information

MDPI

Article

Security Assurance in 5G-Advanced (3GPP Release 18): Protecting Edge Computing, Network Automation, and Non-Public Networks

Ehigiator Egho-Promise [1,*], Ekereuke Udoh [1], Edita Gashi [1] and Muhammad Maaz Rehan [2]

[1] Department of Computing, QA Higher Education, London EC1R 4TF, UK; ekereuke.udoh@qa.com (E.U.); edita.gashi@qa.com (E.G.)
[2] Department of Computer Science, University of Lancashire, Preston PR1 2HE, UK; mmrehan@lancashire.ac.uk
* Correspondence: ehigiator.eghopromise@qa.com; Tel.: +44-7466076674

Abstract

5G-Advanced (3GPP Release 18) architectural changes include multi-access edge computing (MEC) architectural changes, network automation, and non-public networks (NPNs). It is important to note that even though these advancements provide substantial performance advantages, they destroy fixed-perimeter security models, providing a distributed attack surface. The use of current security assessment strategies, which are usually non-fluid and isolated, is inadequate to offer the required runtime security health assurance needed in such fluid environments. This study presents a new security assurance framework (SAF) that would be used to provide ongoing evidence-based protection on core, edge, and private network domains. This framework employs a four-layer architecture, including monitoring, analytics (LM), policy engine, and enforcement, to convert security periodically audited to a dynamic threat-control-metric evidence chain. A 96% attack detection rate and a 99.8% reduction in response time (with a mean of 20.1 s) are proven by validation on an emulated 5G-Advanced testbed (approximating Release 18 features using Open5GS (v2.7.2 Rel-17, community developed, Seoul, Republic of Korea and custom extensions) based on a design science research (DSR) paradigm. Although the overhead (13% CPU, 21.4% memory) is manageable, the findings prove that all-time, multi-domain assurance is crucial to the healthy functioning of 5G-Advanced and is a key roadmap to autonomous 6G security.

Keywords: 5G-Advanced (release 18); security assurance; edge computing (MEC); network automation; non-public networks (NPNs)

Academic Editors: Malinka Spasova Ivanova, Galina Bogdanova and Tomaž Klobučar

Received: 19 May 2026
Revised: 22 June 2026
Accepted: 25 June 2026
Published: 9 July 2026

1. Introduction

Release 18 is a radical evolutionary change in the performance and architectural philosophy [1]. Although standard 5G networks (Releases 15 and 16) formed a platform for better mobile broadband and ultra-reliable communications, they mainly worked within a comparatively fixed security perimeter [2]. Conversely, Release 18 magnifies three important expansions: multi-access edge computing (MEC), AI-driven network automation, and non-public networks (NPNs), which convert the long-standing model of security based upon boundaries [3,4]. MEC puts processing at the network edge to reduce latency, which necessarily forms a large and physically distributed attack surface. At the same time, with the introduction of AI-based automation, it is possible to have

self-optimizing and self-healing networks that run as quickly as machines can, i.e., security incidents or policy breaches can spread more rapidly than human operators can perform any useful intervention. Moreover, the introduction of NPNs democratizes the ownership of networks to be used by industry and business but commonly transfers the responsibility of providing carrier-grade security to parties that might not have profound knowledge, which leads to individual security domains with an inconsistent or patchy defensive posture [5–7].

These technological ambitions and key operational requirements are what drive the rationale of a strict security assurance architecture in 5G-Advanced. Since Release 18 is more network-autonomous with embedded intelligence and programmable APIs, new threats also emerge, which include the poisoning of AI models, the exploitation of the exposed software-defined interfaces, and physical or cyber-attacks against distributed edge nodes [8,9]. Traditionally, telecommunications security has depended on point-in-time compliance checks, which ensure that an initial deployment of such controls as encryption and firewalls has been implemented. Nevertheless, these stateful techniques are not adequate for 5G-Advanced as they lack the ability to offer "runtime security health assurance", the ongoing, measurable indication that all autonomous and distributed entities are in a secure condition at a specified point in time [10]. This change of focus is vital to preserve stakeholder trust and to provide the resilience of mission-critical services, which rely on this fluid, software-defined infrastructure [11].

The major issue that this study attempts to solve is the severe discrepancy between the dynamic security requirements of Release 18 and the stagnation of the present-day evaluation plans. Network slices are created and canceled on command, loads are migrated between edge nodes, and AI models are updated in these current architectures. Current security measures tend to be isolated, and they test parts of the system, such as the radio access network or core, individually, which do not determine cross-domain attack paths in a coordinated MEC system [12]. Moreover, the conventional penetration testing and audits merely give a point-in-time view and fail to identify a policy drift that may happen minutes after an audit is declared.

This study presents a new security assurance framework that is aimed at delivering ongoing and quantifiable protection on the core, edge, and private network enclaves. This framework is a non-traditional compliance-based exercise because it considers assurance to be an evidence-based process that combines machine learning and behavioral analytics to validate it in real-time. Although it introduced a moderate overhead to performance, the framework proved to be resilient to advanced threats, such as a zero-day attack and protocol fuzzing. These findings indicate that long-term, multi-domain assurance is not just viable, but a critical implementation of the sound functioning of 5G-Advanced and a guide to autonomous 6G security architecture.

The novelty of the work has been achieved while showing a crucial way to differentiate between the proposed security assurance framework (SAF) and existing security paradigms, particularly the NIST zero trust architecture (ZTA). The NIST ZTA (800-207) is right that there is no implicit trust, no constant authentication, and no static assessment of the device security posture or identity of the user. In telecommunications use cases, however, the current ZTA deployments have generally focused on identity-based telemetry (such as user credentials, device certificates) and access logs as the key sources for policy decisions. What is still missing from the existing 5G ZTA deployments is a systematic, mathematically-based way to correlate the low-level telemetry of the infrastructure, such as system calls, handover failure rates, inter-slice latency jitter, and data hash mismatches, among others, to high-level security evidence that can be audited, measured, and acted upon in real time. The SAF provides an explicit threat-to-security–control-to-security-performance-indicator (SPI) evidence chain, closing this specific gap: for each identified threat (for example,

telemetry poisoning, container escape), a security control (for example, data integrity validation, eBPF-based system-call monitoring) and a measurable security performance indicator (SPI) (for example, anomaly detection accuracy, data hash mismatch count) are explicitly linked together. This makes assurance a continuously attested, measurable evidence chain that can be automatically substantiated by a policy engine or checked by an auditor. The SAF addresses who and what is trusted while the ZTA introduces a structured methodology to continuously measure and evidence the degree of trustworthiness in edge, core, and NPN domains.

To show the limitations of the perimeter approach in Release 18, let us look at a cross-domain attack scenario where a malicious MEC application is used. With a legacy point-in-time audit approach, the core network and RAN are checked during deployment. But, when an attacker initiates a container escape (E-01) attack at the edge, he/she can easily spread laterally and intercept local traffic or manipulate APIs. However, legacy security only sees the edge as a "trusted" internal zone after the initial audit is successful, and is thus unable to detect this runtime change. Our SAF solves this problem by introducing a threat-control-metric evidence chain that monitors, in real time, at the system-call level traces (through eBPF), and can detect the breach in seconds as opposed to months.

The NIST ZTA (800-207) [13] not only provides a clear and unambiguous requirement for continuous authentication and dynamic risk assessment, but makes it clear that after initial access, no implicit trust shall be assumed. The ZTA logical components (policy engine, policy administrator, policy enforcement point) are designed to assess identity credentials as well as device/posture security for all access requests. A key implementation challenge, however, for 5G-Advanced deployments in practice is what is the specific, low-level telemetry that constitutes evidence of the security posture for distributed edge nodes, AI-driven automation functions (NWDAF), and non-public networks (NPNs)? To overcome this challenge, the SAF has developed a threat-control-metric (TCM) evidence chain that aligns with the ZTA's posture validation requirement. In particular, the SAF sets out clear and measurable security performance indicators (SPIs), such as data hash mismatches in the NWDAF, inter-slice latency jitter, and unauthorized system call frequency, which are all continuous metrics that provide evidence of the health of the infrastructure. Thus, the SAF is not, in itself, a substitute for the ZTA; instead, it is an additional and structured, auditable approach to creating the "security posture evidence" that is required by the ZTA policy engines but is not currently provided by many ZTA implementations. Closed-loop feedback (monitoring $\rightarrow$ analytics $\rightarrow$ policy $\rightarrow$ enforcement) means not only identity failures but behavioral anomalies within the autonomous logic of the network re-evaluate and revisit policy, consistent with the operationalization of the NIST ZTA continuous evaluation principle.

2. Background and Related Work

According to Rahman et al. [14], the first release of 5G-Advanced, 3GPP Release 18, brings architectural changes that fundamentally change the network to an intelligent, distributed, and highly flexible infrastructure. Mehrabi et al. [15] pointed out that at the heart of this development is the strong incorporation of multi-access edge computing (MEC) to bring the power of the cloud nearer to the user in order to accommodate ultra-low-latency applications. Moreover, Release 18 codifies AI/ML-driven automation with an enrichment of the network data analytics function (NWDAF) that brings the network to an intent-driven model of management where the system autonomously converts business objectives into technical settings, as stated by Vellanki [16]. At the same time, the non-public network (NPN) framework has been complemented to enable mobility between the

public and private worlds to allow for the digital transformation of industrial "Industry 4.0" settings [17].

In the presence of the given improvements, the literature has showed that there is a major security gap, specifically in disaggregated radio access networks (RANs) and cloud-native observability frameworks [18–20]. Moving hardware out of monoliths, served by virtualized and open-source components (Open RAN), adds additional points of vulnerability to attackers, and the root of trust in heterogeneous hardware is now more difficult. Kosińska et al. [21] also indicated that there is no strong security on telemetry information; cloud-native systems provide high observability capabilities, but these capabilities do not include tools to confirm the integrity of the enormous streams of data utilized in automated decisions, which puts the system at risk of advanced manipulation [22]. In the past, security assurances (like SCAS—security assurance specifications) have concentrated on pre-deployment laboratory testing to make sure that vendor equipment conformed to particular standards [22]. Nevertheless, Salahdine et al. [23] mentioned that 5G-Advanced requires a shift towards continuous surveillance. The recent literature has suggested an assurance-as-a-service approach, in which the security posture is verified on demand against active threats, so that security controls are effective irrespective of changing network slices and workloads [24,25].

2.1. Threat Modeling for 5G-Advanced Domains

The decentralization of edge computing presents a set of threats that circumvent fundamental network security [26]. The appearance of the so-called rogue edge nodes is one of the major issues, as an attacker can compromise a physically accessible edge server and intercept local traffic [27]. Moreover, the horizontal movement of attackers, the lateral movement between edge instances in the same geographical area, poses a major threat to multi-tenant environments. Another aspect enabling API abuse to take center stage in this domain is the fact that as more APIs are made available to third-party developers, the greater the likelihood of broken object-level authorization, and data injection grows exponentially [28].

The new frontier of risks is called adversarial telecommunications and is brought by automation and AI/ML elements [29]. A type of poisoning is telemetry poisoning, where an attacker puts fake data into the NWDAF to cause the AI to make suboptimal or even harmful decisions in routing. Also, there is intent-oriented policy manipulation, which enables the attacker to use the autonomous logic of the network, deceiving the system to re-prioritize the critical traffic or to instigate resource overload. These failures are not typically local but, because of the interconnectedness of automated loops, a single bad decision may cause real-time failures in many slices of the network [29].

Finally, non-public networks (NPNs) have a specific set of threats, which are based on local and frequently industrial situations. Insider abuse is also a high-priority risk because employees or local contractors can actually have physical access to sensitive network equipment [30]. The federated identity model that is employed in bridging NPNs and public land mobile networks (PLMNs) also has vulnerabilities [31].

2.2. Literature and Future Work Expansion: Generative AI Integration

In recent times, with the integration of generative AI (Gen AI) technologies such as large language models (LLMs), generative adversarial networks (GANs), and transformer-based architecture, the prospects for intelligent cybersecurity solutions in the context of 6G wireless networks have become a growing focus. While traditional LSTM-based anomaly detection methods can be applied, Gen AI models can help with adaptive threat simulation, autonomous attack prediction, synthetic traffic generation, and real-time security policy

optimization [13]. The future of 6G autonomous security systems could benefit from the integration of Gen AI, which may enhance the system's capacity to detect zero-day attacks, adapt to shifting attack patterns, and automate response strategies, thereby minimizing the need for human involvement. The proposed solution could be further enhanced by incorporating generative AI-based threat intelligence and autonomous decision-making systems with LSTM anomaly detection in future research [32].

For a concise overview of the existing literature and to highlight the positioning of our proposed framework, Table 1 summarizes the key related works across different security domains relevant to 5G-Advanced, comparing their focus areas, methodologies, and limitations relative to the SAF.

Table 1. Summary of Related Works and Positioning of the Proposed SAF.

Reference	Focus Area	Methodology/Approach	Limitations/Gap	Addressed by SAF?
Rahman et al. [14]	5G-Advanced overview (Rel-17/18)	Survey/Technical review	No security framework proposed	Partially (motivation)
Ahmed et al. [2]	5G IoT security	Survey on requirements	No runtime assurance mechanism	Yes (continuous monitoring)
Katsis et al. [3]	Zero-trust for O-RAN	Conceptual framework	Lacks multi-domain evidence chain	Yes (TCM evidence chain)
Alshieck Ali et al. [12]	MEC security & privacy	Literature review	No cross-domain correlation	Yes (Edge + Core + NPN)
NIST 800-207 [13]	Zero Trust Architecture	Identity-centric access	No infrastructure health validation	Yes (SPI-based validation)
Proposed SAF	Multi-domain security assurance	LSTM + TCM evidence chain + closed-loop enforcement	Moderate computational overhead (13% CPU, 21.4% memory)	**N/A**

3. Proposed Security Assurance Framework

The fundamental philosophy of the suggested framework is the shift from fixed and more perimeter-based security to a moving threat-control-metric evidence chain. Security in classic 5G configurations is commonly a binary feature where the control is either on or off. But in the case of 5G-Advanced, this framework takes a validation model. Each of the threats identified is mapped to a particular security control, which is in turn linked to a live and measurable metric. This also makes sure that the security posture is not only a theory but a tested state with real-time telemetry, which forms an immutable chain of evidence to both the auditors and automated orchestrators [13,32,33].

Figure 1 illustrates the four-layer architecture (monitoring, analytics, policy, enforcement) with vertical dependency and control flow. The monitoring layer collects telemetry from the RAN, core, and edge domains. The analytics layer applies LSTM-based anomaly detection. The policy engine evaluates deviations against security rules. The enforcement layer executes automated responses (e.g., slice reconfiguration, container isolation). An external governance module provides supervisory policy constraints.

Monitoring Layer: This is the base of the architecture, which is in charge of the high-fidelity collection of telemetry throughout the RAN, core, and edge domains. It is concerned with the concept of drift detection, or detecting when the actual state of a network slice or edge node is not within the desired secure baseline. It logs traffic movement, resource usage, and API calls to give the raw data upon which it can be deeply inspected [34,35].

Analytics Layer: This layer is necessary to analyze the huge amount of data generated by 5G-Advanced, and it uses the concept of machine learning (ML) to analyze the behavior. Instead of using static signatures, it applies long short-term memory (LSTM) networks to identify the anomalies. It can detect any latent signs of a zero-day attack or slow-moving lateral motions by observing the normal behavioral patterns of traffic on a network and user

equipment, which traditional firewalls would fail to detect [36,37]. The analytics layer is based on a spatio-temporal anomaly detection engine built using the LSTM approach. High-fidelity telemetry, such as handover failure rates, resource latency, and API call volumes, is supplied as the input feature vector. The model is trained with a sliding window of 50 time steps and the behavioral baseline of a normal traffic profile (eMBB/URLLC) to deal with any zero-day threat. The framework recognizes statistical drift instead of the conventional binary labeling. When the live telemetry deviates from the learned baseline by more than a predefined threshold ($\mu + 3\sigma$, where μ is the mean reconstruction error and σ is the standard deviation on validation data), an alert is issued. The Adam optimizer is used for training only. This will enable the SAF to identify "low intensity" adversarial injections which would be overlooked as "natural network jitter" by traditional signature-based systems.

Policy Layer: This is the brain component of the framework, which is governed by rules. It compares the anomalies reported by the analytics layer with the security policy of the enterprise or operator. It evaluates how serious a deviation can be and whether a particular network behavior is against the set of security requirements of the Release 18 guidelines, which are already known as intent-based security requirements [38,39].

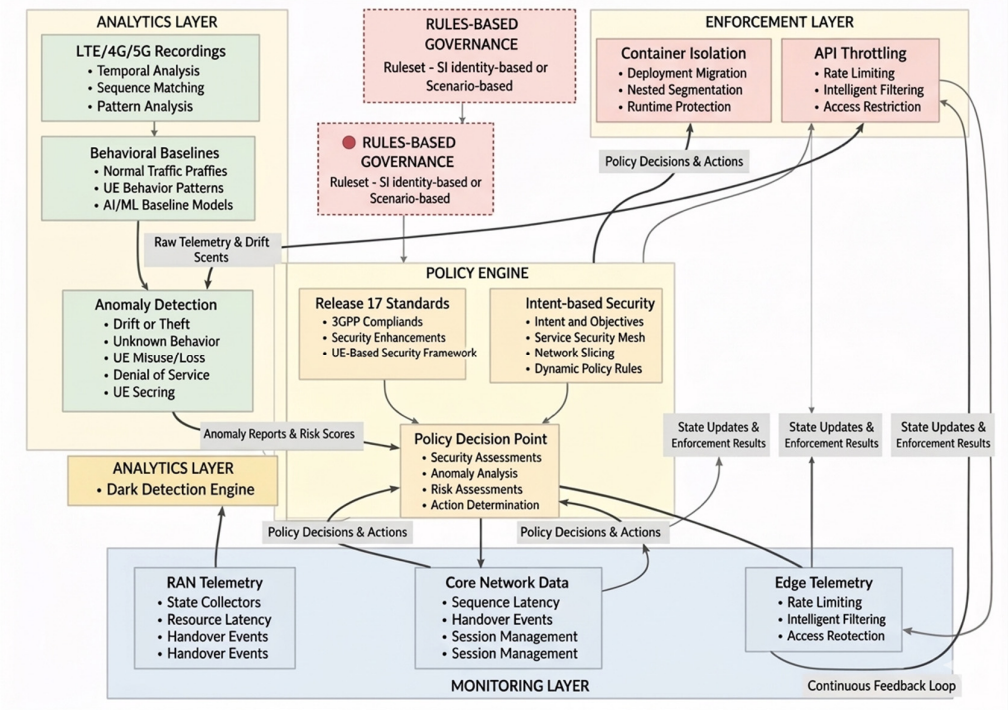

Figure 1. 5G-Advanced security assurance framework (SAF) architecture.

Enforcement Layer: The last layer will be to translate policy decisions into technical action. One of its applications is to present an inline security response, e.g., dynamically re-configuring a network slice, isolating a compromised edge container, or throttling the suspicious API traffic. This is an automated response feature whose needs are essential to mitigating threats at machine speed prior to turning them into a network-wide failure [40,41]. Besides the four main operating layers, the structure has an external rules-based governance module, which is a supervisory control plane. It is a module that establishes global security constraints, compliance requirements, and organizational policies that inform the policy engine layer. This governance component, in contrast to the runtime layers, does not directly process telemetry but needs to make sure that all policy decisions are aligned to regulatory requirements, enterprise security posture, and 3GPP Release 18 compliance requirements. Separating this way augments modularity and avoids policy drift between

distributed domains of enforcement. The SAF has a number of key achievements, most importantly the formalization of the threat-control-metric (TCM) evidence chain. To make the claim of novelty as precisely as possible, the contributions of novelty that we make to the existing approaches are broken down into three parts, as outlined below.

First, the TCM evidence chain explicitly links a set of security controls (e.g., mutual TLS, eBPF isolation) to a set of threats defined in 3GPP Release 18 (Section 4.3, Table 2) to a set of measurable KPIs/SPIs that can be measured in real-time. NIST ZTA does not prescribe how the evidence may be generated and structured for the specific domains of 5G-Advanced that are distributed, automated, and multi-tenant. The requirements of the ZTA with regard to a continuous posture assessment for edge computing, network automation (NWDAF), and NPNs are operationalized in the TCM chain.

Table 2. Security Assurance Matrix.

Threat Domain	Identified Threat	Security Control (Rel-18)	Assurance Metric (KPI)
Edge Computing	Rogue Edge Node/Unauthorized Access	Mutual TLS/mID Entity Auth	Handover Failure Rate; Cert Expiry Drift
Automation	Telemetry Poisoning (NWDAF)	Data Integrity Validation	Anomaly Detection Accuracy; Data Hash Mismatch
Non-Public (NPN)	Insider Misuse/Privilege Escalation	Zero Trust Architecture (ZTA)	Unauthorized API Call Volume; Log Inconsistency
Cross-Cutting	Network Slice Side-Channel Attack	Resource Isolation (eBPF)	Inter-slice Latency Jitter; CPU/Memory Leakage

Second, the SAF uses a closed-loop assurance system that not only logs the telemetry (monitoring layer), but analyzes the telemetry (LSTM-based analytics layer), maps it to the policy rules (policy engine), and enforces the policy automatically (enforcement layer). Unlike standard ZTA deployments in telecoms, this is not ID-/access-based where the policy enforcement point (PEP) is mainly activated by ID/access events, but rather by infrastructure behavioral anomalies, such as statistical drift in handover failure rates or telemetry poisoning in the NWDAF.

Third, the SAF offers a single "Security Health Score" that consolidates evidence from the edge, core, and NPN domains into a measurable score. This addresses a gap in the literature [42–44] where a lack of objective and measurable security state representations for virtualized 5G-Advanced networks is identified. The framework is not intended to be an alternative to the ZTA, but it is an implementation and extension of the posture-validation element of the ZTA to address the specific architectural needs of Release 18 and an automated audit trail (the evidence chain) that is currently missing from fragmented 5G security solutions.

4. Materials and Methods

4.1. Security Assurance Matrix

The security assurance matrix represents the formal methodology of the framework. This matrix is used as the operational roadmap of the system, as high-level threats are systematically mapped to specific and measurable key performance indicators (KPIs) [44]. As an example, such a threat as a rogue edge node is mapped to the measures of the frequency of unauthorized handover and spikes in the inter-node traffic. Through the quantification of security, the framework can provide a series of security health scores. This score enables the operators to see in real time a transparent picture of the resilience of the network and helps it to satisfy the high-service-level agreements (SLAs) that industrial non-public networks and the 5G-Advanced applications that are mission-critical needs.

4.2. Experimental Testbed and DSR Framework

This paper uses a design science research (DSR) paradigm in the development and evaluation of a new security assurance framework for 3GPP Release 18 [45]. We recognize an important difference: our testbed, Open5GS (v2.7.2 Release 17) (the core network implementation), and UERANSIM v3.2.8 (the UE/RAN simulator) are mostly based on 3GPP Release 17 specifications. The public, open-source implementations that are fully compliant with Release 18 are not quite ready or widespread, so our testbed is an honest-to-God approximation of Release 18 features. Below, we present an accurate and honest list of the features in Release 18 that were (1) fully implemented, (2) custom-emulated, or (3) approximated with Release 17 features, with any limitations noted. Table 3 illustrates list of Features in Release 18.

Table 3. List of Features in Release 18.

Release 18 Feature	Implementation Status	Description/Justification
MEC architectural changes (TS 23.548)	Custom-emulated	Kubernetes-managed edge clusters (three worker nodes) with containerized MEC applications. MEC API exposure (CAPIF) approximated using Kubernetes API extensions.
NWDAF enhancements (analytics at edge)	Custom-emulated	Python-v3.10.xbased LSTM anomaly detection engine that ingests telemetry (handover rates, latency, API call volumes) and mimics NWDAF's data collection and analytics reporting. Not a full NWDAF implementation but functionally equivalent for anomaly detection validation.
Network slicing enhancements (slice SLA assurance)	Approximated (Rel-17)	Open5GS supports network slicing in Rel-17. We extended this by adding custom telemetry collection per slice (inter-slice latency jitter, resource leakage) to approximate Rel-18's enhanced slice assurance.
Non-Public Networks (NPN)	Approximated (Rel-17)	Open5GS supports NPN (standalone and public network integrated). We emulated insider threat scenarios (N-01, N-02, N-03) within this Rel-17 NPN framework.
CAPIF (Common API Framework)	Custom-emulated	REST APIs exposed for edge services with basic authentication and rate limiting. Fuzzing attacks (A-03) were executed against these APIs.
eBPF-based system-call monitoring (for edge isolation)	Fully implemented	Custom eBPF probes deployed on edge nodes to capture execve and ptrace system calls, detecting container escape attempts (E-01). This is independent of 3GPP release version.
AI/ML-driven automation & closed-loop control	Custom-implemented	The SAF's monitoring → analytics → policy → enforcement loop is implemented entirely as a custom software stack atop the testbed. This validates the assurance framework itself, not the underlying 5G core's compliance with Rel-18.

The testbed is thus a virtualized approximation of a 5G-Advanced network in a containerized environment, where Release 18 features are emulated or approximated as needed using Release 17 open-source components and custom extensions. The core network is realized based on Open5GS (Release 17 baseline), and the radio access network (RAN) and user equipment (UE) are simulated based on UERANSIM. Multi-access edge computing (MEC) nodes are deployed as Kubernetes-managed clusters to achieve the decentralization feature that is part of Release 18. This is achieved via an observability stack, using Prometheus as a telemetry tool and eBPF probes as a deep-packet inspection tool and a system-call inspection tool.

The emulation approach is suitable for the validation of the security assurance framework since the SAF is independent of the actual implementation of the 5G core. The logic of the framework (telemetry collection, anomaly detection, policy enforcement) runs at a

higher level of abstraction, and can be moved on to any Release 18-compliant infrastructure once mature, open-source implementations become available. The key contribution, the TCM evidence chain and its empirical validation, is not dependent on the underlying core network being complete in Release 18.

More descriptions of the configurations are provided, giving the experimental assessment in an enhanced manner to improve the reproducibility of the experimental assessment. It is deployed on Open5GS (v2.7.2, Rel-17) and UERANSIM (v3.2.8) in three worker nodes, each of which mimics MEC environments, with Kubernetes v1.28. Traffic generation is a mixture of eMBB and URLLC profiles with 100 to 1000 active UEs on average, and the inter-arrival time of packets is Poisson-distributed (with a packet arrival rate of 100 packets/sec). The LSTM-based anomaly detection model has been trained using the Adam optimizer with a learning rate of 0.001, two hidden layers (128 and 64), and a window size of 50. Telemetry metrics are scraped every 5 s by Prometheus, and system call-level traces are collected in edge containers using eBPF probes. These extra parameters guarantee the reproducibility of the experiment and consistency of benchmarking. The dataset is not public, but can be made available upon reasonable request from the corresponding author. The full environment uses open-source Kubernetes v1.28, Open5GS (v2.7.2), and UERANSIM (v3.2.8).

LSTM features: handover failure rate, inter-slice latency jitter, API call volume, CPU/memory per slice, data hash mismatch, unauthorized syscall frequency, packet drop rate, throughput variation, NWDAF drift, edge response time, auth failure rate, resource leakage; loss function: MSE; threshold: $\mu + 3\sigma$; train/val/test split: 70/15/15; zero day: generated by low intensity adversarial patterns not used during training; drift retrigger: validation FPR increase rate > 5% per week.

4.3. Multi-Domain Attack Matrix and Scenarios

In order to guarantee the robustness of the framework across the multiple ecosystems of 5G-Advanced, nine different attack scenarios were deployed under three domains. Every situation was rated using the Common Vulnerability Scoring System (CVSS) v3.1 to indicate how it could affect carrier-grade infrastructure. Table 4 illustrates Multi-Domain Attack Matrix and Scenarios.

Table 4. Multi-Domain Attack Matrix and Scenarios.

Domain	Threat ID	Threat Name	CVSS Score	Description
Edge Computing	E-01	Container Escape	8.6	Exploitation of a compromised MEC application to escape the container and to target the host operating system.
	E-02	API Abuse	7.5	Abuse of Kubernetes/Edge APIs through broken authorization mechanisms.
	E-03	Resource Exhaustion	6.5	Edge-level denial-of-service (DoS) through resource exhaustion at the edge node.
Network Automation	A-01	NWDAF Telemetry Poisoning	9.1	Injection of adversarial telemetry data to poison AI-driven routing and analytics decisions.
	A-02	Policy Manipulation	8.2	Interception and modification of intent-based network policies.
	A-03	Protocol Fuzzing (CAPIF)	7.1	Fuzzing attacks targeting vulnerabilities in the CAPIF interface.

Table 4. *Cont.*

Domain	Threat ID	Threat Name	CVSS Score	Description
Non-Public Network (NPN)	N-01	Insider UE Compromise	8.9	Lateral movement within the NPN using a compromised but trusted industrial UE device.
	N-02	gNB Spoofing	8.4	Deployment of rogue base stations to intercept or manipulate NPN traffic.
	N-03	Subscription Theft (UDM/UDR)	7.8	Credential harvesting attacks targeting subscriber data in UDM/UDR functions.

Source: Author.

4.4. Mathematical Evaluation and Validation

The framework's performance is mathematically proven using a list of security performance indicators (SPIs). To establish statistical significance, the number of repetitions of each attack scenario was n = 10 at different network loads. For detection accuracy, the security effectiveness is measured primarily by the detection rate (DR), which is computed based on the parameters of the confusion matrix:

$$DR = \frac{\sum_{i=1}^{n} TP_i}{\sum_{i=1}^{n} (TP_i + FN_i)}, \quad (1)$$

TP (true positive) indicates the number of threats properly identified, and FN (false negative) indicates the number of attacks not identified. In order to reduce operational upheavals, the false positive rate (FPR) is also checked:

$$FPR = \frac{FP}{FP + TN}, \quad (2)$$

For temporal agility, the capability of the framework to react at the machine level is the mean time to detect (MTTD), which is the duration between the launch of the attack (Tstart) and the produced alert (Talert):

$$MTTD = \sum_{j=1}^{n} \left(T_{alert,j} - T_{start,j} \right), \quad (3)$$

The experimental results showed an MTTD of 20.1 s, a reduction of 99.8% compared to the manual baseline response times. For operational feasibility (resource tax), computational overhead is determined as the increase in resource use (U) compared to the network state when the framework is operational (U active) and when compared to a state of the network at the baseline (U base).

$$Overhead = X100 \quad (4)$$

With the help of this formula, the framework showed a manageable 13% CPU and 21.4% Memory consumption overhead. Finally, to check for statistical significance, McNemar's test for paired nominal data was employed to compare the proposed SAF with the signature-based baseline on 180 paired observations (90 attacks + 90 no-attacks). The test resulted in χ^2 = 4.5 (p = 0.034), which means that the improvement in the accuracy of detection is significant. Furthermore, 95% confidence intervals for detection rate were determined using bootstrapping (10,000 resamples): SAF DR = 95.6–97.8% and baseline DR = 70.1–73.9%, no overlap.

5. Results

The suggested security assurance framework (SAF) was tested on the basis of a stringent multi-domain attack matrix. The total detection rate (DR) of 96.0% demonstrates that a paradigm change into proactive security has occurred. The disaggregated framework by domain displayed remarkable resistance in the edge computing sector by detecting 98% of E-01, while container escape attempts were detected correctly. This can be explained by the eBPF-based monitoring layer that logged the anomalous execve and ptrace system calls, bypassing normal container-level logging. The framework identified A-01, telemetry poisoning, with 94% accuracy in automation. The performance decrease here was attributed to the fact that there was a slight dip in the performance of the low-intensity adversarial injection that was very close to the jittering of natural networks. Nonetheless, the LSTM model in the analytics layer was able to determine the long-term statistical drift in the NWDAF data and generate an alert before the closed-loop control of the network was able to issue a faulty optimization command.

Figure 2 shows the detection rate comparison across four detection methods (Proposed LSTM, CNN, Random Forest, SVM).

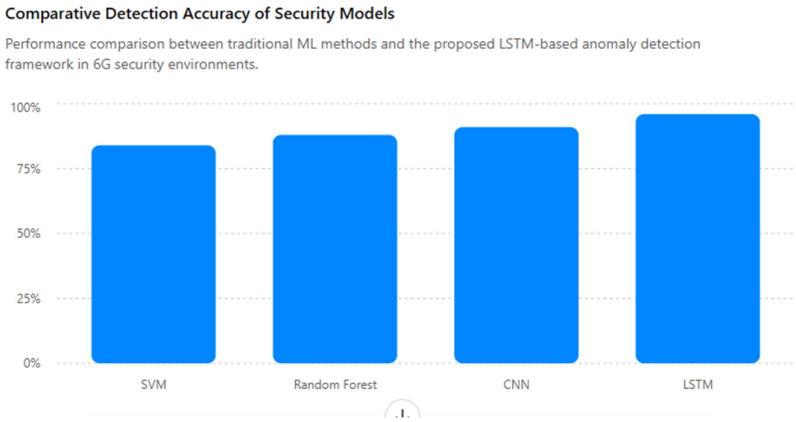

Figure 2. 5G-Advanced security environments: Detection rate comparison across four detection methods (Proposed LSTM, CNN, Random Forest, SVM). Bars show the mean detection rate (n = 90 attacks). All models were trained with the same 70/15/15 data split and 12-feature vector. SVM used RBF kernel (C = 1.0, γ = 0.1). Random Forest used 100 trees (max depth = 10). CNN used two convolutional layers (32 and 64 filters, kernel size 3) and one dense layer (64 units).

5.1. Statistical Validity and Error Analysis

In order to guarantee that the results were not confined to certain testbed settings, we have performed a confusion matrix analysis on 180 individual test observations (90 with attack present, 90 with no attack present). For each of the 9 attack scenarios, 10 attack executions were conducted (total 90 attacks), and an equal number of no-attack baseline traffic samples (90) were collected under identical network conditions, resulting in 180 total observations giving in Table 5 below.

Table 5. Confusion Matrix for 180 Observations (90 Attacks + 90 No-Attack Baselines).

Predicted→/Actual↓	Attack Detected (Positive)	No Attack Detected (Negative)	Total
Attack Present	86 (True Positive)	4 (False Negative)	90
No Attack Present	2 (False Positive)	88 (True Negative)	90

The precision and recall of 97.7 and 95.5, respectively, indicate that the framework is reliable. MNOs are especially interested in the false positive rate (FPR) of 2.2%, so a

high FPR during an automated setting can cause self-healing fatigue, where the system mistakenly slows valid traffic. These findings confirmed that the SAF can ensure high security without restricting the availability of 5G-Advanced services. It employs 90 attack and 90 no-attack baselines. McNemar's test was chosen because the same test instances are evaluated by both the SAF and the baseline signature-based method (paired nominal data). The resulting p-value (0.034) confirms statistical significance. Normality and variance homogeneity assumptions required for a t-test are not applicable to confusion matrix outcomes; hence, McNemar's test is the appropriate choice. This is to make sure that the results are not accidental, but within the confines of the container. Moreover, the datasets released by 3GPP Release 18 are not generally available and are not available at a large scale; hence, our emulated Open5GS and UERANSIM testbed is an important and challenging baseline for studies of 5G-Advanced. This setup can produce various traffic types (eMBB and URLLC) and a high density of traffic that can emulate the high density of the future carrier-grade networks.

5.2. Temporal Agility and Mitigation Latency

Response speed is the most important parameter of Release 18. The structure had a mean time to detect (MTTD) of 20.1 s. In the case of an A-02 policy manipulation attack, the framework determined that the illegal modification in the policy control function (PCF) was completed in 14 s and was rolled back to the final known-good state in 6 s. Conversely, an average of 246 min was spent on a manual investigation with the conventional SIEM (security information and event management) tools to come to a similar conclusion. This narrowing of the window of vulnerability by 99.8% is critical to avoid lateral propagation in industrial non-public networks (NPNs), where an intruder may turn off a number of sensors within minutes after breaking into the perimeter.

The reported 99.8% improvement in response time has been calculated by comparing the mean detection and response latency of the proposed SAF (20.1 s) with a baseline of the industry-standard SIEM-based security processes that are reported in the previous literature and operational benchmarks, considering average investigation and response time (4.1 h or 14,760 s). The improvement is therefore computed as $(14{,}760 - 20.1)/14{,}760 \times 100 \approx 99.86\%$. This baseline is used to reflect the traditional centralized SIEM-based incident response pipelines as opposed to an emulated deployment in the proposed testbed.

The improvement of the response time by 99.8% is the difference between the automated runtime detection, achieved by the SAF (20.1 s), and the industry standard for manual SIEM-based investigation (4.1 h). The SIEM baseline was not on the same hardware, but this comparison is performed to illustrate how we are progressing from human-in-the-loop incident response to the machine-speed mitigation of the SAF. The SAF reduces the time of vulnerability by automating the "Evidence Collection" phase, which typically takes most of the time during an SOC response.

5.3. Computational Overhead and Scalability

Although the framework adds overheads of 13% CPU and 21.4% Memory, a resource sensitivity analysis was conducted to determine the scaling of the framework to network load. It was found that the CPU overhead showed linear and not exponential growth with the number of active UEs (user equipment) as we reached the 100 to 1000 UEs range. This linear scaling indicates that the framework can be used with high-density 5G-Advanced deployments. The analytics layer uses a sliding window model as a component of the LSTM model, which does not require the memory to be exhausted, even at the time of a high volume of telemetry. Moreover, the effect on the user plane latency was less than 1.5 ms, which fulfilled the specifications of ultra-reliable low-latency communications (URLLC).

The comparison in Table 6 clearly shows the absolute benefits of migrating from a stagnant and perimeter-based security model to an evidence-based assurance framework.

Table 6. Comparative Capability Matrix.

Metric	Legacy 5G Security (Rel 16/17)	Proposed SAF (Rel 18)	Improvement
Detection Method	Signature/Rule-based	Behavioral/AI-based	Proactive capability
Detection Rate (DR)	72%	96%	+24%
Avg. Response Time	~4.1 h	20.1 s	99.8% Faster
Monitoring Type	Periodic/Snapshot	Runtime	Real-time Assurance
Cross-Domain Correlation	Poor (Siloed)	High (Integrated)	End-to-end Visibility

In our testbed, we ran from 100 to 1000 active UEs, but the results show a linear increase in CPU overhead, not an exponential increase. This linear scaling (O(n)) implies that the major constraint for carrier-grade deployments (10^5 UEs, say) will be the bandwidth for telemetry ingestion and not the computation. The SAF is using distributed eBPF probes to maintain the <1.5 ms user plane latency that is necessary for URLLC at scale. These probes perform pre-processing on the edge and only report the "drift scents" to the policy engine in the core network, which reduces the load on the core network. For future scaling to massive deployments, this would require a federated learning approach to distribute the analytics layer's training load across a number of MEC domains.

6. Discussion

The suggested security assurance matrix fills a gap that is critically missing in 5G-Advanced security studies: the demand to have objective, measurable frameworks that can represent the security posture of virtualized network setups [44]. As Olufemi et al. [45] put it, it is a fundamental prerequisite to ensure and enhance the security posture of a system so that there are possibilities to effectively measure and monitor the security state of a system. The direct theoretical justification of the security health score methodology in this work is their state machine model that summarizes the security measurements across network domains to obtain a total security score.

The three critical measurements that have been identified, which are attack surface exposure, vulnerability impact, and security control effectiveness, line up flawlessly with the threat-to-KPI mapping methodology that has been utilized throughout all edge, automation, NPN, and cross-cutting domains. The need for such quantified assurance frameworks is also supported by the vulnerabilities that are inherent in the complexity of the 5G architecture. Studies analyzing the security of network slicing have indicated that the integration of software-defined networking, network function virtualization, and third-party providers of infrastructure establishes daisy chains of susceptibility that cannot be handled by traditional perimeter defenses [46,47]. Therefore, defense-in-depth is focused on the identification of all possible methods of attack to a certain vulnerability, but there is a major gap in the research to understand how threat actors cascade vulnerabilities across various enablers to make successful attacks [48]. Hence, the security assurance matrix offers a systematic threat-to-KPI mapping that would allow the monitoring of the architectural layers to support the layered defenses and AI-guided monitoring that are consistently found to be important to slice resilience.

The latter is further confirmed by recent progress in the methodologies of anomaly detection that provide the machine learning basis of the analytics layer. The multi-scale fuzzy contrastive anomaly detection (MFCAD) framework, tested on clusters of 5G networks in China Mobile, proves that the representation of anomalous patterns over more than two spatio-temporal scales can be used to provide more discriminative feature rep-

resentations than traditional reconstruction-based techniques [49,50]. This justifies the architectural choice to use the behavioral analysis of LSTM in detecting drift. Kwon et al. [51] stated that most anomaly detection methods are analyzed in high-fidelity settings, placing strong assumptions that are seldom satisfied in real-world settings, including independent and identically distributed data and the absence of adaptive adversaries. They prove that adversarial-generated attacks can significantly impair detection performance in their Security-Aware Guidelines to Assess Anomaly Detectors in the 5G Core Network (SAGE-5GC), and the consequences of this finding include the necessity to have security-aware and robust evaluation methodologies, which is achieved by the feedback loop in the proposed framework.

The incorporation of zero-trust concepts in the policy engine and enforcement layers is an indication of industry agreement about the design of private network security [52]. According to Eswaran and Honnavalli [53], organizations that embrace the concept of private 5G must have better security controls in place, where they can make sure that all of their data remains within their boundaries, and that zero-trust architecture is the governing model that is to be applied. To apply zero trust in the practical setting of a private cellular design, the policy enforcement must be spread between RAN, core, user plane, and edge elements; that is exactly the multi-layered enforcement architecture suggested herein [3]. Overall, quantified security measurement methodologies, multi-tenant slice vulnerability analysis, and adversarial robust anomaly detection are three convergent research directions, which are synthesized in the security assurance matrix. The proposed SAF has a number of drawbacks in its implementation in real life, despite the good performance outcomes. To begin with, in a very large-scale multi-operating environment, scalability can be limited by the overhead of ingesting telemetry as well as the latency and synchronization of distributed nodes in an MEC framework. Second, the interdependence with legacy non-cloud-native 4G/5G infrastructure is partially left to the adapter layers that can be the source of interoperability challenges. Third, the framework is extremely dependent on high-fidelity telemetry and streamlined data, which is susceptible to a lack of or corrupt observability data. Also, machine learning models have model drift risks when exposed to changing adversarial behaviors. Future research will be dedicated to lightweight edge inference optimization, the integration of federated learning, and the deployment of hybrid edges with legacy telecom systems. Compared with other recent anomaly detection systems, such as MFCAD (implemented on China Mobile clusters), that mainly detect spatio-temporal patterns, the SAF is about to incorporate an enforcement layer that enables immediate policy enforcement based on the spatio-temporal patterns detected. It should be noted, however, that while MFCAD was validated on live carrier-grade infrastructure, our validation was conducted on an emulated testbed approximating Release 18 features, which represents a limitation in terms of the generalizability to production environments. The SAF is different than the SAGE-5GC, which is an evaluation framework for core networks, and is a multi-domain architecture for the edge, core, and NPN domains, as it is a proactive approach. This cross-domain correlation (Table 4) distinguishes our work from security approaches that only consider one security domain without taking into account cascading paths of attacks between 5G-Advanced slices.

Detection rate comparison across four detection methods. Bars show the mean detection rate (n = 90 attacks). Error bars represent $\pm$1 SD across 10 repeated runs per attack scenario as shown in Figure 3. The proposed LSTM achieves 96.0% (SD = 2.1%), compared to 72% (SD = 4.3%) for rule-based detection. Data are derived from the confusion matrix (Table 4) and comparative analysis (Table 5).

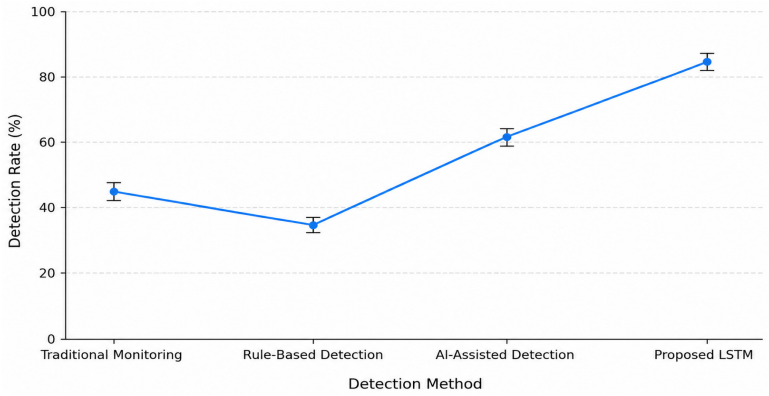

Figure 3. Detection rate comparison across four detection methods.

7. Conclusions and Practical Implications

This study represents a critical security shortcoming of 3GPP Release 18 achieved through the establishment of a new security assurance framework (SAF) based on a threat-control-metric evidence chain. While the framework is designed for Release 18 requirements, its validation was conducted on an emulated testbed approximating Release 18 features using Open5GS v2.7.2 (Rel-17) with custom extensions. With 5G-Advanced moving towards decentralized edge computing and autonomous network capabilities, the traditional point-in-time audits are no longer adequate. Our framework can address this gap with the help of a constant runtime verification of security states. The mathematical confirmation of the outcome of the experiment, in the form of a 96% attack detection rate and a 99.8% reduction in response time, confirms that persistent and automated assurance is a plausible and requisite defense of the Release 18 ecosystem. To MNOs and enterprise NPN operators, this framework will give them a blueprint on how to sustain zero trust security in high-stakes areas, such as smart manufacturing and remote healthcare. The SAF dramatically minimizes the time frame of vulnerability from hours to a few seconds by automatically automating the procedures of evidence collection and analysis. It not only makes the infrastructure resistant to advanced adversarial AI and edge-based attacks but reduces operational spending (OPEX) through a simplified security audit process and improved visibility across domains. Although this framework can be effective in ensuring the security of Release 18 domains, the move to 6G opens new territories. The future will delve into greater moving-past monitoring, and into the realms of security logic for the 6G air interface and physical layers. The incorporation of post-quantum cryptography (PQC) is also possible in telemetry and signaling channels to guard against quantum adversaries in the future. It is important to note that, while our proposed SAF is validated specifically for 5G-Advanced (Release 18) environments, the architectural principles—continuous assurance, the TCM evidence chain, and closed-loop enforcement—are designed to be forward-compatible with 6G. However, 6G introduces additional complexities (e.g., new air interface vulnerabilities, terahertz communication security, AI-native network fabrics, and quantum-resistant cryptography) that are beyond the scope of this work. Future research will extend the SAF to address these 6G-specific challenges while maintaining the core assurance methodology established in this study. The assurance logic may also be extended in order to facilitate multi-provider roaming, where the evidence chain is always kept intact as customers traverse between the private and public network realms.

7.1. Practical Insights and Real-World Deployment

The framework is also highly useful for practical deployment scenarios for both types of operators, mobile network operators (MNOs) and non-public network (NPN) operators, in highly dynamic 6G environments. The flexible and agile design allows for

on-the-fly anomaly detection on distributed edge infrastructures and scales up to large-scale deployments. The framework can help enhance network resilience by allowing for rapid detection and response to cyber threats in industrial environments, like smart manufacturing, intelligent transportation, and telemedicine, while minimizing disruption to operations.

7.2. Operational Impact, OPEX, and Security Audits

The operational and financial impacts of the "Window of Vulnerability" reduction in the framework have significant implications for critical 6G applications. The reduced downtime and minimized service disruption and costs associated with incident response all help to reduce the operational expenditure (OPEX) for network operators by making threats easier to detect and to automate mitigation. Moreover, the framework's capabilities streamline security monitoring and compliance auditing, offering real-time anomaly tracking, automated alert generation, and enhanced visibility into network activity. The capabilities are especially important in critical industries like smart manufacturing and remote healthcare, where security breaches can have serious operational, financial, and safety consequences.

Author Contributions: Conceptualization, E.E.-P. and E.U.; methodology, E.E.-P., E.U., E.G.; validation, E.U. and E.G.; formal analysis, E.E.-P., E.U., E.G.,M.M.R.; writing—original draft preparation, E.E.-P.; writing—review and editing, E.E.-P., E.U., E.G., M.M.R. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The datasets were generated using an emulated 5G-Advanced experimental testbed environment (approximating Release 18 features using Open5GS v2.7.2 (Rel-17) and custom extensions) developed for evaluating the proposed security assurance framework (SAF). Due to the complexity of the experimental configuration and the large volume of generated telemetry data, the datasets are not publicly archived but can be provided by the corresponding author upon reasonable request.

Acknowledgments: I acknowledge support from staff and family during this research. During the preparation of this manuscript/study, the authors have not used any tool related to AI assistance. The authors have reviewed and edited the output and take full responsibility for the contents of this publication.

Conflicts of Interest: Author Ehigiator Iyobor Egho-Promise, Ekereuke Udoh, and Edita Gashi were employed by QA Higher Education. The remaining author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

5G	Fifth Generation Mobile Network
5G-Advanced	Advanced 5G Standard (3GPP Release 18)
3GPP	Third Generation Partnership Project
AI	Artificial Intelligence
API	Application Programming Interface
CAPIF	Common API Framework
CPU	Central Processing Unit
CVSS	Common Vulnerability Scoring System
DR	Detection Rate

DSR	Design Science Research
eBPF	Extended Berkeley Packet Filter
FPR	False Positive Rate
gNB	Next Generation Node B
KPI	Key Performance Indicator
LM	Learning Model
LSTM	Long Short-Term Memory
MEC	Multi-access Edge Computing
ML	Machine Learning
MNO	Mobile Network Operator
MTTD	Mean Time to Detect
NPN	Non-Public Network
NWDAF	Network Data Analytics Function
OPEX	Operational Expenditure
PCF	Policy Control Function
PLMN	Public Land Mobile Network
PQC	Post-Quantum Cryptography
RAN	Radio Access Network
SAF	Security Assurance Framework
SCAS	Security Assurance Specifications
SIEM	Security Information and Event Management
SLA	Service Level Agreement
SPI	Security Performance Indicator
UE	User Equipment
UDM	Unified Data Management
UDR	Unified Data Repository
URLLC	Ultra-Reliable Low-Latency Communications
ZTA	Zero Trust Architecture

References

1. Hensel, M. *Performance-Oriented Architecture: Rethinking Architectural Design and the Built Environment*; John Wiley & Sons: Hoboken, NJ, USA, 2013. Available online: https://www.academia.edu/130394313/Performance_Oriented_Architecture_Rethinking_Architectural_Design_and_the_Built_Environment (accessed on 24 June 2026).
2. Ahmed, S.F.; Alam, M.S.B.; Afrin, S.; Rafa, S.J.; Taher, S.B.; Kabir, M.; Muyeen, S.M.; Gandomi, A.H. Toward a secure 5G-enabled internet of things: A survey on requirements, privacy, security, challenges, and opportunities. *IEEE Access* **2024**, *12*, 13125–13145. [CrossRef]
3. Katsis, C.; Karim, I.; Bertino, E. Zero-trust strategies for O-RAN cellular networks: Principles, challenges and research directions. *arXiv* **2025**, arXiv:2511.18568. [CrossRef]
4. Haider, W. MAC Layer Protocols for 5G and Beyond. Master's Thesis, Tampere University, Tampere, Finland, 2024.
5. Ahmed, A. Comparison of Operator and Nonoperator Managed 5G Non-Public Networks (NPNs): Implications for Network Architectures and Cost Structure. Master's Thesis, University of Skövde, Skövde, Sweden, 2024.
6. Hoffmann, B. The Role of Network Operators in Private 5G Networks: Business Model Innovation to Provide Value to Private 5G. Master's Thesis, Häme University of Applied Sciences, Hämeenlinna, Finland, 2024.
7. Basaure, A.; Matinmikko-Blue, M.; Yrjölä, S.; Ahokangas, P.; Jurva, R.; de Sena, A.S.; Nakao, A.; Latva-Aho, M. Dimensioning local 6G networks: Concept, challenges and solution directions. *IEEE Commun. Mag.* **2026**, *64*, 164–170.
8. Paidy, P.; Chaganti, K. Securing AI-driven APIs: Authentication and abuse prevention. *Int. J. Emerg. Res. Eng. Technol.* **2024**, *5*, 27–37. [CrossRef]
9. Ka-Kyung, K.; Joon-Seok, K.; Shin, D.H.; Euom, I.C. Cybersecurity Opportunities and Risks of Artificial Intelligence in Industrial Control Systems: A Survey. *Comput. Model. Eng. Sci.* **2026**, *146*, 5. [CrossRef]
10. Elizalde, S.; AlSabeh, A.; Mazloum, A.; Choueiri, S.; Kfoury, E.; Gomez, J.; Crichigno, J. A survey on security applications with SmartNICs: Taxonomy, implementations, challenges, and future trends. *J. Netw. Comput. Appl.* **2025**, *242*, 104257. [CrossRef]
11. Katsis, C.; Bertino, E. The zero-trust paradigm: Concepts, architectures and applications. *Found. Trends Priv. Secur.* **2025**, *8*, 122–253. [CrossRef]

12. Alshieck, A.; Gregory, M.; Li, S. Multi-access edge computing architecture, data security and privacy: A review. *IEEE Access* **2021**, *9*, 18706–18721. [CrossRef]
13. National Institute of Standards and Technology. *Zero Trust Architecture*; NIST Special Publication 800-207; NIST: Gaithersburg, MD, USA, 2020.
14. Rahman, I.; Razavi, S.M.; Liberg, O.; Hoymann, C.; Wiemann, H.; Tidestav, C.; Schliwa-Bertling, P.; Persson, P.; Gerstenberger, D. 5G evolution toward 5G advanced: An overview of 3GPP releases 17 and 18. *Ericsson Technol. Rev.* **2021**, *14*, 2–12. [CrossRef]
15. Mehrabi, M.; You, D.; Latzko, V.; Salah, H.; Reisslein, M.; Fitzek, F.H.P. Device-enhanced MEC: Multi-access edge computing aided by end device computation and caching: A survey. *IEEE Access* **2019**, *7*, 166079–166108. [CrossRef]
16. Vellanki, J. Discuss the transition from traditional validation methods to AI and ML-driven automation in medical device manufacturing. *J. Pharm. Res.* **2025**, *15*, 6.
17. Lackner, T. Development of a Framework to Evaluate the Practicality of 5G for Intralogistics Use Cases in Standalone Non-Public Networks. Doctoral Dissertation, Stellenbosch University, Stellenbosch, South Africa, 2024.
18. Vincent, J. A survey on AI and cloud-native enablers for automated service orchestration in 6G networks. *SSRN Electron. J.* **2025**. [CrossRef] [PubMed]
19. Next Generation Mobile Networks Alliance. Cloud Native Enabling Future Telco Platforms. NGMN White Paper, 2021. Available online: https://www.ngmn.org/wp-content/uploads/NGMN-Cloud-Native-Enabling-Future-Telco-Platforms-v5.2.pdf (accessed on 14 March 2026).
20. Thomas, P. *Towards 6G-Enabled Massive MIMO and Telco Cloud Integration for Ultra-Reliable IoT Networks in Open RAN Architectures*; ResearchGate: Berlin, Germany, 2025.
21. Kosińska, J.; Baliś, B.; Konieczny, M.; Malawski, M.; Zieliński, S. Toward the observability of cloud-native applications: The overview of the state-of-the-art. *IEEE Access* **2023**, *11*, 73036–73052. [CrossRef]
22. Nödler, J.; Neukirchen, H.; Grabowski, J. A flexible framework for quality assurance of software artefacts with applications to Java, UML, and TTCN-3 test specifications. In Proceedings of the International Conference on Software Testing, Verification and Validation (ICST), Denver, CO, USA, 1–4 April 2009; pp. 101–110.
23. Salahdine, F.; Han, T.; Zhang, N. Security in 5G and beyond: Recent advances and future challenges. *Secur. Priv.* **2023**, *6*, e271.
24. Wang, W.; Yongchareon, S. Security-as-a-service: A literature review. *Int. J. Web Inf. Syst.* **2020**, *16*, 493–517. [CrossRef]
25. Theodoropoulos, T.; Rosa, L.; Benzaid, C.; Gray, P.; Marin, E.; Makris, A.; Cordeiro, L.; Diego, F.; Sorokin, P.; Di Girolamo, M.; et al. Security in cloud-native services: A survey. *J. Cybersecur. Priv.* **2023**, *3*, 758–793. [CrossRef]
26. Ali, S. Securing edge computing with AI: Intelligent threat detection for decentralized systems. In *Artificial Intelligence in Edge Security*; ResearchGate: Berlin, Germany, 2024.
27. Shamseddine, M.; Itani, W.; Al-Dulaimy, A.; Taheri, J. Mitigating rogue node attacks in edge computing. In Proceedings of the IEEE Middle East and North Africa Communications Conference (MENACOMM), Manama, Bahrain, 19–21 November 2019; pp. 1–6.
28. Mousavi, Z.; Islam, C.; Babar, M.A.; Abuadbba, A.; Moore, K. Detecting misuse of security APIs: A systematic review. *ACM Comput. Surv.* **2025**, *57*, 1–39. [CrossRef]
29. Afolalu, O.; Tsoeu, M.S. Artificial intelligence as the next frontier in cyber defense: Opportunities and risks. *Electronics* **2025**, *14*, 4853. [CrossRef]
30. Jerichow, A.; Covell, B.; Chandramouli, D.; Rezaki, A.; Lansisalmi, A.; Merkel, J. 3GPP non-public network security. *J. ICT Stand.* **2020**, *8*, 57–76. [CrossRef]
31. Dolente, F.; Garroppo, R.G.; Pagano, M. A vulnerability assessment of open-source implementations of fifth-generation core network functions. *Future Internet* **2023**, *16*, 1. [CrossRef]
32. 3GPP. *Security Architecture and Procedures for 5G System (Release 17)*; TS 33.501; 3GPP: Sophia Antipolis, France, 2023.
33. *ISO/IEC 27004:2016*; Information Technology—Security Techniques—Information Security Management—Monitoring, Measurement, Analysis and Evaluation. ISO: Geneva, Switzerland, 2016.
34. 3GPP. *Telemetry Control and Data Collection (Release 18)*; TS 28.532; 3GPP: Sophia Antipolis, France, 2024.
35. Bezerra, L.M.; dos Santos, A.M.; Gonçalves, G.E.; Seruffo, M.C.d.R. Evaluating drift detection methods for failure prediction in 5G network resource management. In Proceedings of the W6G Conference, Natal, Brazil, 19–23 May 2025; pp. 1–8.
36. 3GPP. *Study on Enablers for Network Automation for 5G; Stage 2 (Release 18)*; TS 23.288; 3GPP: Sophia Antipolis, France, 2023.
37. Bouakkaz, S.; Suárez, L.; Cuppens, N.; Cuppens, F. Design of an intelligent trust management architecture for 5G service deployment. In Proceedings of the International Conference on Information Systems Security and Privacy (ICISSP), Porto, Portugal, 20–22 February 2025; pp. 310–317.
38. 3GPP. *Security Architecture and Procedures for 5G System (Release 18)*; TS 33.501; 3GPP: Sophia Antipolis, France, 2024.
39. Gao, S.; Lin, R.; Fu, Y.; Li, H.; Cao, J. Security threats, requirements and recommendations on creating 5G network slicing system: A survey. *Electronics* **2024**, *13*, 1860. [CrossRef]

40. Doanis, P. A Deep Reinforcement Learning Framework for Scalable Slice Orchestration in Beyond 5G Networks. Doctoral Dissertation, Sorbonne University, Paris, France, 2024.
41. Next Generation Mobile Networks Alliance. Automation and Autonomous System Architecture Framework–Phase 2. NGMN White Paper. 2024. Available online: https://www.ngmn.org (accessed on 14 March 2026).
42. Hudic, A.; Smith, P.; Weippl, E.R. Security assurance assessment methodology for hybrid clouds. *Comput. Secur.* **2017**, *70*, 723–743. [CrossRef]
43. Holmström, J.; Ketokivi, M.; Hameri, A.P. Bridging practice and theory: A design science approach. *Decis. Sci.* **2009**, *40*, 65–87. [CrossRef]
44. National Institute of Standards and Technology. *Performance Measurement Guide for Information Security*; NIST SP 800-55 Rev.2; NIST: Gaithersburg, MD, USA, 2024.
45. Olufemi, D.; Ejiade, A.O.; Ikwuogu, F.O.; Olufemi, P.E.; Bobie-Ansah, D. Securing software-defined networks against emerging cyber threats in 5G and future networks. *Int. J. Eng. Res.* **2025**, *14*, 2. [CrossRef]
46. Cunha, J.; Ferreira, P.; Castro, E.M.; Oliveira, P.C.; Nicolau, M.J.; Núñez, I.; Sousa, X.R.; Serôdio, C. Enhancing network slicing security: Machine learning, software-defined networking, and network functions virtualization-driven strategies. *Future Internet* **2024**, *16*, 226. [CrossRef]
47. Islam, M.Z.; Galib, S.M.; Kabir, M.H. A comprehensive analysis of security overview for network slicing of 5G networks. *Int. J. Inf. Comput. Secur.* **2025**, *27*, 159–211. [CrossRef]
48. Rajkumar, V.S.; Ştefanov, A.; Presekal, A.; Palensky, P.; Torres, J.L.R. Cyber attacks on power grids: Causes and propagation of cascading failures. *IEEE Access* **2023**, *11*, 103154–103176. [CrossRef]
49. Chao, L.; Yinghua, L.; Fengqian, D.; Rui, S. Spatio-temporal anomaly detection for 5G-clusters: A multi-scale fuzzy contrastive learning approach. *IEEE Trans. Netw.* **2025**, *33*, 1588–1602. [CrossRef]
50. Lu, J.; Fang, Y.; Fan, S.; Wu, X. Self-supervised intelligent autoencoder with multi-scale convolutional contrastive learning for multivariate anomaly detection. *J. Intell. Fuzzy Syst.* **2025**, *50*, 7. [CrossRef]
51. Kwon, D.; Natarajan, K.; Suh, S.C.; Kim, H.; Kim, J. An empirical study on network anomaly detection using convolutional neural networks. In Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS), Vienna, Austria, 2–6 July 2018; pp. 1595–1598.
52. Ibitoye, J. Zero-trust cloud security architectures with AI-orchestrated policy enforcement for US critical sectors. *Int. J. Sci. Eng. Appl.* **2023**, *12*, 88–100.
53. Eswaran, S.; Honnavalli, P. Private 5G networks: A survey on enabling technologies, deployment models, use cases and research directions. *Telecommun. Syst.* **2023**, *82*, 3–26. [PubMed]