Future Technology
Open Access Journal
ISSN 2832-0379

August 2026| Volume 05 | Issue 03 | Pages 276-286

Journal homepage: https://fupubco.com/futech

https://doi.org/10.55670/fpll.futech.5.3.24

Future Publishing LLC

Article

Securing national connectivity infrastructure through identity resilience: implications for zero trust–aligned telecom security

Shiva Kumara[1]*, Maunik Shah[2]

[1]Independent Researcher, Kuvempu University, Karnataka, India
[2]Independent Researcher, Indian Institute of Technology Bombay, Maharashtra, India

ARTICLE INFO

Article history:
Received 07 January 2025
Received in revised form
05 May 2026
Accepted 22 June 2026

Keywords:
Identity resilience, Zero Trust security,
Telecom control plane,
National connectivity infrastructure

*Corresponding author
Email address:
reachkumaras@gmail.com

DOI: 10.55670/fpll.futech.5.3.24

ABSTRACT

The importance of identity-centric controls for securing national connectivity infrastructure in cloud-native telecom environments is increasingly recognized. Modern telecom control planes are built on software-defined and service-based architectures. Identities are both a trust boundary and a significant attack surface. This study evaluates the effects of identity compromises on security and operational behavior in a simulated cloud-native telecom control plane. In this paper, we describe a scenario-based experimental approach to assessing three security postures: (i) perimeter-based, (ii) Zero Trust-based, and (iii) Zero Trust-based with basic identity-resilience mechanisms. Our findings demonstrate that perimeter-based security was bypassed in all evaluated attack scenarios and that it provided broad control-plane reachability. Zero Trust aligned security reduced attack success to less than 15% and limited lateral propagation. The attack success rate dropped to zero across all tested scenarios when identity resilience mechanisms were added. The average blast radius reduced from more than five services under perimeter security to near zero with identity-resilient Zero Trust. The measured request-success rate during attack and containment windows decreased from 100% under the perimeter baseline to 0% under the Zero Trust and identity-resilient configurations for unauthorized or quarantined requests. This decrease was primarily due to intentional policy-based denial rather than infrastructure failure. The results in the simulated environment show that identity resilience can enhance Zero Trust by reducing the persistence of compromised identities. The results also show the security-availability trade-offs, which must be further validated in telecom environments at production scale.

1. Introduction

National connectivity infrastructure is a cornerstone of our modern society, facilitating emergency response, financial transactions, health care delivery, government communication, and other vital digital services. Telecom networks are widely regarded as essential infrastructure for their contributions to economic stability, public safety, and integrated digital services [1]. Therefore, security failures in telecom environments can have consequences that go beyond localized technical disruption. Telecom networks are moving to cloud-native and software-defined paradigms, a major architectural shift. Previously, core network functions were implemented as tightly coupled hardware systems. Increasingly, these core network functions are deployed as virtualized and containerized services orchestrated through software platforms. Although this transition enhances agility, scalability, and automation, it also enlarges the attack surface, especially in the control and management planes of 5G networks [2]. Modern telecom cores are based on service-oriented architectures that expose APIs and control protocols, paving the way for new possibilities for adversarial interaction. The control plane has become a major security concern since a breach at this level can affect subscriber management, service orchestration, and interconnection functions. Large-scale disruption can be caused by attacks on authentication, authorization, and orchestration interfaces without exploiting radio-access or user-plane protocols [3], as demonstrated by previous studies. These results indicate a need to re-evaluate security assumptions in cloud-native telecom environments, where network location alone does not imply trust. Zero Trust Architecture has become the leading security model for distributed and cloud-based

systems. It stresses continuous validation, least-privilege access, and explicit authorization of every interaction, dismissing implicit trust based on network location [4]. Zero Trust is especially relevant in telecom systems, where trust relationships increasingly rely on dynamic interactions among services, users, and automation components [5]. However, most existing Zero Trust research focuses on architectural principles, scalability, or policy enforcement rather than on empirical evaluation of compromise outcomes in telecom-specific settings.

A key open issue is what happens to compromised identities once authentication succeeds. Systematic reviews have found limited operational validation of Zero Trust mechanisms, particularly regarding identity misuse and post-compromise behavior [6]. This problem is particularly relevant in telecom control planes where identities include human operators, workload identities, and automated orchestration components. Recent research also points to identity governance and lifecycle management as immature areas of Zero Trust implementation, particularly in highly automated service-to-service environments [7]. Future telecom systems, including 6G and disaggregated multi-vendor architectures, are expected to increase the complexity of identity relationships. Programmable interfaces and autonomous management functions make identity-centric controls increasingly important, while conventional perimeter defenses become less effective in dynamic cloud-native environments [8]. Despite this recognition, system-level evidence remains limited on how identity compromise affects security and availability in telecom control planes. This study addresses this gap by evaluating identity-centric security behavior in a simulated cloud-native telecom control-plane environment. It compares perimeter-based security, Zero Trust-aligned security, and Zero Trust enhanced with basic identity-resilience mechanisms through repeatable, scenario-driven experiments. The research is guided by the following objectives, which are explicitly supported by the results presented in this study:

- To assess the effects of identity compromise on security outcomes in a telecom control-plane environment with a cloud-native architecture.
- To compare perimeter-based and Zero Trust-aligned security architectures in terms of limiting the success of identity-driven attacks.
- To analyze the operational and availability implications of introducing basic identity-resilience mechanisms in Zero Trust-aligned telecom security.

Based on these objectives, the study evaluates the following hypotheses.

H1: Zero Trust-aligned security reduces identity-driven attack success compared with perimeter-based security.

H2: Zero Trust-aligned security reduces post-compromise blast radius compared with perimeter-based security.

H3: Adding identity-resilience mechanisms further reduces the likelihood of sustained compromise compared with standard Zero Trust.

H4: Identity-resilience mechanisms introduce measurable availability trade-offs during containment due to deliberate access restriction.

This work contributes empirical evidence to the debate on securing the national connectivity infrastructure by studying the observable system behavior and identity lifecycle effects. The results are intended to enable researchers and practitioners to understand the role of identity resilience as a key complement to Zero Trust in modern telecom environments.

2. Literature review

Zero Trust security has emerged as a leading paradigm for securing distributed and software-defined systems by eliminating implicit trust and requiring continuous verification. Surveys on Zero Trust security describe the evolution from conceptual trust models to practical mechanisms grounded in authentication, authorization, and least-privilege access control [9]. These principles are particularly relevant to telecom networks, which use service-based architectures with exposed control-plane interfaces rather than dedicated network paths. However, much of the Zero Trust literature is still focused on architectural principles and policy frameworks. There is limited empirical analysis on the impact of identity compromise on system behavior once access has been granted. Identity management is at the core of cloud-native telecom security, since human operators, network functions, and orchestration components all depend on identities for authentication and authorization. Studies of 5G core identity protection show that identity exposure and misuse remain persistent concerns with the increasing programmability and automation of telecom networks [10]. Containerized network functions and microservices increase the number of active identities within the control plane and, as a consequence, the attack surface associated with credential compromise.

Empirical evaluations of open-source 5G core implementations indicate that identity-related vulnerabilities often stem from misconfiguration, over-trust among network functions, and inadequate isolation of the control plane [11]. Once a control-plane identity is compromised, attackers can access multiple internal services unless effective containment mechanisms are in place. These results support identity compromise as a primary attack vector in telecom environments. The majority of the existing studies on 5G authentication and key agreement have focused on the cryptographic strength and authentication efficiency. For initial access procedures, hybrid and blockchain-based authentication schemes have been proposed to improve trust establishment [12]. Comparative evaluations of authentication and key agreement approaches provide useful protocol-level insights, but typically abstract away post-authentication identity behavior [13]. Therefore, current identity-management research offers limited insights into how compromised identities behave over time in cloud-native telecom control planes.

Network function virtualization and service function chaining also bring new security dependencies between virtualized components. NFV and service function chaining reviews show that while virtualization improves flexibility and scalability, it also creates new attack vectors and failure modes that challenge traditional security assumptions [14]. In these cases, enforcement decisions can directly affect service availability, particularly when control-plane components rely on tightly coupled orchestration workflows. Telecom virtualization that is cloud native makes it even more difficult to manage east-west traffic and control service dependencies. Virtual networking for cloud-native NFV environments has been studied with a focus on the need for fine-grained isolation to prevent unauthorized service interactions without risking unintended service disruption [15]. This is a good example of the conflict between security enforcement and operational continuity in telecom systems. Recent Kubernetes security research also shows that security responses can impact performance and availability.

Cloud-based denial-of-service attacks can be detected using kernel-level observability methods, such as extended Berkeley Packet Filter instrumentation. However, if the response actions are misaligned with operational requirements, they may cause latency or reduced service availability [16]. Our results suggest that resilience-oriented security design must consider both the effectiveness of enforcement and the consequences for availability. Overall, the reviewed literature indicates three gaps. First, studies of Zero Trust in telecom contexts are largely conceptual or architectural rather than empirical studies of identity compromise. Second, much of the research in identity management has focused on authentication rather than post-compromise behavior. Third, research on resilience and availability includes enforcement effects but rarely discusses identity theft as a trigger. This study addresses these gaps by assessing perimeter-based security, Zero Trust-aligned security, and Zero Trust with basic identity resilience in a simulated cloud-native telecom control plane using scenarios.

3. Methodology

This study uses a controlled simulation-based approach to evaluate identity-based security behavior in a cloud-native telecom control-plane environment. It compares perimeter-based security, Zero Trust-aligned security, and Zero Trust with basic identity-resilience mechanisms through scripted, repeatable attack scenarios. Structured logs are used to measure attack success, lateral propagation, containment behavior, and availability effects. Figure 1 summarizes the workflow from system scope definition and identity modeling to attack execution, metric extraction, comparative analysis, and interpretation.

3.1 Experimental scope and system boundary

The study adopts a simulation-based approach to the study of identity-centric security behavior in national connectivity infrastructure, with a particular focus on the cloud-native telecom control and management plane.

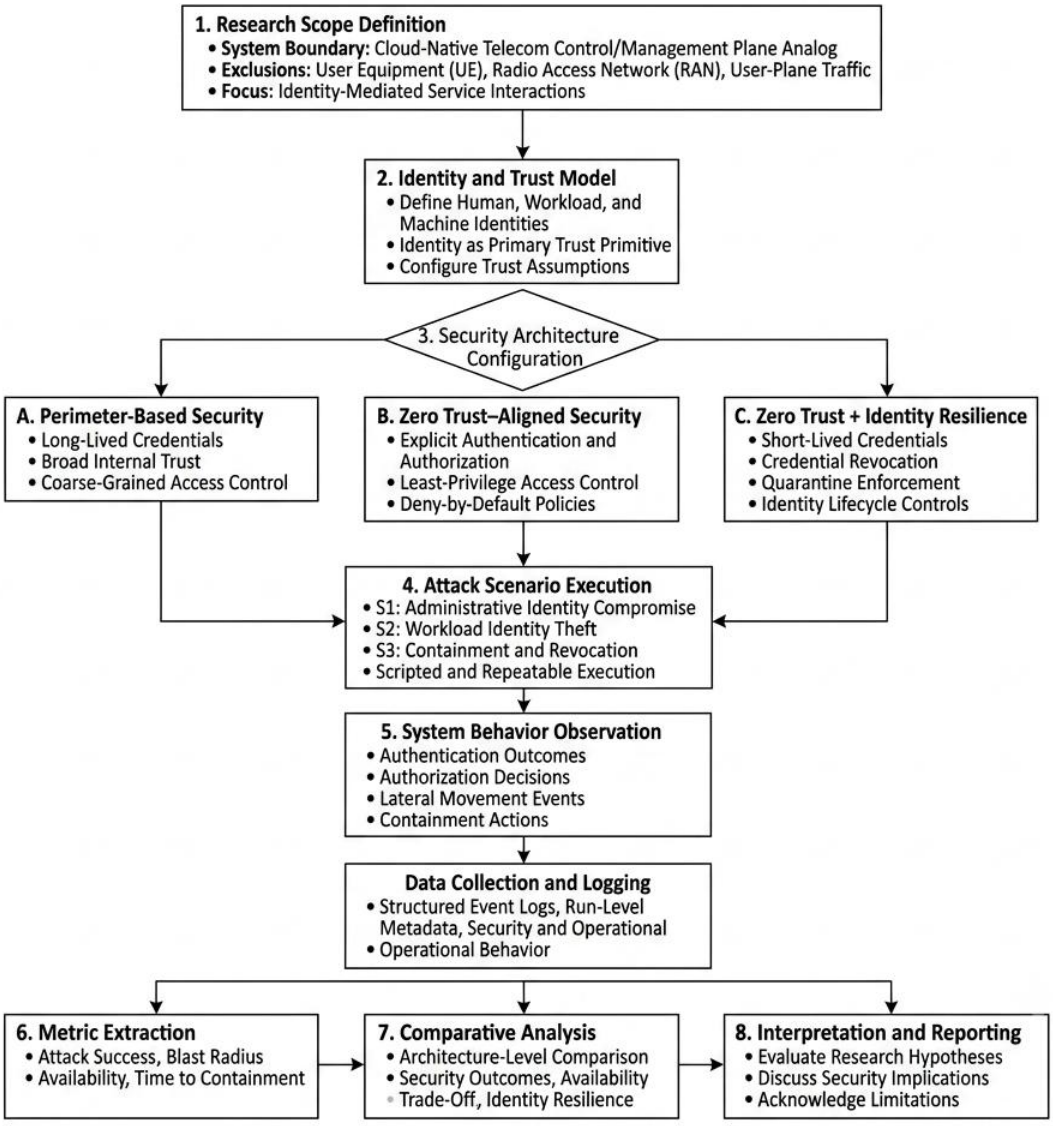

Figure 1. Simulation workflow for evaluating identity resilience in Zero Trust-aligned telecom security

Service interactions are mediated by identity, and the system boundary includes control-plane equivalent services and orchestration interfaces. We selected this scope because compromising the control or management plane can lead to service disruption at the national scale. The study considers identity-driven trust relationships but not protocol-layer or physical-layer vulnerabilities, user equipment, radio access networks, or high-volume user-plane traffic. The simulated environment is a telecom control-plane analog consisting of interconnected microservices that model interdependencies between cloud-native network functions. These services communicate via authenticated service-to-service requests in accordance with service-based telecom architectures. The approach focuses the analysis on identity-mediated interactions, thereby enabling a targeted evaluation of the roles of identity compromise and identity resilience in shaping security outcomes and control-plane behavior.

3.2 Simulation environment implementation

The experimental testbed was implemented as a local Kubernetes-based cloud-native environment to simulate identity-mediated interactions in a telecom control plane. We deployed containerized services on a kind cluster consisting of one control-plane node and two worker nodes. Container images were managed via Docker Engine 29.1.3, and deployment and interaction with the cluster were done through the kubectl client v1.34.1. We used Istio service mesh enforcement to support mutual authentication, authorization policy enforcement, and access control restrictions between services. The simulation control plane consisted of eight interconnected microservices representing control-plane functions of a telecom control plane, including a protected critical service. These services communicated via HTTP-based service-to-service requests, enabling controlled evaluation of identity compromise, lateral movement, authorization enforcement, and containment behavior across different security architectures. To support repeatability, deployment, attack execution, and metric collection were automated using scripts and Kubernetes manifests. Each architecture-scenario pair was run 30 times, leading to 270 total experimental runs. The simulation scripts, Kubernetes manifests, run-level logs, metric extraction files, and generated outputs are packaged as supplementary replication materials to enable reproducibility and describe the deployment configuration, attack execution, and data-processing workflow.

3.3 Identity model and trust assumptions

The methodology considers identity as the primary trust primitive in the simulated telecom environment. We consider three identity planes: human identities for administrators and operators, workload identities for control-plane services, and machine identities for automation and orchestration components. Authentication and authorization decisions of these identity planes govern access to control-plane services and communication between services. Management actions are taken by human identities that are either administrator or operator accounts. Workload identities are the running instances of services (like the simulated control-plane microservices). Machine identities are automation and orchestration components (Kubernetes service accounts, deployment scripts, service-mesh enforcement mechanisms, etc.). These distinctions are necessary because administrative abuse, workload identity theft, and automated containment involve different trust relationships, authorization paths, and enforcement points. Trust relationships are transient and contextual. Perimeter-based configurations enable long-lived identities to be broadly trusted within the environment. Zero Trust-based configurations explicitly authenticate identities for every interaction and control access through least-privilege authorization. Identity-resilience mechanisms further alter trust assumptions, e.g., through time-limited credentials, revocation, and containment workflows. These assumptions establish common ground for evaluating how identity compromise behaves under various security postures.

3.4 Security architectures evaluated

Three security architectures were compared to observe the effects of different security postures on identity-driven compromise, lateral movement, containment, and availability. The same simulated control-plane services, attack scenarios, and measurement procedures were applied to all architectures, so that any differences in the results can be attributed to security configuration rather than system composition. The first architecture is a perimeter-based baseline. It employs coarse-grained trust boundaries, limited internal segmentation, permissive internal access, and long-lived credentials. This configuration indicates an older security posture in which, once initial access is achieved, services within the trusted environment are generally available. The second architecture is a Zero Trust-aligned security posture. It requires explicit service-to-service authentication, deny-by-default authorization, and least-privilege access control. Service communication is limited to defined trust relationships, and requests are refused by policy rather than implicitly trusted based on network location. The third architecture extends the Zero Trust configuration with mechanisms for identity resilience.

This configuration maintains the same least-privilege authorization structure as the Zero Trust architecture, but adds short-lived credentials, token validation at the protected critical service, revocation logic, and quarantine enforcement. The controls have been implemented to reduce the time a compromised identity remains active after it has been abused. Table 1 indicates the configuration parameters used in the three evaluated architectures. This configuration table explains the implementation of the “basic identity resilience” in the experiment. In particular, the identity-resilient configuration reduced the credential lifetime from 3600 seconds to 15 seconds, required token validation at the protected critical service, rejected credentials issued before the revocation event, and enforced quarantine via an Istio authorization policy that denied traffic from the attacker context.

The assessed Zero Trust configurations were developed to be in accordance with the fundamental tenets of Zero Trust Architecture as defined in NIST SP 800-207 [17]. Zero Trust and identity-resilient architectures, in particular, enforce explicit checks of service interactions, access control according to least-privilege principles, authorization policies based on deny-by-default, and the elimination of implicit trust based on network location. Identity-resilient configuration extends these principles further with identity lifecycle controls such as short-lived credentials, revocation, and quarantine mechanisms. The goal of this study is not to reproduce the complete NIST reference architecture, but to explore the effects of identity compromise and identity-resilience mechanisms on security outcomes and operational behavior in a cloud-native telecom control-plane environment.

Table 1. Security architecture configuration parameters

Configuration parameter	A: Perimeter baseline	B: Zero Trust aligned	C: Zero Trust with identity resilience
Internal trust model	Broad internal trust	Explicit verification	Explicit verification with lifecycle control
Service authentication	Limited internal authentication	Istio service-mesh authentication	Istio service-mesh authentication
Authorization model	Coarse-grained/permissive	Deny-by-default, least -privilege	Deny-by-default, least -privilege
Credential lifetime	3600 seconds	3600 seconds	15 seconds
Protected service token validation	Not enforced	Not enforced at critical endpoint	Enforced at critical endpoint
Credential revocation	Not automated	Not automated	Revocation epoch logic
Quarantine enforcement	Not applied	Not applied	Istio DENY authorization policy
Lateral movement control	Minimal internal restriction	Authorization-policy based restriction	Authorization-policy restriction plus quarantine
Containment behavior	Reactive and limited	Policy-based denial	Revocation and quarantine after misuse trigger

3.5 Attack scenarios and execution strategy

The experimental evaluation was conducted using three identity-based attack scenarios to assess the security behavior of the simulated telecom control plane. The first scenario simulated the compromise of an administrative identity, in which an attacker attempts to perform unauthorized control-plane actions and access protected services. This scenario is designed to test the effect of misusing a human or operator identity. The second scenario was designed to simulate a workload identity theft, where an attacker-controlled service attempted to communicate east-west across the service graph without authorization. This scenario illustrated the risk of service identity abuse in cloud-native telecom systems, where workload credentials can be targeted without exploiting lower-layer telecom protocols. The third scenario assessed containment after identity misuse. Detection was treated as an external trigger rather than as a detection algorithm evaluated in this study. Following the trigger, revocation or quarantine controls were applied, and the system was monitored to see if the compromised identity could still access protected services. Detection latency was not modeled as an independent variable, but delayed detection could lead to longer exposure times in operational deployments. All scenarios were executed using scripted, deterministic procedures within each security architecture to ensure repeatability and consistent comparison.

The threat model assumes an attacker operating from a separate attacker namespace using an attacker tooling pod. The attacker can issue HTTP requests and perform limited Kubernetes API actions, depending on the permissions available in each evaluated architecture. The attacker is not assumed to compromise Kubernetes nodes, the container runtime, service source code, or the Istio control plane. Therefore, the study focuses on identity-mediated control-plane misuse rather than radio-access attacks, kernel exploits, or user-plane compromise. The three scenarios were selected to represent the main identity-related risks in cloud-native telecom control planes: administrative identity misuse, workload identity theft, and post-detection containment. Detection in the containment scenario is treated as an external trigger rather than as a detection algorithm evaluated in this study. The threat model, attacker capabilities, and success criteria for the evaluated attack scenarios are summarized in Table 2.

This scenario selection intentionally isolates identity-driven control-plane behavior. Broader threats, including privilege escalation, credential chaining, insider activity, supply-chain compromise, and AI-assisted attacks, are addressed as limitations and future research directions rather than evaluated directly in the present experiment.

3.6 Data collection and metric extraction

During each simulation run, the deployed services generated structured event logs that recorded authentication decisions, authorization outcomes, service access attempts, lateral movement attempts, request outcomes, and containment actions. These logs served as the primary data source for metric extraction and analysis. No external datasets were used; all measurements were derived from controlled experimental executions. Three categories of metrics were extracted from the collected logs. First, security outcome metrics measured whether an attacker obtained unauthorized access to a protected control-plane service. Second, propagation metrics measured the extent of unauthorized service reachability after identity compromise. Third, operational metrics measured service availability during attack and containment windows using request-success outcomes. The key metrics were formally defined as follows:

$$Attack\ Success\ Rate = \frac{Successful\ Attack\ Runs}{Total\ Experimental\ Runs} \quad (1)$$

$$Blast\ Radius = Number\ of\ distinct\ unauthorized\ services\ reached \quad (2)$$

$$Availability = \frac{Successful\ Service\ Requests}{Total\ Service\ Requests} \quad (3)$$

$$Time\ to\ Containment = T_{first\ denial} - T_{containment\ action} \quad (4)$$

Attack success was recorded when a compromised identity successfully reached a protected control-plane service without authorization. Blast radius was calculated as the number of distinct downstream services reached through unauthorized lateral movement. Availability was calculated as the proportion of successful service requests during the relevant attack or containment execution window. Time to containment was calculated as the elapsed time between the application of a revocation or quarantine action and the first observed denial of the compromised identity.

Table 2. Formal definition of attack scenarios

Scenario	Identity plane	Attacker capability	Success criterion	Failure criterion
S1: Administrative identity compromise	Human/admin identity	Performs Kubernetes API discovery and attempts access to the protected critical service	Unauthorized access to svc-critical is recorded as allowed	API request or critical-service access is denied
S2: Workload identity theft	Workload identity	Issues east-west HTTP requests across the simulated service graph	Unauthorized lateral movement reaches downstream services or svc-critical	Authorization policy blocks requests and denial events are recorded
S3: Containment and revocation	Workload/machine identity lifecycle	Probes protected service before and after quarantine or revocation	Compromised identity continues accessing the protected service after containment	Access is denied after quarantine or revocation enforcement

Each metric was linked directly to observable log events and aggregated uniformly across all architectures and attack scenarios. This approach ensured that comparisons reflected differences in security posture and identity-resilience behavior rather than differences in measurement procedure.

3.7 Experimental repetition and analysis approach

All architecture-scenario combinations were executed 30 times to minimize the effects of transient execution and enable consistent comparison across security postures. The study considered three security architectures for three attack scenarios. Therefore, the entire data set consisted of 270 runs for analysis. Results were aggregated at the architecture-scenario level to compare security and operational behavior across configurations. The analysis was not performed to optimize absolute performance values, but to explore architecture-level differences in attack success, propagation behavior, containment, and availability. The analysis is descriptive and based on repeated experimental runs. Because the scripted scenarios produced deterministic outcomes in most cases, the results are reported as repeated-run aggregate values rather than inferential statistical significance tests. Visual representations of the studied security architectures were used to illustrate trends. The run-level dataset from the 30 repetitions for each architecture–scenario pair was used to descriptively compare attack success, blast radius, authorization-denial events, containment behavior, and availability across architectures. Because the scripted attack and containment scenarios produced highly consistent outcomes, the manuscript reports repeated-run aggregate values, including means, proportions, and request-success rates, rather than inferential significance tests.

3.8 Methodological limitations

The methodology deliberately simplifies several aspects of production telecom environments. The testbed models identity-mediated interactions among cloud-native control-plane services, but does not fully capture the scale, traffic diversity, protocol complexity, vendor heterogeneity, or operational dependencies of production telecom networks. To keep the evaluation focused on control-plane identity behavior, user equipment, radio access networks, roaming interfaces, and high-volume user-plane traffic were excluded. The attack scenarios we study deliberately focus on identity-mediated control-plane misuse and do not model the full range of modern telecom threats. More complex attack paths, such as privilege escalation, credential chaining, insider misuse, supply-chain compromise, adversarial automation, and AI-assisted reconnaissance, were outside the scope of the experiment. Although these exclusions limit the generalizability of the findings, they allow the study to isolate the specific effects of identity compromise, zero-trust enforcement, and identity-resilience controls. The study also assumes the availability of well-structured logging and scripted attack execution. Human-in-the-loop security operations, detection uncertainty, and independently validated detection mechanisms were not directly modeled. In addition, containment behavior was evaluated after an external detection trigger rather than through a detection system evaluated within the experiment. These assumptions may limit the transferability of the results to multi-vendor or operator-grade environments such as O-RAN deployments. Therefore, the findings should be interpreted as controlled simulation evidence of architectural behavior rather than as direct proof of production-network performance. Future work should validate the approach in larger testbeds with heterogeneous network functions, realistic traffic loads, variable detection delays, recovery workflows, latency and throughput measurements, inter-operator scenarios, and compound attack models.

4. Results

This section presents the simulation results for identity-centric security behavior in the cloud-native telecom control-plane environment. The results are organized to cover the three confirmed research objectives: (i) the analysis of the impact of the identity compromise, (ii) the comparison of the perimeter-based and Zero Trust-aligned architectures, and (iii) the analysis of the operational implications of the introduction of basic identity resilience mechanisms. All results are based on multiple experimental runs, summarized in a canonical data set constructed from structured logs and metadata.

4.1 Experimental coverage and dataset summary

Table 3 summarizes the experimental design and dataset characteristics. Each security architecture was evaluated under the same three identity-focused attack scenarios. Every architecture–scenario combination was repeated 30 times to ensure consistency and reduce the influence of transient execution effects, resulting in a total of 270 runs analyzed.

4.2 Impact of identity compromise on security outcomes

The first objective assessed the impact of identity compromise on security outcomes. Table 4 reports the repeated-run aggregate attack-success rate for each security architecture across the three identity-focused attack scenarios. Figure 2 complements these numerical results by illustrating the observed attack reachability and policy-

enforcement behavior across the evaluated architectures. The repeated-run aggregate results showed consistent differences across the evaluated architectures. Attack success was highest under the perimeter-based architecture, substantially lower under the Zero Trust architecture, and fully suppressed in the identity-resilient configuration. Because the scripted scenarios produced deterministic outcomes in most cases, these values are interpreted as descriptive repeated-run aggregates rather than inferential statistical significance results. The perimeter-based configuration allowed unauthorized reachability to downstream services and the protected critical service, while the Zero Trust and identity-resilient configurations restricted propagation through authorization denial, revocation, and quarantine enforcement. Under the perimeter-based architecture, identity compromise consistently resulted in successful access to the protected control-plane service across all evaluated scenarios. This indicates that, in the absence of fine-grained internal authorization, compromised identities can directly translate into control-plane security failure.

Table 3. Experimental design overview

Parameter	Value
Security architectures	A perimeter, B zero trust, C zt identity resilience
Attack scenarios	S1 admin compromise, S2 workload identity theft, S3 containment revocation
Repetitions per case	30
Total analyzed runs	270
Data sources	Structured service logs, per-run metadata
Key metrics	Attack success, blast radius, availability, containment observability

Table 4. Repeated-run aggregate attack success rate

Architecture	S1: Admin Compromise	S2: Workload Identity Theft	S3: Containment Scenario
A perimeter	1.00	1.00	1.00
B zero trust	0.03	0.00	0.00
C zt identity resilience	0.00	0.00	0.00

In contrast, the Zero Trust-aligned architecture substantially reduced the success of attacks. A small fraction of successful outcomes remained in the administrative identity scenario, but successful compromise was suppressed in the workload identity theft and containment scenarios. These results indicate a lower observed compromise rate under Zero Trust-aligned enforcement compared with the perimeter baseline. The identity-resilient Zero Trust configuration further reinforced this effect. No successful compromise was observed in the evaluated scenarios, suggesting that short credential lifetimes, revocation logic, and quarantine enforcement limited compromised identities' ability to sustain access within the scope of the experiment.

4.3 Attack propagation and policy enforcement behavior

The second objective examines how security architectures affect the propagation of attacks following an identity compromise. Table 5 reports repeated-run aggregate values for blast radius and authorization-denial events under the workload identity theft scenario. The results show that the perimeter-based architecture enabled broad internal propagation, while the Zero Trust-aligned and identity-resilient configurations reduced the observed blast radius to zero by enforcing authorization policies.

Under workload identity theft, the perimeter-based architecture allowed broad internal propagation, with compromised service identities reaching multiple downstream services. This behavior indicates that limited internal segmentation can allow identity misuse to spread once access is obtained. Both Zero Trust-aligned architectures prevented compromised identities from reaching additional services, resulting in a blast radius of zero. However, attempts at lateral movement still occurred and were systematically denied by authorization policies. These results indicate that Zero Trust reduces the impact of identity compromise primarily through policy enforcement rather than by preventing attacks altogether.

4.4 Operational and availability implications of identity resilience

Time-to-containment results were limited to cases where containment behavior was both observable and instrumented. A measurable containment delay appeared only under the perimeter-based architecture, where compromised identities persisted long enough to require a reactive containment response.

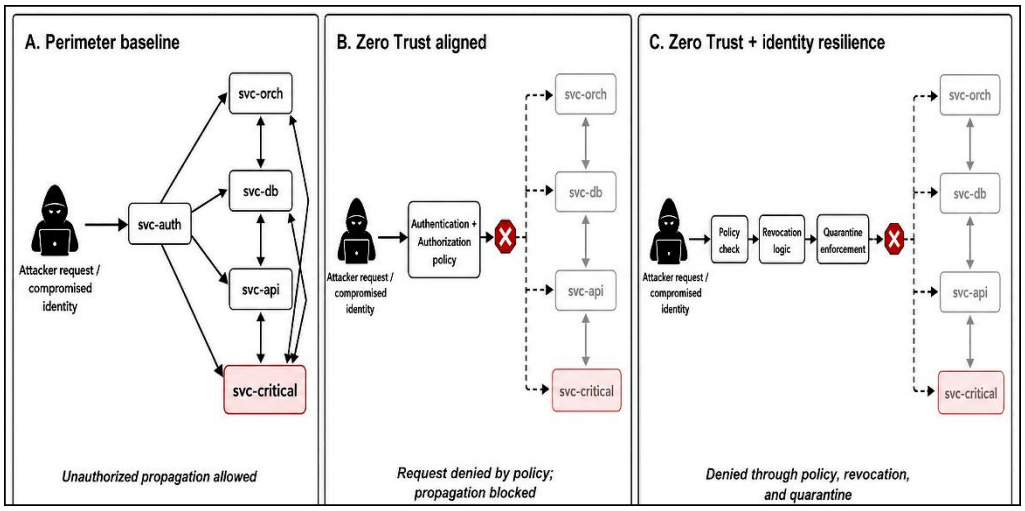

Figure 2. Attack success rate across security architectures and identity-focused attack scenarios

In the Zero Trust-aligned and identity-resilient configurations, identity misuse was denied immediately by authorization policies, leaving no sustained compromise interval to measure. Reporting these cases as zero seconds would be misleading, as the observed behavior reflects immediate prevention rather than delayed containment. Therefore, the reported containment value should be interpreted as indicative rather than comparative. Table 6 summarizes operational metrics and containment observability.

Table 5. Repeated-run aggregate propagation and enforcement indicators

Architecture	Scenario	Blast Radius (Services)	Authorization Denials
A perimeter	S2 workload identity theft	8.0	0
B zero trust	S2 workload identity theft	0.0	8.0
C zt identity resilience	S2 workload identity theft	0.0	8.0

Figure 3 illustrates the containment sequence observed in the identity-resilient Zero Trust configuration. The sequence shows how a compromised identity attempt was evaluated, marked for revocation or quarantine, and then denied in subsequent access attempts. This supports the interpretation that identity-resilience controls prevented sustained compromise after misuse was triggered. The sequence shows how a compromised identity is evaluated, revoked or quarantined, denied in subsequent access attempts, and prevented from achieving sustained compromise.

Even when compromise was successful, availability remained high under the perimeter-based architecture because there was no enforcement mechanism that limited request processing. However, this apparent availability was coupled with a complete loss of security, highlighting the limitations of interpreting availability in isolation. In the Zero Trust-aligned architectures, availability was affected during attack and containment windows as unauthorized requests were blocked per policy. This effect was most visible in the identity-resilient setup during containment, where revocation and quarantine actions were intentionally used to limit access. These results show a security–availability trade-off which is especially important for national connectivity infrastructure.

Table 6. Operational metrics and containment observability

Architecture	Scenario	Availability / request-success rate	Availability degradation	Containment observed
A perimeter	All scenarios	100%	0%	Limited
B zero trust	Attack scenarios	0%	100%	Limited
C zt identity resilience	Containment scenario	0%	100%	Observed

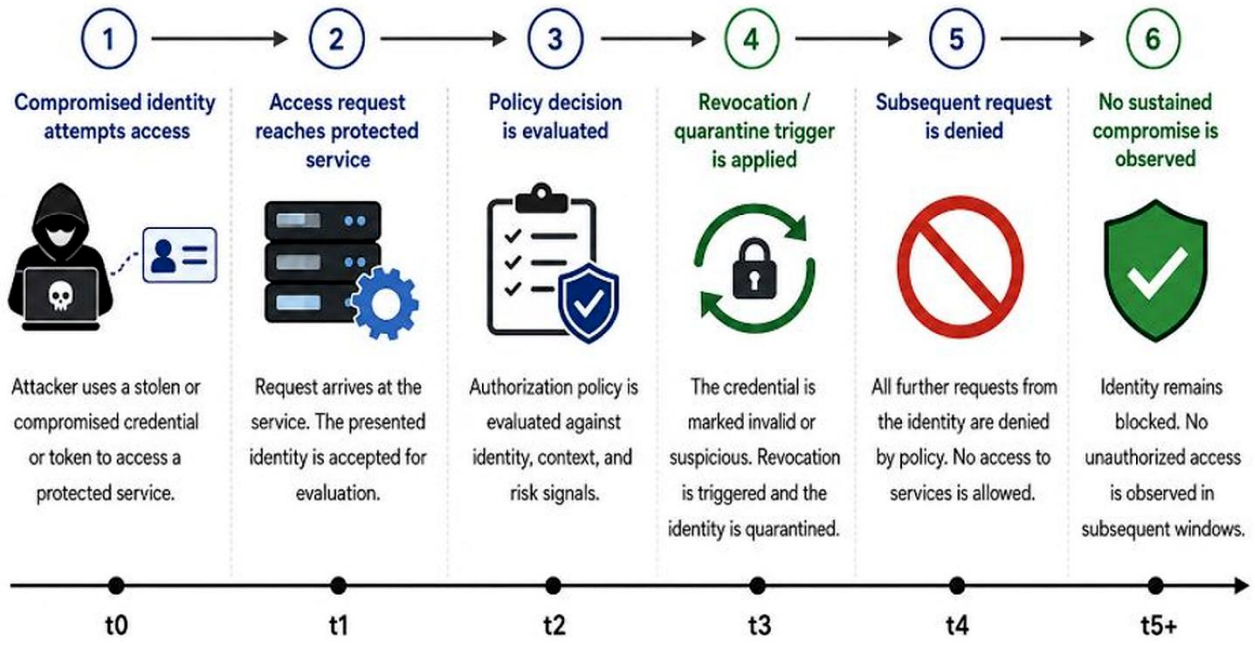

Figure 3. Identity misuse containment timeline

Here, availability was measured as the request-success rate during the attack and containment execution windows. The perimeter-based architecture maintained a 100% request success rate, corresponding to 0% availability degradation, because no enforcement mechanism restricted request processing. The Zero Trust-aligned configuration reduced request success to 0% in attack scenarios, corresponding to a 100% degradation in request success for unauthorized requests. The identity-resilient configuration also showed a 0% request success rate during the containment scenario, corresponding to a 100% degradation in request success for quarantined or unauthorized requests. This reduction should be interpreted as intentional policy-driven denial rather than infrastructure failure.

4.5 Summary of results

Overall, the results support H1 and H2, demonstrating that Zero Trust-aligned enforcement decreased both attack success and lateral propagation compared to the perimeter-based baseline. Further, the results are consistent with H3: the identity-resilient configuration did not allow sustained compromise across the evaluated scenarios. H4 is partially supported because Zero Trust and identity-resilience controls reduced request-success rates during enforcement windows. The observed request-success degradation was 100% for both the Zero Trust-aligned and identity-resilient containment configurations for unauthorized or quarantined requests. This trade-off resulted from intentional policy restriction rather than infrastructure failure.

5. Discussion

The results are discussed with respect to the four hypotheses and the overall goal of evaluating identity resilience as a complement to Zero Trust in cloud-native telecom control-plane environments. This comparison of perimeter-based and Zero Trust-aligned architectures shows that applying enforcement at the identity level greatly affects the outcome of a compromise. This is consistent with previous work on Zero Trust deployment in telecom and 6G scenarios, which focuses on authentication, authorization, scalability, and orchestration complexity but generally lacks empirical evidence on how compromised identities affect attack success, propagation, and service behavior. The results are consistent with H1, in the sense that Zero Trust-aligned security reduced identity-driven attack success relative to the perimeter-based baseline. In all the scenarios we considered, compromised identities could access protected control-plane services under the perimeter model. Zero Trust-aligned enforcement, in contrast, reduced successful compromise by requiring explicit authentication and authorization for service interactions. But the residual success observed in the administrative compromise scenario shows that Zero Trust doesn't completely eliminate identity risk under static credential lifecycles. This is in line with theoretical discussions that describe Zero Trust as a model of continuous verification and containment, rather than as a complete prevention mechanism.

The findings also provide further support for H3 by demonstrating that identity-resilience mechanisms prevented persistent compromise in the scenarios tested. The identity-resilient configuration used shorter credential lifetimes, revocation logic, token validation, and quarantine enforcement to reduce the window of opportunity for compromised identities to be used. This implies that the effect of identity lifecycle governance can be material on compromise consequences in the simulation environment evaluated. The current results suggest that relatively simple resilience-oriented controls can provide practical security value without additional architectural complexity, compared with adaptive Zero Trust approaches that use machine learning or blockchain-based policy optimization [18]. The results also support H2, showing that Zero Trust-aligned enforcement reduced lateral movement after a compromise. In the workload identity theft scenario, the perimeter-based architecture had overall internal reachability, while the Zero Trust configurations reduced the observed blast radius to zero. This is consistent with criticism already leveled against perimeter-based security in software-defined telecom environments, where internal trust assumptions can be unsafe once access is granted. Existing work on SDN-based 5G security focuses on centralized control and fine-grained enforcement. However, these approaches may still be limited if internal identities are implicitly trusted [19].

These observations are relevant for open and disaggregated telecom architectures such as O-RAN and future 6G systems. Multi-layered and multi-vendor trust boundaries are the basis for hybrid Zero Trust models for O-RAN, while the increasing complexity of trust relationships in open RAN ecosystems is reflected in federated learning-based approaches [20,21]. The survey and threat analysis on O-RAN also identify openness, exposed interfaces, and orchestration complexity as important risk factors [22,23]. We provide system-level evidence here that identity compromise can impact reachability and enforcement behavior, and that identity-resilience controls can mitigate sustained exposure in cloud-native control-plane settings. The results partially support H4, showing that Zero Trust and identity-resilience controls reduced request success during enforcement windows. The observed request-success degradation was 100% under the Zero Trust-aligned and identity-resilient containment configurations for unauthorized or quarantined requests. This reduction should be viewed as policy-driven denial, not infrastructure failure. Thus, the availability results also show a security-availability trade-off: more containment reduces attacker reachability but may also reduce service accessibility during enforcement. Similarly, prior cross-layer security work on 5G/6G network slicing points out that the SDN, NFV, and AI-based controls need to balance enforcement with operational continuity [24].

Table 7 compares representative studies with the present work to put these findings into context of the existing Zero Trust and telecom security research. The comparison indicates that prior work has predominantly focused on conceptual trust models, architectural surveys, or adaptive policy optimization. This study presents scenario-driven evidence on the impact of identity compromise on attack success, lateral movement, and availability behavior in a simulated telecom control-plane environment. The results have potential implications for national connectivity infrastructure, but should be interpreted in the context of the simulated telecom control-plane environment's controlled scope. Telecom control planes support functions critical to economic continuity and public safety; therefore, identity compromise in such environments could pose systemic risk. But the results cannot be interpreted as direct proof of production-network behavior. Instead, they argue that Zero Trust access control can benefit from explicit identity-resilience mechanisms that limit the temporal and functional scope of compromised identities. The availability effects observed during containment also underscore the need for resilience-conscious operational models, such as scoped revocation, degraded-mode operation, redundant service paths, and controlled recovery procedures.

Table 7. Comparison of identity-centric Zero Trust approaches in telecom security and their evaluation scope

Literature	Primary Focus	Treatment of Identity & Zero Trust	Evaluation Approach	Key Results / Comparative Insight
Yongjun Ren et al. [5]	Zero Trust architectures for future telecom and 6G networks	Controller-based authentication and authorization with scalability focus	Architectural survey and systems discussion	Indicates architectural needs of Zero Trust in telecom but does not assess how identity compromise relates to attacks and service disruption.
Hongzhaoning Kang et al. [9]	Zero Trust theory and trust modelling	Identity as continuous verification; compromise remains possible	Conceptual survey and theoretical analysis	Finds that Zero Trust limits but does not stop identity abuse; does not empirically validate compromise outcomes.
Abdulrahman K. Alnaim [18]	Adaptive Zero Trust using AI and blockchain	Identity treated as an optimizable policy element	Simulation with ML-driven metrics	Shows that adaptive policy mechanisms can improve enforcement efficiency, but does not directly evaluate identity lifecycle resilience or availability trade-offs.
This study	Identity-centric security behaviour in telecom control planes	Identity as the primary trust primitive, extended with lifecycle resilience	Scenario-driven simulation of a cloud-native telecom control-plane analogy	Shows that (i) perimeter security allows for complete compromise, (ii) Zero Trust substantially reduces attack success and propagation of attacks, and (iii) basic identity resilience mechanisms prevent sustained compromise with observable trade-offs in availability in the process of being contained.

The evaluation was performed in a simplified cloud-native simulation environment, so the results should be considered as controlled experimental evidence, not as direct evidence of production telecom behavior. The testbed captures identity-mediated service interaction, policy enforcement, lateral movement, and containment behavior. It does not reproduce the full scale, heterogeneity, vendor diversity, roaming dependencies, regulatory constraints, or operational complexity of national telecom networks. Therefore, results are best interpreted as evidence of architectural trends in controlled conditions, with real-world deployment requiring further validation in production-scale or operator-grade telecom testbeds. From a policy and regulatory perspective, the findings suggest that identity lifecycle management should be considered alongside traditional network security controls when assessing the resilience of national connectivity infrastructure. And for cloud-native telecom control planes, resilience guidance may need to emphasize short-lived credentials, fast revocation, scoped quarantine, least-privilege service identities, and auditable policy enforcement. Since this study is based on a controlled simulation, these implications are considered design considerations rather than prescriptive regulatory requirements. In future work, the approach should be validated in larger and more heterogeneous telecom environments. Key extensions are inter-operator and roaming scenarios, production-scale traffic conditions, variable detection delays, deterministic recovery workflows, latency and throughput measurements, and integration with adaptive enforcement mechanisms.

6. Conclusion

This study investigated identity-based security behavior in a simulated cloud-native telecom control plane environment. The evaluation analyzed the behavior of compromised identities under perimeter-based security, Zero Trust-aligned security, and Zero Trust with basic identity-resilience mechanisms. The findings suggest that the repercussions of identity compromise are substantially different depending on the security posture adopted. The evaluated simulation showed that perimeter-based architecture enabled compromised identities to access protected control-plane services while the Zero Trust-aligned architecture greatly reduced attack success and lateral propagation through explicit authentication and authorization. In the tested scenarios, sustained compromise was further constrained by the inclusion of basic identity-resilience mechanisms, such as shorter credential lifetimes and revocation or quarantine controls. This suggests that identity resilience can be a significant complement to Zero Trust in cloud-native telecom control-plane environments. But the results should be considered as controlled simulation evidence rather than direct evidence of behavior in production national telecom infrastructure. The availability effects observed during containment also suggest that security enforcement can have operational trade-offs, with 100% request-success degradation observed during Zero Trust and identity-resilient enforcement windows for unauthorized or quarantined requests. Therefore, future work should evaluate identity resilience in larger and more heterogeneous telecom environments, including inter-operator and roaming environments, production-scale testbeds, deterministic containment and recovery models, latency and throughput effects, and integration with adaptive enforcement policies. Operationally, the results indicate that telecom security architectures should consider identity lifecycle management, Zero Trust enforcement, and continuity-preserving containment mechanisms as interconnected design concerns.

Ethical issue

The authors are aware of and comply with best practices in publication ethics, specifically regarding authorship (avoidance of guest authorship), dual submission, manipulation of figures, competing interests, and compliance with research ethics policies. The authors adhere to publication requirements that the submitted work is original and has not been published elsewhere.

Data availability statement

The manuscript contains all the data. However, additional data will be provided by the corresponding author upon reasonable request.

Conflict of interest

The authors declare no potential conflict of interest.

References

[1] P. Scalise, M. Boeding, M. Hempel, H. Sharif, J. Delloiacovo, and J. Reed, "A systematic survey on 5G and 6G security considerations, challenges, trends, and research areas," Future Internet, vol. 16, no. 3, p. 67, 2024. https://doi.org/10.3390/fi16030067

[2] Q. Tang, O. Ermis, C. D. Nguyen, A. De Oliveira, and A. Hirtzig, "A systematic analysis of 5G networks with a focus on 5G core security," IEEE Access, vol. 10, pp. 18298–18319, 2022. doi: 10.1109/ACCESS.2022.3151000

[3] R. Patil, Z. Tian, M. Gurusamy, and J. McCloud, "5G core network control plane: Network security challenges and solution requirements," Computer Communications, vol. 229, p. 107982, 2025. 10.1016/j.comcom.2024.107982

[4] N. F. Syed, S. W. Shah, A. Shaghaghi, A. Anwar, Z. Baig, and R. Doss, "Zero trust architecture (ZTA): A comprehensive survey," IEEE Access, vol. 10, pp. 57143–57179, 2022. doi: 10.1109/ACCESS.2022.3174679

[5] Y. Ren, Z. Wang, P. K. Sharma, F. Alqahtani, A. Tolba, and J. Wang, "Zero trust networks: Evolution and application from concept to practice," Computers, Materials & Continua, vol. 82, no. 2, 2025. https://doi.org/10.32604/cmc.2025.059170

[6] M. L. Gambo and A. Almulhem, "Zero trust architecture: A systematic literature review," Journal of Network and Systems Management, vol. 34, no. 1, p. 25, 2026. https://doi.org/10.1007/s10922-025-09998-x

[7] S. Mushtaq, M. Mohsin, and M. M. Mushtaq, "A systematic literature review on the implementation and challenges of zero trust architecture across domains," Sensors, vol. 25, no. 19, p. 6118, 2025. https://doi.org/10.3390/s25196118

[8] N. Nahar, K. Andersson, O. Schelén, and S. Saguna, "A survey on zero trust architecture: Applications and challenges of 6G networks," IEEE Access, 2024. doi: 10.1109/ACCESS.2024.3425350

[9] H. Kang, G. Liu, Q. Wang, L. Meng, and J. Liu, "Theory and application of zero trust security: A brief survey," Entropy, vol. 25, no. 12, p. 1595, 2023. https://doi.org/10.3390/e25121595

[10] P. Scalise, M. Hempel, and H. Sharif, "A survey of 5G core network user identity protections, concerns, and proposed enhancements for future 6G technologies," Future Internet, vol. 17, no. 4, p. 142, 2025. https://doi.org/10.3390/fi17040142

[11] F. Dolente, R. G. Garroppo, and M. Pagano, "A vulnerability assessment of open-source implementations of fifth-generation core network functions," Future Internet, vol. 16, no. 1, p. 1, 2023. https://doi.org/10.3390/fi16010001

[12] F. F. Ashrif and R. Ahmad, "A secure and efficient hybrid approach for 5G-AKA in blockchain smart contracts," Computer Networks, p. 111761, 2025. 10.1016/j.comnet.2025.111761

[13] Z. Benfarhi, O. Gemikonakli, and M. A. Mobarhan, "Evaluation of authentication and key agreement approaches of 5G networks," in Proc. Int. Conf. Artificial Intelligence and Applied Mathematics in Engineering, Cham, Switzerland: Springer Nature, Nov. 2023, pp. 194–221. https://doi.org/10.1007/978-3-031-56322-5_15

[14] H. U. Adoga and D. P. Pezaros, "Network function virtualization and service function chaining frameworks: A comprehensive review of requirements, objectives, implementations, and open research challenges," Future Internet, vol. 14, no. 2, p. 59, 2022. https://doi.org/10.3390/fi14020059

[15] L. F. Gonzalez, I. Vidal, F. Valera, R. Martin, and D. Artalejo, "A link-layer virtual networking solution for cloud-native network function virtualisation ecosystems: L2S-M," Future Internet, vol. 15, no. 8, p. 274, 2023. https://doi.org/10.3390/fi15080274

[16] Sadiq, H. J. Syed, A. A. Ansari, A. O. Ibrahim, M. Alohaly, and M. Elsadig, "Detection of denial of service attack in cloud-based Kubernetes using eBPF," Applied Sciences, vol. 13, no. 8, p. 4700, 2023. https://doi.org/10.3390/app13084700

[17] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero trust architecture," NIST Special Publication, vol. 800, no. 207, pp. 1–52, 2020. https://doi.org/10.6028/NIST.SP.800-207

[18] K. Alnaim, "Adaptive zero trust policy management framework in 5G networks," Mathematics, vol. 13, no. 9, p. 1501, 2025. https://doi.org/10.3390/math13091501

[19] J. Yao, Z. Han, M. Sohail, and L. Wang, "A robust security architecture for SDN-based 5G networks," Future Internet, vol. 11, no. 4, p. 85, 2019. https://doi.org/10.3390/fi11040085

[20] M. Hashem Eiza, B. Akwirry, A. Raschella, M. Mackay, and M. K. Maheshwari, "A hybrid zero trust deployment model for securing O-RAN architecture in 6G networks," Future Internet, vol. 17, no. 8, p. 372, 2025. https://doi.org/10.3390/fi17080372

[21] M. El-Hajj, "Secure and trustworthy open radio access network (O-RAN) optimization: A zero-trust and federated learning framework for 6G networks," Future Internet, vol. 17, no. 6, p. 233, 2025. https://doi.org/10.3390/fi17060233

[22] W. Azariah, F. A. Bimo, C. W. Lin, R. G. Cheng, N. Nikaein, and R. Jana, "A survey on open radio access networks: Challenges, research directions, and open source approaches," Sensors, vol. 24, no. 3, p. 1038, 2024. https://doi.org/10.3390/s24031038

[23] M. K. Motalleb, C. Benzaid, T. Taleb, M. Katz, V. Shah-Mansouri, and J. Kim, "Towards secure intelligent O-RAN architecture: Vulnerabilities, threats and promising technical solutions using LLMs," Digital Communications and Networks, 2025. https://doi.org/10.1016/j.dcan.2025.05.001

[24] Z. Allaw, O. Zein, and A. M. Ahmad, "Cross-layer security for 5G/6G network slices: An SDN, NFV, and AI-based hybrid framework," Sensors, vol. 25, no. 11, p. 3335, 2025. https://doi.org/10.3390/s25113335

This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).