZKP) enable the verification of computations without revealing underlying data, facilitating trust across administrative domains [23]. Finally, Trusted Execution Environments (TEEs) provide hardware-based isolation and attestation mechanisms to ensure the integrity and confidentiality of sensitive operations [19].

These mechanisms are complementary and can be combined to enforce privacy and trust at different stages of the AI-driven security pipeline. In particular, they play a critical role in enabling secure model training, aggregation, and deployment across distributed edge–cloud environments.

DG currently supports MPC, HE, ZKP and DP to ensure the data privacy and security of the data, especially of the model weights that are being shared through the network and that contain a bigger exposure to potential data breaches. Moreover, Tomas et al. [22] describe some preliminary work reporting on the impact of adding PETs to the DG foundations, the HSPF framework.

4 Towards a Unified Edge–Cloud Security Architecture

This section proposes a unified architecture that integrates the individual components from Section 3, including traffic capture, hardware-accelerated threat detection, distributed intelligence, and privacy-preserving mechanisms, into a cohesive Edge–Cloud security framework that relies on the key principles of continuous training, detection, and reaction. This architecture, Fig. 4, is organised into four functional domains: monitoring and telemetry collection via eBPF, AI-driven threat detection, distributed intelligence through federated learning and privacy-enhancing technologies, and system observability with policy enforcement. These components are connected through continuous data and control flows, enabling a closed-loop security pipeline across the Edge–Cloud continuum. At the data plane level, the system utilises an eBPF-based traffic capture

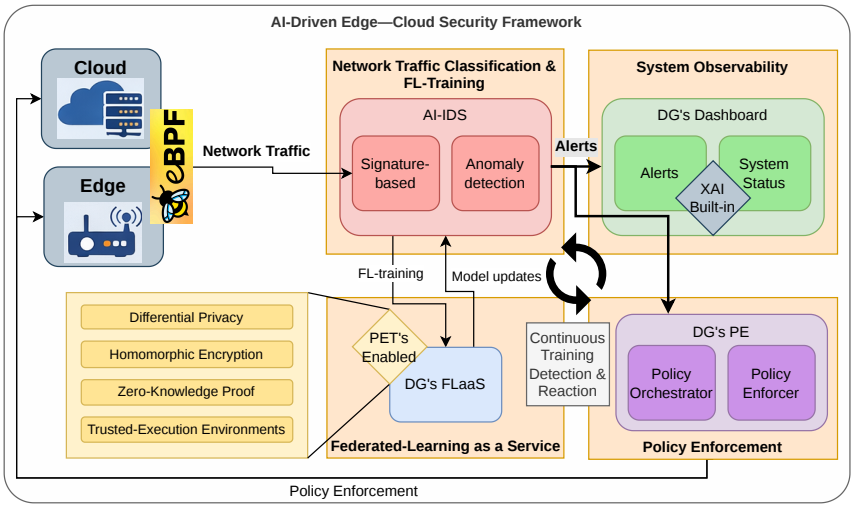

Fig. 4: Unified Edge AI-based security framework

mechanism (Section 3.1) to intercept traffic at the transport level directly within the operating system kernel. Captured traffic is then forwarded to the AI-IDS framework (Section 3.2), which is in charge of performing network traffic classification, via encrypted Kafka brokers supporting authentication and authorisation mechanisms. Leveraging its underlying FPGA support, the AI-IDS enables deterministic, line-rate inspection through hardware-accelerated signature matching while simultaneously incorporating adaptive machine learning models to detect both known and previously unseen threats.

To support large-scale and distributed deployments, the framework employs the DG Federated Learning as a Service (FLaaS), an on-demand federated and distributed intelligence mechanism, which enables edge nodes to collaboratively train intrusion detection models while keeping raw data local. The FLaaS mechanism also supports hierarchical aggregation via intermediate edge aggregators, which improves scalability and reduces communication overhead.

Privacy and trust within the framework are reinforced by the integration of PETs, including DP, HE and ZKP, which protect sensitive data and model updates during both training and inference. Additionally, TEEs provide secure enclaves for critical operations such as secure aggregation and policy enforcement, utilising remote attestation to establish trust across the Edge-Cloud continuum.

The DG Dashboard provides system observability by presenting generated alerts complemented with known vulnerabilities linked to the source systems being protected. To ensure transparency, the Dashboard incorporates Explainable AI (XAI) by leveraging post-hoc methods to provide interpretability over the algorithms' decisions. The Dashboard also enables Human-in-the-Loop interaction within this architecture by enabling the operator to override the algorithm's decision, which is also considered during the next federated training round. Lastly, the DG Policy Enforcement (PE) enables the framework's reaction capabilities by providing intelligent policy enforcement mechanisms and ensuring deployment both on Edge devices and on Cloud infrastructures.

5 Deployment Perspectives, Application Domains, and Business Implications

The practical realisation of the unified Edge–Cloud security architecture requires careful consideration of deployment constraints, resource heterogeneity, and domain-specific operational requirements. In real-world environments, AI-driven security mechanisms must be flexibly mapped across distributed infrastructures while maintaining performance, scalability, and trust.

5.1 Deployment and Validation Scenarios

Edge–Cloud infrastructures are inherently heterogeneous, encompassing resource-constrained gateways, edge servers with hardware acceleration capabilities, and cloud platforms offering high computational capacity and confidential computing features. This diversity necessitates adaptive deployment strategies in which security functions are placed according to available resources and operational requirements.

In the considered framework, monitoring and low-latency threat detection are performed at the edge, where FPGA-based acceleration enables real-time processing of network traffic. More computationally intensive tasks, such as model aggregation and global training, are executed in cloud environments, which provide the required scalability and processing capabilities. Policy enforcement mechanisms are deployed at gateway nodes, enabling control over data flows at trust boundaries. Confidential computing technologies, such as TEEs, are selectively employed to protect sensitive operations depending on data confidentiality requirements and trust assumptions. These deployment principles have been validated in smart factory scenarios, including equipment health monitoring and cross-site energy optimisation. In such settings, sensor data is processed locally at edge nodes, while collaborative model training is performed across multiple sites without transferring raw data. This approach enables low-latency inference

while preserving data privacy. Furthermore, the use of heterogeneous TEE technologies (e.g., AMD SEV-SNP and Intel TDX) demonstrates that trust can be established across administrative domains through remote attestation, even in environments with diverse hardware platforms.

5.2 Application Domains and Business Implications

The proposed architectural approach is applicable across multiple domains that share requirements for distributed processing, real-time decision-making, and privacy preservation. In industrial environments, such as smart manufacturing, the framework supports predictive maintenance and operational optimisation through distributed analytics. Similarly, telecommunications infrastructures, including 5G and 6G networks, can leverage the framework to support distributed, AI-driven security functions across operator domains. In this context, projects such as 6G-PATH provide testbed-driven environments for validating these capabilities in realistic conditions, enabling the experimentation of Edge–Cloud security mechanisms across multiple vertical use cases and deployment scenarios.

From a business perspective, this architecture enables new operational models for security deployment. Network operators and service providers can offer security-as-a-service solutions, where intrusion detection, model management, and policy enforcement are orchestrated across distributed infrastructures. By centralising intelligence while preserving data locality, the framework reduces operational complexity and enables scalable deployment of advanced security capabilities across diverse environments.

6 Open Challenges and Research Directions

The proposed framework highlights several open challenges that must be addressed to enable scalable and operational deployment in real-world Edge–Cloud environments.

At the monitoring layer, scalability remains a primary concern. The combination of eBPF-based telemetry, AI-driven detection, and federated learning generates large volumes of data and coordination overhead. This necessitates efficient mechanisms for filtering, feature extraction, and data reduction to prevent overwhelming downstream components. Additionally, tracing-based eBPF instrumentation introduces portability and maintainability challenges, as probes attached to internal kernel functions are sensitive to changes across kernel versions. Ensuring robustness therefore requires continuous adaptation to evolving system environments.

At the detection layer, hardware-accelerated AI-IDS solutions provide high performance but reduce flexibility in heterogeneous deployments, highlighting the need for adaptive hardware–software co-design. Similarly, federated learning introduces challenges in dynamic and adversarial environments, including robustness to poisoning attacks, handling non-IID data distributions, and ensuring reliable model convergence across distributed participants.

From a trust and privacy perspective, establishing secure collaboration across administrative domains remains a key challenge. Mechanisms such as TEEs and remote attestation are essential for ensuring integrity, but their integration into large-scale distributed systems introduces additional complexity. Furthermore, privacy-enhancing technologies (e.g., differential privacy, homomorphic encryption, and secure computation) introduce trade-offs between privacy, performance, and detection accuracy, requiring context-aware optimisation.

Finally, system-level orchestration of distributed security functions remains an open issue. Coordinating monitoring, detection, learning, and enforcement across heterogeneous environments demands advanced and adaptive orchestration strategies, potentially leveraging AI-driven control loops.

Addressing these challenges will be critical to enabling robust, scalable, and trustworthy AI-driven security architectures for future Edge–Cloud and B5G/6G systems, and represents a key direction for ongoing research and standardisation efforts.

7 Conclusion

This paper presented an integration-oriented perspective on AI-driven security in the Edge–Cloud continuum, addressing the increasing complexity and distributed nature of B5G/6G environments. Rather than treating monitoring, threat detection, distributed intelligence, and privacy-preserving mechanisms as isolated components, the paper positioned them as complementary elements of an end-to-end security pipeline.

Building on this view, we outlined a unified architectural framework that connects telemetry collection, AI-driven detection, federated intelligence, and trusted orchestration within a coherent Edge–Cloud security architecture. The relevance of this approach was further illustrated through deployment perspectives and validation insights, highlighting its applicability in realistic environments. At the same time, several challenges remain open, particularly in terms of scalability, adversarial robustness, explainability, privacy–utility trade-offs, and trust across heterogeneous infrastructures. Addressing these challenges is essential for moving from fragmented solutions towards fully integrated and operational AI-driven security systems.

Overall, this work provides a structured foundation for advancing AI-driven security in the Edge–Cloud continuum and supports future research towards scalable, trustworthy, and deployable security architectures for next-generation distributed systems.

Acknowledgments. This work was supported in part by: Smart Networks and Services Joint Undertaking (SNS JU) under the 6G-PATH (Grant No. 101139172) and ELASTIC (Grant No. 101139067) projects; the Horizon Europe Research and Innovation programme HORIZON-CL3-2022-CS-01-01 under the MIRANDA project (Grant No. 101168144); PNRR-NGEU, which has received funding from the MUR - DM 117/2023. The views expressed in this work are those of the authors and do not necessarily reflect the official position of the projects or the funding bodies.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Al-Quraan, M., Mohjazi, L., Bariah, L., Centeno, A., Zoha, A., Arshad, K., Assaleh, K., Muhaidat, S., Debbah, M., Imran, M.A.: Edge-native intelligence for 6g communications driven by federated learning: A survey of trends and challenges. IEEE Transactions on Emerging Topics in Computational Intelligence **7**(3), 957–979 (2023). https://doi.org/10.1109/TETCI.2023.3251404
2. Albshaier, L., Almarri, S., Albuali, A.: Federated learning for cloud and edge security: A systematic review of challenges and ai opportunities. Electronics **14**(5), 1019 (2025)
3. European Union Agency for Cybersecurity (ENISA): Artificial intelligence cybersecurity challenges. Tech. rep., ENISA (2020),
4. Feng, Y., Qi, Y., Li, H., Wang, X., Tian, J.: Leveraging federated learning and edge computing for recommendation systems within cloud computing networks. In: Third International Symposium on Computer Applications and Information Systems (ISCAIS 2024). vol. 13210, pp. 279–287. SPIE (2024)
5. Imani, F., Kargar, M., Assadzadeh, A., Bayani, A.: Integrating cnn-lstm networks with statistical filtering techniques for intelligent iot intrusion detection. In: 2024 8th International Conference on Smart Cities, Internet of Things and Applications (SCIoT). pp. 189–195. IEEE (2024)
6. Javadpour, A., Ja'Fari, F., Taleb, T., Benzaid, C., Rosa, L., Tomás, P., Cordeiro, L.: Deploying testbed docker-based application for encryption as a service in kubernetes. In: 2024 International Conference on Software, Telecommunications and Computer Networks (SoftCOM). pp. 1–7. IEEE (2024)
7. John, W., Karapantelakis, A., Mouradian, C.: Supporting ai-driven mobile applications with a 6g ai compute continuum. Tech. rep., Ericsson Research (2026)
8. Kaur, N., Gupta, L.: Securing the 6G-IoT environment: A framework for enhancing transparency in artificial intelligence decision-making through explainable artificial intelligence. Sensors (Basel) **25**(3) (Jan 2025)
9. Kostopoulos, S., Papatsaroucha, D., Kefaloukos, I., Markakis, E.K.: eidps: A real-time ebpf-based and machine learning-powered network intrusion detection and prevention solution. In: 2025 6th International Conference in Electronic Engineering & Information Technology (EEITE). pp. 1–8. IEEE (2025)
10. Kumar, R., Dutta, J., Vamsi, N., Varri, U., Puthal, D.: Next-generation security in the 6g era: The role of ai in safeguarding future networks. IEEE Access **PP**, 1–1 (01 2026). https://doi.org/10.1109/ACCESS.2025.3650208
11. Li, H., Ge, L., Tian, L.: Survey: federated learning data security and privacy-preserving in edge-internet of things. Artificial Intelligence Review **57**(5), 130 (2024)
12. Lim, S.Y., Prasad, T., Han, X., Pasquier, T.: Safebpf: Hardware-assisted defense-in-depth for ebpf kernel extensions. In: Proceedings of the 2024 on Cloud Computing Security Workshop. pp. 80–94 (2024)
13. Liyanage, M., Pham, Q.V., Dev, K., Bhattacharya, S., Maddikunta, P.K.R., Gadekallu, T.R., Yenduri, G.: A survey on zero touch network and service management (zsm) for 5g and beyond networks. Journal of Network and Computer Applications **203**, 103362 (2022)

14. Marcolla, C., Sucasas, V., Manzano, M., Bassoli, R., Fitzek, F.H.P., Aaraj, N.: Survey on fully homomorphic encryption, theory, and applications. Proceedings of the IEEE **110**(10), 1572–1609 (2022). https://doi.org/10.1109/JPROC.2022.3205665
15. Ouadrhiri, A.E., Abdelhadi, A.: Differential privacy for deep and federated learning: A survey. IEEE Access **10**, 22359–22380 (2022). https://doi.org/10.1109/ACCESS.2022.3151670
16. Papadogiannaki, E., Chrysos, G., Georgopoulos, K., Ioannidis, S.: A reconfigurable ids framework for encrypted and non-encrypted network data in supply chains. In: 2023 International Conference on Engineering and Emerging Technologies (ICEET). pp. 1–6. IEEE (2023)
17. Pérez Palma, N., Skarmeta Gómez, A., Bisson, P.: Innovative approaches for 6G security: Challenges, solutions,and impact (2025). https://doi.org/10.5281/zenodo.14619619
18. Rabbi Ahmed, K., Hosien, M., Tawkir Nesar, S., Khan, M., Karim, M., Chowdhury, M.A.R., Bazan-Antequera, R.: Ai-enhanced adaptive network security for 6g and edge computing. In: 2025 International Conference on Quantum Photonics, Artificial Intelligence, and Networking (QPAIN). pp. 1–6 (07 2025). https://doi.org/10.1109/QPAIN66474.2025.11172162
19. Tariq, A., Serhani, M.A., Sallabi, F.M., Barka, E.S., Qayyum, T., Khater, H.M., Shuaib, K.A.: Trustworthy federated learning: A comprehensive review, architecture, key challenges, and future research prospects. IEEE Open Journal of the Communications Society **5**, 4920–4998 (2024). https://doi.org/10.1109/OJCOMS.2024.3438264
20. Tomas, P.R., Felix, P., Rosa, L., Gomes, A.S., Cordeiro, L.: A novel approach for continual and federated network anomaly detection. In: Arai, K. (ed.) Proceedings of the Future Technologies Conference (FTC) 2024, Volume 4. pp. 212–225. Springer Nature Switzerland, Cham (2024)
21. Tomas, P.R., Silva, S., Neto, M., Proença, J., Rosa, L., Cordeiro, L., Taleb, T., Cruz, T.: Network policy enforcement in cloud-native environments. In: Papaleonidas, A., Pimenidis, E., Papadopoulos, H., Chochliouros, I. (eds.) Artificial Intelligence Applications and Innovations. AIAI 2025 IFIP WG 12.5 International Workshops. pp. 195–208. Springer Nature Switzerland, Cham (2025)
22. Tomás, P., Kamali Poorazad, S., Benzaïd, C., Rosa, L., Proença, J., Taleb, T., Cordeiro, L.: Enhancing federated learning with homomorphic encryption and multi-party computation for improved privacy. In: 2024 IEEE Future Networks World Forum (FNWF). pp. 891–896 (2024). https://doi.org/10.1109/FNWF63303.2024.11028749
23. Xing, Z., Zhang, Z., Zhang, Z., Li, Z., Li, M., Liu, J., Zhang, Z., Zhao, Y., Sun, Q., Zhu, L., Russello, G.: Zero-knowledge proof-based verifiable decentralized machine learning in communication network: A comprehensive survey. IEEE Communications Surveys & Tutorials **28**, 985–1024 (2026). https://doi.org/10.1109/COMST.2025.3561657
24. Yalli, J., Hasan, M., Badawi, A.: Internet of things (iot): Origins, embedded technologies, smart applications and its growth in the last decade. IEEE Access **12**, 91357 – 91382 (06 2024). https://doi.org/10.1109/ACCESS.2024.3418995
25. Zhou, I., Tofigh, F., Piccardi, M., Abolhasan, M., Franklin, D., Lipman, J.: Secure multi-party computation for machine learning: A survey. IEEE Access **12**, 53881–53899 (2024). https://doi.org/10.1109/ACCESS.2024.3388992