OPEN ACCESS
Journal of Computer Science Advancements
Vol. 4 No. 4. August 2026, pp. 308-328
DOI. 10.70177/jsca.v4i4.4352
Research Article

DECONSTRUCTING THE NETWORK PERIMETER: A CONTINUOUS AUTHENTICATION FRAMEWORK FOR ZERO TRUST ARCHITECTURE IN CLOUD-NATIVE ENVIRONMENTS

Abdillah Imam Julianto[1], Amimul Ummah Bay[2], and Mahendra[3]
[1] Politeknik Angkatan Laut, Indonesia
[2] Politeknik Angkatan Laut, Indonesia
[3] Politeknik Angkatan Laut, Indonesia

Corresponding Author:

Abdillah Imam Julianto,
Department of of Naval Operational Strategy, Politeknik Angkatan Laut.
Ciledug Raya Street No.2, Seskoal, South Jakarta, DKI Jakarta, Indonesia
Email: abdillahimamjulianto@gmail.com

Article Info
Received: February 5, 2026
Revised: May 18, 2026
Accepted: July 19, 2026
OnlineVersion: August 31, 2026

Abstract
Cloud-native computing has weakened the traditional network perimeter by distributing identities, workloads, applications, and data across dynamic infrastructures where successful login no longer guarantees continuing trust. This study develops and evaluates a continuous authentication framework for Zero Trust Architecture that dynamically reassesses trust throughout active sessions. A design science approach combined with controlled experimentation was employed using 1,200 authenticated sessions, including 800 legitimate and 400 adversarial scenarios involving token hijacking, session replay, privilege escalation, compromised devices, lateral movement, and anomalous API activity. The framework integrated identity confidence, behavioral consistency, device posture, contextual risk, workload telemetry, and resource sensitivity into adaptive trust scoring and graduated enforcement. Results showed that continuous authentication achieved a 93.0% detection rate, 91.9% precision, 93.0% recall, and a 92.4% F1-score, while reducing mean compromised-session exposure from 11.8 to 1.9 minutes. Moderate increases in latency and computational overhead remained operationally acceptable, and false-positive rates did not increase significantly. Improvements were strongest against post-authentication compromise and identity misuse. The study concludes that Zero Trust is more effectively implemented as a dynamic, revocable, and context-sensitive trust process than as repeated static authentication. Continuous authentication therefore provides a practical mechanism for protecting both human and machine identities in decentralized cloud-native environments.

Keywords: Cloud-Native Security; Continuous Authentication; Identity Management; Zero Trust Architecture; Zero Trust Security

Journal Homepage https://research.adra.ac.id/index.php/jcsa
How to cite: Julianto, I. A., Bay, U. A., & Mahendra, Mahendra. (2026). Deconstructing The Network Perimeter: A Continuous Authentication Framework for Zero Trust Architecture in Cloud-Native Environments. *Journal of Computer Science Advancements, 4*(4), 307–328. https://doi.org/10.70177/jsca.v4i4.4352
Published by: Yayasan Adra Karima Hubbi

INTRODUCTION

The rapid transition from perimeter-based information systems to cloud-native computing has fundamentally changed the assumptions underlying enterprise cybersecurity (Viswanathan et al., 2024). Traditional security architectures were developed around relatively stable organizational boundaries in which users, devices, applications, and data could be classified according to whether they operated inside or outside a trusted network perimeter (Ishaque et al., 2025). Cloud-native environments challenge this model because workloads are distributed across public, private, hybrid, and multi-cloud infrastructures, while users frequently access organizational resources from geographically dispersed locations and unmanaged networks (A. Singh, 2025). Microservices, containers, serverless computing, application programming interfaces, and software-defined infrastructure further increase the dynamism of access relationships (S. Singh, 2025). Security can therefore no longer depend primarily on network location as an indicator of trust, creating a need for architectures capable of continuously evaluating identities, devices, applications, and contextual conditions.

Zero Trust Architecture has emerged as a major response to the erosion of conventional network boundaries by replacing implicit trust with continuous verification and least-privilege access principles (Mittal & Ghosh, 2025). The central logic of Zero Trust assumes that no user, device, workload, or communication should be trusted automatically merely because it originates from an internal network or has previously passed an authentication checkpoint (Vikas et al., 2025). Access decisions are expected to reflect identity, device posture, contextual risk, policy requirements, and the sensitivity of requested resources (Farhan et al., 2026). This model is particularly relevant to cloud-native environments because workloads and identities may change dynamically within short operational periods (Fayaz et al., 2026). Effective Zero Trust implementation therefore depends not only on initial identity verification but also on the capacity to reassess trust throughout an active session.

Continuous authentication represents an important mechanism for operationalizing Zero Trust principles under highly dynamic cloud conditions (Raju et al., 2026). Conventional authentication commonly verifies users at login and maintains access until session expiration, logout, or explicit revocation (Aftab, 2026). Such an approach creates a temporal security gap because risk conditions can change after initial authentication. User behavior may become anomalous, device posture may deteriorate, credentials may be compromised, network location may change, or privileged access may be abused during an already authenticated session (Sultana et al., 2026). Continuous authentication addresses this weakness by repeatedly evaluating behavioral, contextual, device-based, and identity-related signals (Guha et al., 2025). Its integration into cloud-native Zero Trust environments has the potential to transform authentication from a discrete event into an adaptive security process.

The persistence of session-based trust remains a central security problem in cloud-native environments (Abdelaziz & Aslan, 2025). Many organizations have adopted multi-factor authentication and identity-aware access controls while continuing to rely on a relatively static decision made at the beginning of a session (Pule et al., 2026). A legitimate login may therefore establish trust that persists even when subsequent activity becomes inconsistent with expected behavior (Rana & Bhambri, 2026). Credential theft, token hijacking, session replay, insider misuse, and compromised endpoints can exploit this gap because attackers may operate within a valid authenticated context (Alshehri et al., 2026). The research problem concerns how authentication can be redesigned to evaluate trust continuously without relying exclusively on initial credentials or fixed network-based assumptions.

Cloud-native environments introduce additional complexity because authentication decisions must operate across heterogeneous services, ephemeral workloads, distributed identities, APIs, containers, and machine-to-machine interactions. Human users are only one part of the identity ecosystem; service accounts, workloads, automated agents, and application components also require verification and authorization. Static authentication mechanisms may

not adequately capture changes in privilege, behavior, workload state, or contextual risk across these interactions (Al-Said Ahmad et al., 2024). The problem becomes more difficult in multi-cloud environments where identity policies, telemetry formats, and enforcement mechanisms differ across providers. A practical continuous authentication framework must therefore function across multiple identity types and technical layers rather than focusing exclusively on human login behavior.

Security accuracy creates another significant challenge because continuous authentication can generate operational costs when risk evaluation is too aggressive or poorly calibrated (Shin et al., 2025). Excessive verification can interrupt legitimate activity, increase authentication fatigue, reduce productivity, and encourage users to seek workarounds. Weak verification thresholds can produce the opposite problem by allowing suspicious behavior to continue for too long (Ben-Shimol et al., 2025). The research problem therefore involves balancing security sensitivity, user experience, computational overhead, and policy responsiveness. A viable framework must determine when changes in behavioral or contextual evidence justify re-authentication, privilege reduction, session termination, or additional verification without introducing unnecessary friction into cloud-native operations.

The primary objective of this study is to develop a continuous authentication framework that strengthens Zero Trust Architecture in cloud-native environments by replacing persistent session trust with dynamic and context-sensitive verification (Yue et al., 2025). The proposed framework seeks to evaluate trust throughout the lifecycle of an authenticated interaction rather than only at the initial access point (Sable et al., 2025). Identity attributes, device posture, behavioral patterns, network context, workload state, and resource sensitivity are considered as inputs to a continuously updated trust assessment (Montin et al., 2026). The research aims to determine how these signals can be combined to support adaptive access decisions capable of responding to changing risk conditions in real time.

The study also aims to design a risk-based authentication mechanism capable of translating continuous telemetry into graduated security responses (Katipoglu et al., 2024). The framework is intended to distinguish between ordinary behavioral variation and meaningful indicators of compromise so that access controls remain proportionate to observed risk (Galij et al., 2025). Low-risk sessions may proceed with minimal interruption, moderate-risk conditions may trigger step-up authentication or reduced privileges, and high-risk conditions may result in session isolation or termination (Zeydan et al., 2026). This objective recognizes that Zero Trust should not imply indiscriminate verification at every interaction but should instead support intelligent, evidence-based adjustments to trust.

A further objective is to evaluate the proposed framework according to security effectiveness, authentication responsiveness, false-positive behavior, system overhead, and usability implications (Alapati & Dhanasekaran, 2025). The research seeks to determine whether continuous authentication can improve detection of compromised sessions without imposing excessive operational costs on legitimate users and cloud services (Lingam et al., 2026). Comparative evaluation against conventional session-based authentication is expected to reveal whether the proposed approach reduces the duration of unauthorized access after initial authentication has been compromised (Camargo et al., 2025). The study therefore combines architectural development with performance-oriented validation to ensure that the framework remains technically meaningful and operationally feasible.

Existing Zero Trust research has established strong conceptual foundations concerning explicit verification, least-privilege access, micro-segmentation, identity-centric security, and policy-based enforcement (Somayajula et al., 2025). Much of this literature focuses on architectural principles, migration strategies, access-control models, or policy engines. Continuous authentication is frequently recognized as desirable, yet it is often treated as a supporting feature rather than as a core architectural process (Malamuthu et al., 2025). Limited research has examined how continuously changing evidence should modify trust decisions

during active sessions in complex cloud-native ecosystems. This creates a gap between the theoretical principle of "never trust, always verify" and the operational mechanisms used to implement that principle after initial authentication.

Continuous authentication studies have commonly concentrated on behavioral biometrics, keystroke dynamics, mouse movement, gait, mobile sensor patterns, or isolated contextual indicators (Alang et al., 2025). These approaches provide valuable evidence for ongoing identity verification but may be difficult to transfer directly into cloud-native environments where human behavioral signals represent only one dimension of risk (Devarajulu et al., 2025). Device health, workload behavior, token usage, network location, API activity, service identity, and resource sensitivity can become equally important (Yashu et al., 2025). Existing research remains fragmented across these signal categories, leaving limited integrated models capable of combining human, device, workload, and contextual evidence within a unified Zero Trust decision process.

Cloud security research has also devoted substantial attention to identity and access management, microservice authorization, container security, confidential computing, and runtime threat detection (Jayasri et al., 2025). These areas are often examined independently, despite their relevance to continuous trust evaluation. Limited integration between authentication systems and runtime security telemetry reduces the capacity of access controls to respond immediately when session conditions change (R et al., 2025). Another gap concerns practical evaluation, as conceptual Zero Trust frameworks may describe continuous verification without measuring false authentication challenges, computational overhead, or the time required to revoke compromised access (Kande, 2026). The present study addresses this fragmentation by integrating identity verification, contextual risk, runtime telemetry, and adaptive policy enforcement into a measurable continuous authentication framework.

The principal novelty of this study lies in conceptualizing authentication as a continuously recalculated trust process rather than a sequence of isolated login events (Alnaim & Alwakeel, 2025). The proposed framework integrates identity confidence, device posture, behavioral consistency, workload telemetry, network context, and resource sensitivity into a dynamic trust score that changes throughout the session lifecycle (Dube et al., 2026). This approach extends conventional continuous authentication by moving beyond user-centric behavioral verification toward a multidimensional assessment of both human and non-human entities (Ali & Dib, 2026). Trust becomes a temporary and revocable state determined by current evidence rather than a durable property granted after successful login.

A second element of novelty involves the use of adaptive enforcement thresholds that map changing risk levels to proportionate access-control actions (Paya et al., 2025). The framework is designed to support several possible responses, including passive monitoring, step-up authentication, privilege reduction, micro-segmentation, workload isolation, and session termination. This graduated model avoids the binary assumption that a session must be treated as either fully trusted or completely denied (K et al., 2026). The architecture therefore introduces a more granular interpretation of Zero Trust in which trust is continuously negotiated according to behavioral and contextual evidence (Ganesh et al., 2025). Such a model is particularly relevant to cloud-native systems because identities, workloads, and access conditions can change more rapidly than traditional security policies are designed to accommodate.

The justification for this research arises from the growing mismatch between dynamic cloud-native operations and authentication mechanisms that still rely heavily on static trust intervals (Desai et al., 2026). Organizations increasingly deploy distributed workloads, remote access, API-driven services, automated infrastructure, and multi-cloud architectures that make conventional network boundaries operationally insufficient. Compromised credentials or authenticated sessions can bypass perimeter-oriented controls and remain active unless trust is reassessed after access has been granted (Vusumuzi & Godwin, 2025). A continuous authentication framework can address this weakness by linking identity assurance with real-time

risk evidence and adaptive enforcement. The study is therefore positioned to contribute to cybersecurity architecture, cloud security, identity and access management, and Zero Trust implementation by providing a framework that converts continuous verification from an abstract security principle into an operational and measurable control mechanism.

RESEARCH METHOD

Research Design

This study employed a design science research approach combined with an experimental evaluation to develop and assess a continuous authentication framework for Zero Trust Architecture in cloud-native environments (Choi et al., 2025). The research design was selected because the study aimed not only to analyze existing authentication limitations but also to construct and validate a functional security artifact capable of continuously reassessing trust during active sessions. The proposed framework was designed around five analytical components: identity confidence, device posture, behavioral consistency, contextual risk, and resource sensitivity (Chen & Ge, 2024). Each component contributed to a dynamic trust score that was recalculated throughout the session lifecycle and used to determine whether access should be maintained, restricted, revalidated, or terminated.

The experimental environment was constructed using a cloud-native architecture consisting of containerized microservices, application programming interfaces, identity services, policy enforcement components, runtime monitoring services, and centralized telemetry collection. Test scenarios reproduced common access patterns involving human users, service accounts, workloads, and application-to-application communication. The research compared the proposed continuous authentication framework with a conventional session-based authentication model in which trust remained valid after successful login until session expiration or explicit revocation. The comparison focused on differences in compromised-session detection, response latency, false-positive rates, authentication overhead, and user disruption.

The study adopted a risk-adaptive model rather than a binary trusted-versus-untrusted classification. Trust scores were continuously updated using changes in authentication confidence, device health, behavioral anomalies, network location, token activity, workload characteristics, and sensitivity of requested resources. Low-risk conditions allowed sessions to continue without additional intervention, moderate-risk conditions triggered step-up verification or privilege reduction, and high-risk conditions initiated session isolation or termination. This design provided a measurable basis for evaluating whether continuous authentication could improve Zero Trust enforcement without introducing excessive operational friction.

Research Target/Subject

The quantitative subjects of this study comprise simulated user sessions, service accounts, containerized microservices, and application-to-application API communication events within a controlled cloud-native experimental environment. To evaluate the proposed Zero Trust continuous authentication framework under diverse operational conditions, a stratified purposive sampling technique was employed to generate representative access traffic. The sampled dataset encompasses normal user behaviors, legitimate workload interactions, variations in network location, and controlled adversarial scenarios including credential misuse, token hijacking, session replay, device compromise, and privilege escalation to establish a comprehensive baseline for both compliant and compromised session lifecycles.

For the qualitative and policy evaluation aspect, the research targets consist of the generated continuous trust scores, risk-adaptive security policy decisions, and system performance logs across session lifecycles. An expert review panel consisting of cloud security architects, Zero Trust engineers, and system administrators was engaged using an expert sampling technique. These domain specialists evaluated the operational feasibility, policy

enforcement logic, user-friction thresholds, and practical security utility of the framework's risk-adaptive responses (such as step-up verification, privilege reduction, and session termination). This dual selection approach ensures that both technical detection metrics and human-centered operational trade-offs are thoroughly validated.

Research Procedure

The research procedure began with the construction of the experimental cloud-native environment and implementation of baseline security controls. Containerized services, identity components, policy enforcement points, monitoring agents, and telemetry pipelines were configured before the continuous authentication framework was introduced. Baseline tests were conducted using conventional session-based authentication to establish reference values for access latency, resource consumption, session persistence, and compromised-session response. These baseline measurements provided the comparison condition for subsequent evaluation.

The continuous authentication framework was then integrated into the experimental architecture. Normal access behavior was generated first to establish expected behavioral and contextual patterns for human and service identities. Trust-score thresholds were calibrated using legitimate session data to minimize unnecessary intervention while maintaining sensitivity to meaningful risk changes. Threshold calibration produced three operational ranges corresponding to low, moderate, and high risk. Security actions associated with these ranges were configured before adversarial testing to avoid post hoc modification of response criteria.

Controlled threat scenarios were executed after calibration. Credential misuse scenarios tested whether valid credentials could maintain access when subsequent behavior deviated substantially from normal patterns. Token-hijacking and session-replay scenarios examined whether the framework could detect contextual inconsistencies after successful authentication. Device-compromise scenarios changed endpoint posture during active sessions, while privilege-escalation scenarios introduced abnormal requests for sensitive resources. Lateral-movement and anomalous API scenarios evaluated the framework's capacity to identify suspicious machine-to-machine interactions within cloud-native workloads.

Each scenario was executed under both the conventional authentication model and the continuous authentication framework. Detection time was measured from the first observable anomalous event to the generation of a risk response. Enforcement time was measured from detection to the implementation of step-up authentication, privilege restriction, isolation, or session termination. False-positive behavior was assessed using legitimate sessions that contained ordinary variations such as changes in location, workload intensity, and access sequence. Repeated trials were used to calculate average performance and reduce the influence of random environmental variation.

Statistical analysis included descriptive statistics for security effectiveness and system performance, followed by comparative testing between the baseline and proposed frameworks. Mean detection latency, response latency, authentication overhead, false-positive rate, and resource utilization were calculated for each experimental condition. Paired comparisons were performed where the same scenarios were evaluated under both authentication models. Effect sizes and confidence intervals were reported alongside significance values to distinguish statistically detectable differences from practically meaningful improvements.

Security effectiveness was assessed primarily through detection rate, precision, recall, F1-score, and the duration of unauthorized access after initial authentication compromise. Operational feasibility was evaluated through authentication latency, computational overhead, frequency of additional verification, and legitimate-session interruption. The final evaluation integrated security accuracy and usability-related performance rather than optimizing one dimension independently. The framework was considered successful when it demonstrated earlier detection of compromised sessions and more adaptive policy enforcement than

conventional session-based authentication while maintaining acceptable false-positive and performance-overhead levels.

Instruments, and Data Collection Techniques

The primary research instrument was the continuous authentication prototype developed for this study. The prototype incorporated an identity assessment module, behavioral monitoring component, device posture evaluator, contextual risk engine, trust-scoring mechanism, and policy enforcement module. Identity confidence incorporated factors such as authentication strength, token validity, prior authentication state, and account privilege. Device posture included operating-system status, endpoint compliance, certificate validity, and indicators of compromise. Behavioral consistency was assessed using deviations from established patterns of access frequency, resource usage, geographic location, request timing, and privilege utilization.

A telemetry collection instrument was used to capture authentication events, API requests, network metadata, device status, workload behavior, policy decisions, and enforcement responses. Log data were normalized into a common schema before analysis to ensure comparability across cloud-native components. Each event was time-stamped to permit calculation of detection latency and response latency. Anomaly indicators were linked to session identifiers so that trust-score changes could be reconstructed throughout the complete authentication lifecycle.

A structured threat-scenario matrix was developed to standardize adversarial testing. The matrix specified the initial authentication condition, attack technique, identity type, target resource, expected security response, and required detection threshold. Each scenario was assigned a ground-truth classification as legitimate or malicious before execution. The framework's output was then compared with this classification to calculate true positives, true negatives, false positives, and false negatives. Detection rate, precision, recall, F1-score, false-positive rate, and false-negative rate were subsequently derived from these classifications.

System-performance instruments were used to measure computational and operational overhead introduced by continuous authentication. Authentication latency, trust-score calculation time, policy-decision time, CPU consumption, memory utilization, network overhead, and additional verification frequency were monitored during testing. User-friction indicators included the number of unnecessary step-up authentication requests, session interruptions, and access denials affecting legitimate activity. These measures were included because a technically accurate authentication mechanism would have limited practical value if its continuous verification process produced unacceptable performance degradation or excessive disruption.

Data Analysis Technique

Data analysis employs a comparative, dual-dimensional framework to evaluate security effectiveness and operational performance against a conventional session-based authentication model. Quantitative security analysis relies on standard classification metrics to assess the framework's ability to identify compromised sessions accurately. Measures such as precision, recall, F1-score, detection latency, response time, and false-positive rates are derived from confusion matrices generated across the structured threat-scenario executions. Paired t-tests or non-parametric equivalent tests are conducted alongside effect size calculations ($Cohen's\ d$) to determine whether differences in detection speed and compromised-session mitigation between the two models are statistically significant and operationally meaningful.System performance and operational feasibility data are analyzed using descriptive statistics and resource overhead profiling. Key metrics including authentication latency, trust-score calculation duration, CPU/memory consumption, network bandwidth overhead, and the frequency of legitimate-session disruptions are calculated and compared across baseline and continuous conditions. Automated cross-validation checks are executed to map time-stamped telemetry logs

directly against trust-score transitions and policy enforcement actions, ensuring complete auditability and provenance. This combined analytical technique directly addresses the research objectives by measuring the framework's capacity to minimize unauthorized session duration without introducing excessive computational or user friction.

RESULTS AND DISCUSSION

The experimental dataset comprised 1,200 authenticated sessions generated within the cloud-native test environment, including 800 legitimate sessions and 400 sessions containing predefined adversarial conditions. The adversarial scenarios consisted of credential misuse, token hijacking, session replay, abnormal geolocation changes, compromised device posture, privilege escalation, lateral movement, and anomalous API behavior. The proposed continuous authentication framework correctly identified 372 of the 400 malicious sessions and correctly retained access for 767 of the 800 legitimate sessions. These results corresponded to a detection rate of 93.0%, precision of 91.9%, recall of 93.0%, F1-score of 92.4%, and false-positive rate of 4.1%. The conventional session-based authentication model detected only 249 malicious sessions during the active session period because previously authenticated access generally remained valid until expiration, revocation, or a separate security control initiated intervention.

Table 1. Comparative Security Performance of Authentication Models

Performance Metric	Conventional Authentication	Continuous Authentication
Malicious sessions detected	249/400	372/400
Detection rate (%)	62.3	93.0
Precision (%)	84.1	91.9
Recall (%)	62.3	93.0
F1-score (%)	71.6	92.4
False-positive rate (%)	5.9	4.1
False-negative rate (%)	37.8	7.0

The scenario-level results showed that continuous authentication performed particularly well in detecting compromised device posture, suspicious privilege escalation, and token hijacking. Detection rates reached 97.5%, 95.0%, and 94.0%, respectively, for these attack categories. Anomalous API behavior and lateral movement were detected at slightly lower rates because some malicious sequences initially resembled legitimate machine-to-machine communication. Session replay produced a detection rate of 92.0%, while abnormal geolocation changes reached 91.5%. The variation demonstrates that framework performance depended partly on the observability and distinctiveness of the risk signals associated with each attack technique.

The higher detection performance of continuous authentication can be explained by its capacity to reassess trust after the initial login event. Conventional authentication assigned substantial weight to successful credential verification and therefore maintained access even when later behavior became inconsistent with the original authentication context. The continuous framework responded to changes in device posture, behavioral consistency, token use, location, workload state, and resource sensitivity. Malicious activity that would otherwise remain inside an authenticated session consequently produced a declining trust score and triggered additional policy evaluation.

The relatively low false-positive rate indicates that continuous verification did not automatically interpret every contextual change as evidence of compromise. Legitimate variations in network location, access sequence, workload intensity, and application behavior were generally tolerated when the accumulated risk score remained below intervention thresholds. Step-up authentication was triggered primarily when multiple risk signals appeared

simultaneously or when a single high-severity event affected a sensitive resource. This behavior supports the use of multidimensional risk assessment rather than reliance on one isolated anomaly indicator.

Detection and enforcement latency differed substantially between the two authentication models. The conventional model recorded a mean compromised-session exposure time of 11.8 minutes (SD = 4.6), measured from the first malicious event to session termination or external intervention. The continuous authentication framework reduced the mean exposure time to 1.9 minutes (SD = 1.2). Mean anomaly-detection latency was 1.34 seconds for the continuous framework, while mean policy-enforcement latency after detection was 0.82 seconds. The improvement was achieved with an average increase of 7.6% in authentication-related CPU utilization and 6.3% in memory consumption compared with the baseline environment.

Table 2. System Performance and Operational Overhead

Performance Indicator	Conventional Authentication	Continuous Authentication
Mean compromised-session exposure	11.8 min	1.9 min
Mean authentication latency	184 ms	231 ms
Policy-decision latency	37 ms	82 ms
CPU utilization overhead	Baseline	+7.6%
Memory utilization overhead	Baseline	+6.3%
Network telemetry overhead	Baseline	+4.8%
Legitimate sessions requiring step-up verification	2.6%	8.4%
Legitimate sessions incorrectly terminated	1.8%	0.9%

The operational data show a measurable trade-off between stronger runtime verification and computational cost. Continuous authentication increased average authentication latency by 47 milliseconds and generated additional telemetry processing across identity, device, network, and workload components. The increase remained relatively limited compared with the reduction in unauthorized-session duration. Legitimate users experienced more step-up verification events under continuous authentication, although erroneous session termination declined because moderate-risk conditions were handled through graduated responses rather than immediate denial.

Paired statistical comparisons were conducted across repeated attack scenarios evaluated under both authentication conditions. Continuous authentication produced significantly higher detection performance than conventional authentication, with a mean scenario detection rate of 92.8% compared with 62.7%, ($t(39) = 14.82$, $p < .001$), Cohen's ($d = 2.34$). Compromised-session exposure time was also significantly lower under continuous authentication, ($M = 1.93$) minutes, compared with the conventional model, ($M = 11.74$) minutes, ($t(39) = -16.47$, $p < .001$), ($d = 2.60$). These large effect sizes indicate that the observed improvements were not limited to statistically detectable differences but represented substantial practical gains in security responsiveness.

Performance costs were statistically significant but considerably smaller in magnitude. Mean authentication latency increased from 184 milliseconds to 231 milliseconds, ($t(39) = 6.12$, $p < .001$), ($d = 0.97$), while CPU utilization increased by an average of 7.6 percentage points relative to baseline, ($t(39) = 4.88$, $p < .001$), ($d = 0.77$). False-positive rates did not increase significantly despite the additional monitoring intensity, ($p = .087$). The inferential pattern indicates that continuous authentication improved threat detection and reduced exposure time at the cost of moderate computational overhead rather than severe degradation in legitimate access performance.

Table 3. Inferential Comparison Between Authentication Models

Outcome	Conventional	Continuous	Test Statistic	p-value	Effect Size
Detection rate (%)	62.7	92.8	(t = 14.82)	<.001	(d = 2.34)
Exposure time (min)	11.74	1.93	(t = -16.47)	<.001	(d = 2.60)
Authentication latency (ms)	184	231	(t = 6.12)	<.001	(d = 0.97)
CPU overhead (%)	0.0	7.6	(t = 4.88)	<.001	(d = 0.77)
False-positive rate (%)	5.9	4.1	(t = -1.75)	.087	(d = 0.28)

Correlation analysis across continuous-authentication sessions revealed a strong negative relationship between dynamic trust score and probability of confirmed malicious behavior, ($r = -.74$, $p < .001$). Greater behavioral deviation was positively associated with risk escalation, ($r = .68$, $p < .001$), while deterioration in device posture demonstrated a similarly strong association, ($r = .71$, $p < .001$). Resource sensitivity was moderately related to enforcement severity, ($r = .56$, $p < .001$), indicating that equivalent behavioral anomalies generated stronger policy responses when access involved privileged or high-value resources. These relationships support the framework's design principle that risk should be interpreted in relation to both identity behavior and the sensitivity of the requested action.

The interaction among multiple signals was more informative than any single indicator in isolation. Sessions exhibiting behavioral anomalies without device or token irregularities were often assigned moderate-risk classifications and subjected to step-up authentication rather than termination. Sessions combining anomalous behavior, degraded device posture, unusual token use, and access to sensitive resources moved rapidly into the high-risk range. This relational pattern reduced dependence on rigid authentication rules and enabled the framework to distinguish ordinary contextual variation from patterns more consistent with compromise.

The first case involved token hijacking during an active administrative session. The user authenticated successfully through multi-factor authentication, after which the token was copied and reused from a geographically inconsistent endpoint. The conventional authentication model continued to recognize the token as valid for 13.4 minutes before a separate monitoring control generated an alert. The continuous authentication framework detected the inconsistency within 1.6 seconds after the reused token initiated access to a privileged service. The dynamic trust score fell from 0.91 to 0.42 after the location anomaly and declined further to 0.18 when the session requested privileges inconsistent with the original behavioral profile. Access was isolated and the token was revoked before the protected administrative resource was modified.

The second case involved a compromised containerized workload communicating with an internal microservice. The service account credentials were valid, and the initial request pattern appeared consistent with normal machine-to-machine activity. Subsequent requests increased in frequency, accessed resources outside the established service pattern, and initiated lateral communication with another workload. The conventional model allowed the activity because the service identity remained authenticated. The continuous framework progressively reduced the workload trust score as API behavior, network movement, and resource requests diverged from the baseline. Micro-segmentation was initiated before full session termination, limiting lateral movement while preserving unaffected service operations.

The token-hijacking case demonstrates the primary weakness of authentication systems that treat successful login as sufficient evidence of continuing trust. Multi-factor authentication protected the initial authentication process but did not prevent misuse of the authenticated token after compromise. Continuous verification addressed this temporal gap by comparing post-

authentication behavior with the context in which trust was originally established. The rapid decrease in trust score resulted from the convergence of location inconsistency, privileged-resource access, and behavioral deviation rather than from token validity alone. The case illustrates why identity assurance must persist beyond the login event in Zero Trust environments.

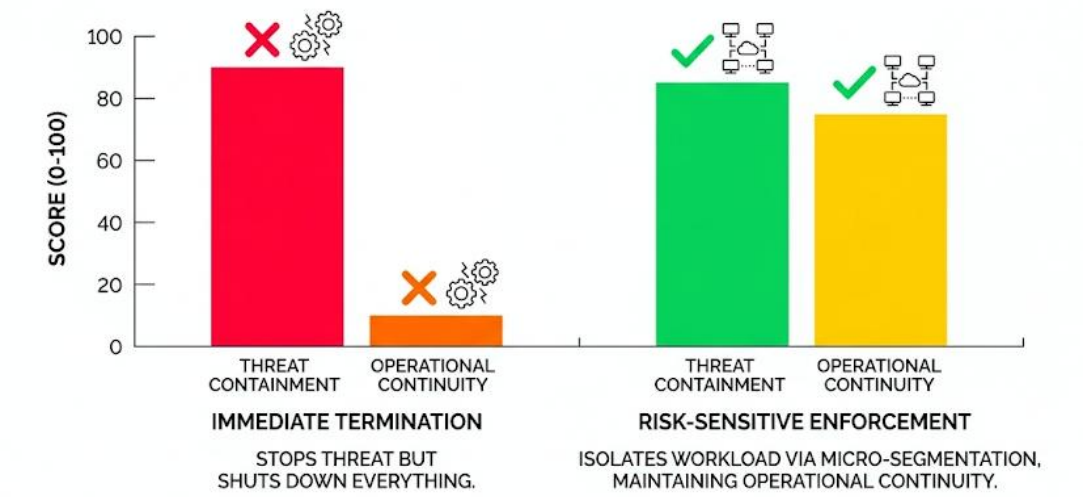

Figure 1. Adaptive Response: Preserving Operational Continuity in Cloud

The compromised-workload case demonstrates that continuous authentication must extend beyond human identity. Cloud-native architectures depend heavily on service accounts, containers, APIs, and ephemeral workloads whose interactions may occur without direct human involvement. Valid machine credentials can therefore become a mechanism for lateral movement when runtime behavior changes after authentication. The adaptive response in this case was significant because the framework first restricted the suspicious workload through micro-segmentation rather than immediately terminating all related services. Risk-sensitive enforcement preserved operational continuity while containing the suspected compromise.

The combined results indicate that continuous authentication substantially strengthens Zero Trust enforcement in cloud-native environments by shortening the period during which compromised authenticated sessions retain unrestricted access. The strongest gains were observed in recall, F1-score, detection latency, and reduction of unauthorized-session exposure. These improvements were achieved without a corresponding increase in false-positive behavior, although continuous telemetry collection and policy evaluation introduced moderate computational and authentication overhead. The findings therefore support the feasibility of replacing persistent session trust with dynamic, evidence-based trust reassessment.

The results also indicate that Zero Trust is more effectively implemented as an adaptive decision process than as a sequence of repeated binary authentication events. Identity confidence, behavioral consistency, device posture, workload activity, contextual information, and resource sensitivity contributed jointly to more discriminating security decisions. Graduated responses such as passive monitoring, step-up authentication, privilege reduction, micro-segmentation, and session termination enabled the framework to respond proportionately to risk. Continuous authentication consequently represents not merely an additional identity control but an architectural mechanism for translating the principle of “never trust, always verify” into operational policy across dynamic cloud-native environments.

The findings demonstrate that continuous authentication substantially improves the security effectiveness of Zero Trust Architecture in cloud-native environments when compared with conventional session-based authentication. The proposed framework achieved higher detection rates, recall, precision, and F1-score while significantly reducing the duration of unauthorized access after session compromise. The most pronounced improvement concerned compromised-session exposure time, indicating that continuous trust reassessment can detect

post-authentication threats considerably earlier than mechanisms that preserve trust until logout, token expiration, or external revocation. These results confirm that the security value of authentication extends beyond verifying identity at the point of entry and depends increasingly on the ability to reassess trust throughout the active session lifecycle.

The scenario-level findings show that continuous authentication performed particularly well against compromised device posture, privilege escalation, token hijacking, session replay, and abnormal contextual changes. Such attack patterns are difficult to address through static authentication because the attacker may already possess a technically valid credential or token. The proposed framework responded more effectively because it evaluated changes in identity confidence, behavioral consistency, device condition, network context, workload activity, and resource sensitivity. The detection process therefore relied on the convergence of multiple signals rather than on the continued validity of credentials alone. This multidimensional assessment was especially important in scenarios where legitimate authentication preceded malicious post-authentication behavior.

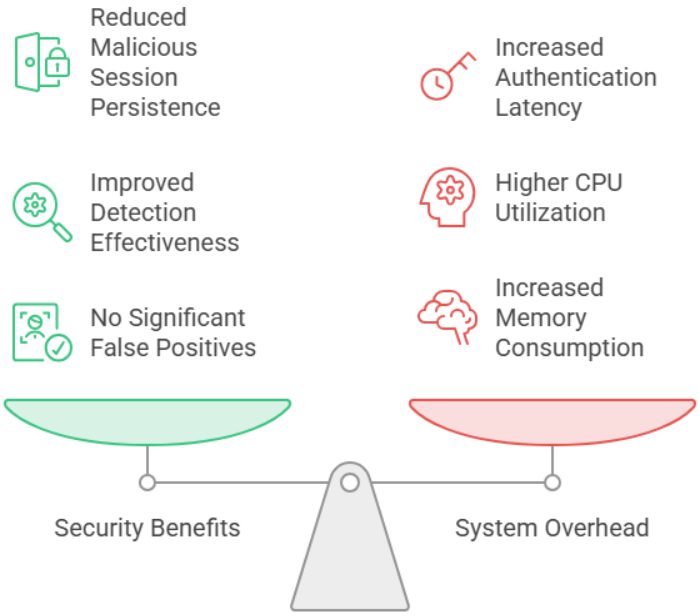

Figure 2. Balancing Security and System Overhead in Continuous Authentication

The operational results also reveal a measurable trade-off between stronger security and system overhead. Continuous authentication increased authentication latency, policy-decision time, CPU utilization, memory consumption, telemetry traffic, and the frequency of step-up verification. The increases remained relatively moderate when compared with the substantial reduction in malicious-session persistence and the improvement in detection effectiveness. False-positive behavior did not rise significantly, suggesting that the framework could achieve greater security sensitivity without automatically treating ordinary contextual changes as hostile activity. This balance is critical because continuous verification becomes operationally unsustainable if it protects systems at the cost of excessive disruption to legitimate users and services.

The case evidence further demonstrates that continuous authentication is applicable to both human and non-human identities. The token-hijacking case illustrated how an authenticated human session could become dangerous after credential verification, while the compromised-workload case showed that service identities and containerized workloads can also develop suspicious behavior despite using valid machine credentials. The framework responded to these scenarios through different graduated controls, including step-up authentication, privilege restriction, micro-segmentation, isolation, and session termination. This flexibility indicates that

continuous authentication should be viewed as an adaptive policy mechanism rather than merely as repeated user verification.

The findings are consistent with the broader Zero Trust literature that challenges the assumption that internal network location or prior authentication should confer durable trust. Existing Zero Trust approaches emphasize explicit verification, least-privilege access, resource-centered protection, identity awareness, and continuous assessment of security conditions. The present findings reinforce these principles by demonstrating experimentally that trust persistence after authentication creates a significant temporal vulnerability. The results extend conventional Zero Trust reasoning by showing that continuous verification is most effective when implemented as ongoing risk reassessment rather than as repeated execution of the same authentication procedure.

The results also correspond with previous continuous-authentication research that emphasizes behavioral and contextual signals as mechanisms for detecting identity misuse after login. Existing approaches frequently rely on behavioral biometrics, device attributes, geolocation, access patterns, or usage anomalies to determine whether an authenticated session remains consistent with expected behavior. The present study broadens this perspective by integrating identity, device, behavioral, workload, network, token, and resource-sensitivity indicators into one dynamic trust process. This integration is particularly relevant to cloud-native environments because authentication targets include not only human users but also APIs, containers, service accounts, workloads, and automated components.

The observed reduction in compromised-session exposure differs from security models that concentrate primarily on strengthening initial authentication through passwords, multi-factor authentication, certificates, or privileged identity controls. Strong initial authentication remains necessary, yet the token-hijacking scenario demonstrates that successful verification cannot guarantee the integrity of subsequent session activity. A valid token can be stolen, an endpoint can become compromised, and an authenticated user can behave maliciously after access has been granted. The findings therefore suggest that initial authentication and continuous authentication solve related but distinct security problems. Entry-point assurance confirms who or what should begin a session, while continuous assurance evaluates whether that trust remains justified over time.

The moderate performance overhead observed in this study also contributes to debates concerning the practical feasibility of continuous authentication. Some security approaches prioritize maximum detection sensitivity without examining the computational, latency, or usability consequences of continuous monitoring. The present results indicate that continuous authentication can remain operationally viable when verification is risk-adaptive and security responses are proportionate to observed conditions. Step-up authentication and micro-segmentation prevented the framework from relying exclusively on disruptive session termination. This finding differentiates adaptive Zero Trust enforcement from highly rigid models in which every anomaly results in immediate denial.

The findings signify that the network perimeter has lost much of its traditional value as a reliable boundary of trust. Cloud-native environments distribute applications, identities, services, and data across dynamic infrastructures that may change location, ownership, and operational context continuously. Security decisions based primarily on whether an entity has crossed an authentication boundary are therefore increasingly disconnected from the actual conditions of risk. The results indicate that trust itself must become dynamic, conditional, and revocable. Authentication should no longer be conceptualized as a gate that an entity passes once, but as an evidence-based relationship that must remain defensible throughout access.

The importance of post-authentication monitoring signifies a shift in the primary security question from “Was this identity authenticated correctly?” to “Does the current behavior still justify the privileges previously granted?” This shift has substantial implications for cybersecurity architecture. Valid credentials increasingly represent only one component of

identity assurance because attackers can operate through stolen tokens, compromised endpoints, hijacked sessions, or legitimate accounts. Continuous trust evaluation enables security systems to respond to changes that occur after authentication rather than assuming that the initial state remains valid. The findings therefore signal a transition from credential-centric security toward behavior- and context-aware access governance.

The successful detection of anomalous workload behavior signifies that Zero Trust must extend beyond human identity management. Modern cloud-native systems depend on machine identities that can outnumber human accounts and communicate continuously through APIs and service-to-service interactions. Service credentials can be abused in ways that resemble legitimate operations, making static authorization particularly vulnerable to lateral movement. The results indicate that workload behavior, communication patterns, requested resources, and runtime context should become part of authentication logic. Machine identity should therefore be treated as an active security subject whose trust level can change over time.

The effectiveness of graduated enforcement signifies that security does not need to operate through a simple permit-or-deny model. Risk evolves continuously, while not every anomaly represents an active compromise. A system that responds to every deviation with immediate termination may generate operational instability and user resistance. The proposed framework demonstrates that moderate-risk conditions can be addressed through additional verification, restricted privileges, or segmented access before stronger intervention becomes necessary. Adaptive enforcement therefore represents a more mature expression of Zero Trust because it combines skepticism with proportionality rather than equating distrust with indiscriminate denial.

The principal architectural implication is that continuous authentication should become a core component of Zero Trust implementation rather than an optional extension of identity and access management. Organizations migrating to cloud-native systems should avoid designing authentication around static sessions whose trust remains unchanged after login. Identity providers, policy engines, runtime monitoring systems, device-management platforms, and workload telemetry should be integrated so that changes in security conditions can modify access decisions dynamically. Such integration converts Zero Trust from a declarative policy principle into an operational control loop.

The identity-management implication concerns the need to expand authentication beyond human credentials. Cloud-native environments require unified approaches capable of evaluating employees, administrators, service accounts, workloads, containers, APIs, automated agents, and other non-human entities. Identity assurance should incorporate not only credential validity but also behavioral consistency, execution context, device or workload integrity, privilege use, and target-resource sensitivity. Organizations that apply sophisticated authentication to users while leaving service identities governed by long-lived static credentials may create significant blind spots. Continuous authentication must therefore evolve alongside machine-identity governance.

The security-operations implication is that telemetry becomes an active component of access control rather than merely a source for retrospective investigation. Runtime signals from endpoint detection, network monitoring, identity platforms, cloud workloads, API gateways, and policy engines can contribute directly to trust reassessment. Security operations centers should consequently integrate threat detection with real-time policy enforcement so that observed anomalies can reduce privileges or isolate sessions before manual investigation is completed. Such integration can shorten the interval between compromise, detection, and containment, which was one of the strongest improvements observed in the experimental results.

The usability implication is that Zero Trust implementation should prioritize adaptive friction rather than universal friction. Repeated authentication for every transaction can create fatigue, reduce productivity, and encourage insecure workarounds. The results suggest that legitimate low-risk activity can proceed with limited interruption when trust scores remain stable, while higher-risk activity receives stronger verification. This approach allows security intensity

to correspond to actual risk rather than treating all sessions identically. Effective Zero Trust design therefore requires simultaneous optimization of security assurance, operational continuity, and user experience.

The substantial improvement in malicious-session detection can be explained by the temporal nature of cyber risk. Authentication conditions are not static because users change locations, devices change posture, applications alter behavior, workloads communicate with new services, and tokens may be stolen after legitimate issuance. Conventional authentication captures only a snapshot of trust at the beginning of a session. Continuous authentication repeatedly updates that snapshot as new evidence becomes available. The framework consequently had more opportunities to identify compromise before malicious activity persisted long enough to cause significant damage.

The strong performance against token hijacking and compromised device scenarios can be explained by contextual inconsistency. Stolen credentials or tokens may remain cryptographically valid even when they are used from an unexpected device, location, behavioral pattern, or resource-access sequence. Static systems typically evaluate validity more strongly than contextual plausibility. The proposed framework reversed this imbalance by treating validity as necessary but insufficient evidence of trust. Contextual contradictions reduced the trust score even when the technical credential remained acceptable, allowing the system to identify forms of compromise that would otherwise appear legitimate.

The slightly lower detection performance for anomalous API behavior and lateral movement can be explained by the complexity of machine-generated activity. Cloud-native workloads regularly produce high-volume, dynamic, and automated communication that can vary according to scaling, deployment changes, service dependencies, or application demand. Malicious lateral movement may therefore resemble legitimate service interaction during its early stages. Human behavioral anomalies are often easier to interpret because expected access patterns are more stable and semantically meaningful. Machine-identity monitoring requires richer baselines and more context-sensitive models to distinguish legitimate operational variation from malicious deviation accurately.

The moderate computational overhead can be explained by the additional processing required to collect, normalize, correlate, and interpret continuous telemetry. Dynamic trust scoring introduces repeated calculations and policy decisions that conventional authentication performs less frequently. The overhead remained manageable because the framework applied risk-based escalation rather than repeating full authentication for every event. Lightweight monitoring was used during low-risk activity, while more expensive verification occurred only after risk increased. This tiered strategy explains how the system achieved stronger security without producing proportionally large performance degradation.

Future Zero Trust implementations should develop interoperable trust-scoring mechanisms capable of consuming telemetry from heterogeneous cloud platforms and security tools. Multi-cloud environments create substantial challenges because identity formats, logging structures, workload metadata, policy systems, and enforcement capabilities differ across providers. A continuous authentication framework will have limited strategic value if it functions only within a single proprietary ecosystem. Standardized trust attributes and policy interfaces could improve portability and allow organizations to maintain consistent risk decisions across distributed infrastructures.

Future research should investigate adaptive trust models that incorporate temporal learning and more sophisticated anomaly detection. Static thresholds may perform adequately in controlled environments but can become less effective as user behavior, workload patterns, and organizational processes change. Machine-learning approaches could assist in identifying subtle deviations, although they would introduce concerns regarding explainability, bias, adversarial manipulation, and model drift. Research should therefore compare interpretable rule-based systems, statistical models, and machine-learning approaches under equivalent threat conditions.

Detection performance should be evaluated alongside transparency and operational governability.

Future experimental work should also examine continuous authentication at substantially larger scale. The present framework can establish proof of feasibility, but production cloud environments may contain millions of authentication events and highly dynamic microservice relationships (Nurtas et al., 2025). Large-scale evaluation should measure telemetry volume, policy-engine throughput, failure recovery, distributed decision latency, and resilience under sustained attack. Testing should include realistic workload bursts, auto-scaling events, service migrations, hybrid-cloud connections, and infrastructure failures to determine whether continuous trust reassessment remains stable under operational pressure.

Future security governance should address the privacy and ethical implications of continuous monitoring (Lee et al., 2024). Behavioral authentication and contextual evaluation may require collection of detailed information about user activity, device status, location, communication patterns, and resource access. Excessive surveillance could create legal, organizational, and human-rights concerns even when implemented for legitimate security purposes. Continuous authentication should therefore follow principles of data minimization, purpose limitation, transparency, proportionality, retention control, and accountable access to telemetry. Zero Trust will be most sustainable when continuous verification strengthens security without transforming organizational environments into unnecessarily intrusive monitoring systems.

CONCLUSION

The most important finding of this study is that continuous authentication substantially strengthens Zero Trust Architecture by reducing reliance on persistent session-based trust in cloud-native environments. The proposed framework demonstrated stronger capability in detecting token hijacking, compromised device posture, privilege escalation, session replay, and anomalous workload behavior because trust was recalculated continuously using identity confidence, behavioral consistency, device condition, contextual information, workload activity, and resource sensitivity. The framework also shortened compromised-session exposure while maintaining a manageable level of computational and authentication overhead. These findings show that successful authentication at login should not be interpreted as durable evidence of trust. Security decisions in cloud-native environments need to reflect changing conditions throughout the entire session lifecycle, particularly when valid credentials or machine identities can be exploited after initial authorization.

The principal contribution of this research is conceptual and architectural, supported by an experimental evaluation framework. The study reconceptualizes authentication as a dynamic and revocable trust process rather than a discrete verification event and extends continuous authentication beyond human users to include service accounts, containers, workloads, APIs, and other machine identities. The proposed architecture also introduces graduated enforcement by mapping changing risk levels to proportionate responses such as passive monitoring, step-up authentication, privilege reduction, micro-segmentation, isolation, and session termination. Methodologically, the comparison between conventional and continuous authentication using controlled adversarial scenarios, detection metrics, exposure time, latency, false-positive behavior, and computational overhead provides a systematic basis for evaluating both security effectiveness and operational feasibility. This contribution helps translate the Zero Trust principle of continuous verification into an actionable control mechanism for cloud-native security environments.

The study is limited by its controlled experimental environment, predefined threat scenarios, restricted workload diversity, and reliance on simulated cloud-native sessions rather than sustained production-scale deployment. Real-world infrastructures may involve

substantially greater telemetry volumes, changing behavioral baselines, multi-cloud interoperability challenges, legacy components, and adversarial techniques that were not represented in the current evaluation. Future research should validate the framework across large-scale hybrid and multi-cloud environments, incorporate longitudinal behavioral baselines, and examine adaptive trust models capable of responding to evolving user and workload patterns. Further studies should also compare rule-based, statistical, and machine-learning approaches to trust scoring while evaluating explainability, adversarial robustness, privacy, and regulatory compliance. Such research would strengthen the development of continuous authentication as a scalable, transparent, and operationally sustainable foundation for Zero Trust security in increasingly decentralized computing environments.

DECLARATION OF AI AND AI ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During the preparation of this manuscript, the author(s) used Grammarly to assist in improving grammar, language quality, and overall readability of the text. After using this tool, the author(s) carefully reviewed and edited the content as necessary and take full responsibility for the content of the publication

AUTHOR CONTRIBUTIONS

Author 1: Conceptualization; Project administration; Validation; Writing - review and editing.

Author 2: Conceptualization; Data curation; In-vestigation.

Author 3: Data curation; Investigation.

DECLARATION OF COMPETING INTEREST

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

REFERENCES

Abdelaziz, O. M., & Aslan, H. K. (2025). A Unified Security Operations Center and Zero Trust Architecture Framework for Adaptive IoT Threat Mitigation. *2025 International Mobile, Intelligent, and Ubiquitous Computing Conference (MIUCC)*, 1–8. https://doi.org/10.1109/MIUCC66482.2025.11196839

Aftab, S. (2026). AI-Augmented Zero Trust Security Framework for Conversational AI Systems: A Behavioral Analytics Approach. *2026 International Conference on Electrical/Electronics, Robotics, Artificial Intelligence, and Informatics (ICERAI)*, 1–6. https://doi.org/10.1109/ICERAI69511.2026.11494527

Alang, K., Peta, S. B., Pai, R. R., & Patil, B. (2025). Scalable Cloud Architectures for Efficient Processing of Multi-Structured Big Data. *2025 Global Conference in Emerging Technology (GINOTECH)*, 1–6. https://doi.org/10.1109/GINOTECH63460.2025.11077101

Alapati, N. K., & Dhanasekaran, S. (2025). Addressing Data Quality and Consistency Issues in Cloud-Based Big Data Environments. *2025 International Conference on Networks and Cryptology (NETCRYPT)*, 458–462. https://doi.org/10.1109/NETCRYPT65877.2025.11102213

Ali, B., & Dib, O. A. (2026). The Next Frontier of Cybersecurity: Zero Trust for Enterprise IoT Ecosystems. In M. Trabelsi, M. Khan, Z. Bouida, & M. Murugappan (Eds.), *Proceedings of the 2nd Symposium on Smart, Sustainable, and Secure Internet of Things* (Vol. 1513, pp. 1–7). Springer Nature Singapore. https://doi.org/10.1007/978-981-95-5136-1_1

Alnaim, A. K., & Alwakeel, A. M. (2025). Zero Trust Strategies for Cyber-Physical Systems in 6G Networks. *Mathematics*, *13*(7), 1108. https://doi.org/10.3390/math13071108

Al-Said Ahmad, A., Al-Qora'n, L. F., & Zayed, A. (2024). Exploring the impact of chaos engineering with various user loads on cloud native applications: An exploratory empirical study. *Computing*, *106*(7), 2389–2425. https://doi.org/10.1007/s00607-024-01292-z

Alshehri, A. M., Atawneh, S. H., Al Bazar, H., & Mallouhy, R. E. (2026). Enhancing the Adoption of Zero Trust in Organizations Using Machine Learning. *Future Internet*, *18*(6), 278. https://doi.org/10.3390/fi18060278

Ben-Shimol, L., Lavi, D., Klevansky, E., Brodt, O., Mimran, D., Elovici, Y., & Shabtai, A. (2025). Detection of compromised functions in a serverless cloud environment. *Computers & Security*, *150*, 104261. https://doi.org/10.1016/j.cose.2024.104261

Camargo, J. S., Rezazadeh, F., Chergui, H., Siddiqui, S., & Liu, L. (2025). Towards Cloud-Native Agentic Protocol Learning for Conflict-Free 6G: A Case Study on Inter-Slice Resource Allocation. *2025 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*, 43–48. https://doi.org/10.1109/EuCNC/6GSummit63408.2025.11036996

Chen, B., & Ge, W. (2024). Design and optimization strategy of electricity marketing information system supported by cloud computing platform. *Energy Informatics*, *7*(1), 67. https://doi.org/10.1186/s42162-024-00366-8

Choi, H.-J., Kim, J.-H., Lee, J.-H., Han, J.-Y., & Kim, W.-S. (2025). Adaptive Microservice Architecture and Service Orchestration Considering Resource Balance to Support Multi-User Cloud VR. *Electronics*, *14*(7), 1249. https://doi.org/10.3390/electronics14071249

Desai, V., H N, S., U, S., Kumar, K. A., C, B., & U, C. (2026). Next-Gen Secure Access Using Intelligent Zero Trust Network Architecture. *2026 IEEE Global Symposium on Emerging and Communication Technologies (GSEACT)*, 1–6. https://doi.org/10.1109/GSEACT68539.2026.11620242

Devarajulu, V. S., Kanipakam, S., & Kavarakuntla, T. (2025). Explainable Anomaly Detection in Cloud-Native Systems via RAG-Enhanced Vector Retrieval. *2025 6th International Conference on Information Science, Parallel and Distributed Systems (ISPDS)*, 181–185. https://doi.org/10.1109/ISPDS67367.2025.11391017

Dube, A., Mpande, B., Dzehonye, F., Chazuza, T., & Ndlovu, B. (2026). Zero Trust Architecture: Redefining Security in the Digital World. In V. Bhateja, J. Anguera, N. Mostafa, & A. Ghosh (Eds.), *Advances in Micro-Electronics, Embedded Systems and IoT* (Vol. 1600, pp. 26–39). Springer Nature Switzerland. https://doi.org/10.1007/978-3-032-20533-9_3

Farhan, M., Rajapakshe, M., Bin Sulaiman, R., Aljaidi, M., Butt, U., Ahmad, M., Hadi, W., & Alshammari, A. A. (2026). ZTF-SF: A Zero-Trust Blockchain-Enabled IIoT Framework for Secure Smart Factories. *2026 2nd International Conference on Computational Intelligence Approaches and Applications (ICCIAA)*, 1–9. https://doi.org/10.1109/ICCIAA68481.2026.11543806

Fayaz, K., Chakradhar, V., & D, U. (2026). API Security with JWT Authentication and Reputation Based Monitoring System. *2026 Fifth International Conference on Power, Control and Computing Technologies (ICPC2T)*, 746–751. https://doi.org/10.1109/ICPC2T68221.2026.11646386

Galij, S., Pawlak, G., & Grzyb, S. (2025). Evaluating the Performance Impact of Data Sovereignty Features on Data Spaces. *Applied Sciences*, *15*(17), 9841. https://doi.org/10.3390/app15179841

Ganesh, C., Garg, R., Sreenivas, N., Saikia, H., Abdulhasan, M. M., & Dhanraj, J. A. (2025). Design and Implementation of AI-Based Decision Support Systems for Enhancing Strategic Business Operations. *2025 IEEE International Conference on Innovate for Humanitarian: Tech Solutions for Global Challenge (ICIH)*, 1–6. https://doi.org/10.1109/ICIH67754.2025.11608700

Guha, D., Saw, S. K., Mukherjee, P., Singh, G., & Satyam. (2025). A Practical Implementation of Zero Trust Architecture for Securing Web-Based Applications. *2025 International Conference on Advancements in Smart, Secure and Intelligent Computing (ASSIC)*, 1–6. https://doi.org/10.1109/ASSIC64892.2025.11158196

Ishaque, M., Albatati, H., & Nofal, M. (2025). AI-Based Zero Trust Architecture for Cognitive City Networks. *2025 2nd International Conference on Advanced Innovations in Smart Cities (ICAISC)*, 01–14. https://doi.org/10.1109/ICAISC64594.2025.10959378

Jayasri, T., Jenis, M. R. A., Aswathy, P. B., Manoranjitham, S., George, C., & Senthilkumar, R. (2025). Using Context-Aware and Identity-First Defense in Zero Trust Security Architecture to Protect Cyberspace. *2025 3rd International Conference on Intelligent Cyber Physical Systems and Internet of Things (ICoICI)*, 351–356. https://doi.org/10.1109/ICoICI65217.2025.11254505

K, R., Sulthana, S., Krishna, R. P., Naik, S. G., & Kavya. (2026). Redefining Trust Boundaries: A Holistic Synthesis of Zero-Trust Architecture Across Emerging Ecosystems. *2026 IEEE International Conference on Emerging Synergy Science and Technology (ICESST)*, 1–5. https://doi.org/10.1109/ICESST69086.2026.11582862

Kande, R. (2026). Zero-Trust Healthcare Clouds: AI-Enhanced Identity, Threat Detection, and Automated Compliance Enforcement. *2026 Third International Conference on Networking and Communications (ICNWC)*, 1–7. https://doi.org/10.1109/ICNWC68145.2026.11518125

Katipoglu, G., Utku, S., Mijailović, I., Mekić, E., Avdić, D., & Milić, P. (2024). Action Research Approach to Analysis of Teaching of Blockchain Web 3.0 Application Based on MACH Architecture. *Applied Sciences*, *14*(23), 11158. https://doi.org/10.3390/app142311158

Lee, D., Park, S., & Lee, B. (2024). Dynamic Traffic Load Rebalancing for Hardware-accelerated 6G UPF Resilient Architecture. *2024 IEEE International Performance, Computing, and Communications Conference (IPCCC)*, 1–7. https://doi.org/10.1109/IPCCC59868.2024.10850096

Lingam, R., Nagi, S., Chigurupati, M., & Myneni, B. T. (2026). Resilient DevSecOps: Self-Healing Cloud-Native Systems via SRE-Driven AI Threat Detection and Response. *2026 International Conference on Smart Futuristic Technology*, 1–6. https://doi.org/10.1109/ICSFT66733.2026.11508049

Malamuthu, B. K., Pandi, V. S., D, S., Jaber, A. H., Giri, J., & P, R. Y. (2025). Examining Multi-Cloud Architectures to Provide Enterprises with Resilient, Scalable, and Economical

Cloud Computing Solutions. *2025 International Conference on Engineering Innovations and Technologies (ICoEIT)*, 970–975. https://doi.org/10.1109/ICoEIT63558.2025.11211532

Mittal, G., & Ghosh, S. (2025). Trust for Critical Power Infrastructure Cybersecurity Framework for Modern OT Environments. *2025 International Conference on Intelligent Digitization of Systems and Services (IDSS)*, 129–134. https://doi.org/10.1109/IDSS67199.2025.11398188

Montin, E., Nguyen, X. T., & Lattanzi, R. (2026). MR Optimum: A web-based open-source tool for standardized signal-to-noise ratio evaluation in MRI. *Computer Methods and Programs in Biomedicine Update*, *9*, 100235. https://doi.org/10.1016/j.cmpbup.2026.100235

Nurtas, M., Nurakynov, S., Altaibek, A., Mergembayeva, A., & Mohammed, M. A. (2025). Satellite based deep learning approaches for detecting environmental disasters across Kazakhstan. *Discover Applied Sciences*, *8*(2), 144. https://doi.org/10.1007/s42452-025-08133-4

Paya, A., Vicente-García, & Gómez, A. (2025). Enhancing software-defined perimeters with integrated identity solutions and threat detection for robust zero trust security. *International Journal of Information Security*, *24*(4), 178. https://doi.org/10.1007/s10207-025-01099-9

Pule, B., Nleya, B., & Sibiya, K. (2026). A Zero Trust Driven Federative Learning Algorithm for Privacy Enhancement. *Applied Sciences*, *16*(8), 3872. https://doi.org/10.3390/app16083872

R, H., Marothia, F., Nag, S., Gangrade, D., Govindasamy, D., & P, P. (2025). A Comprehensive Survey on Zero Trust Architecture in IoT Networks. *2025 IEEE International Conference on Internet of Things and Intelligence Systems (IoTaIS)*, 90–96. https://doi.org/10.1109/IoTaIS67227.2025.11282130

Raju, V., Ujang, A. H., An, Y. D., Bellamkonda, S., Yeruva, L. A., & Chawla, N. (2026). Secure-by-Design Cloud Architectures: Integrating Machine Learning with Zero Trust Security Models. *2026 Innovations in Machine, Engineering, and Digital Conference (IMED)*, 1–6. https://doi.org/10.1109/IMED68921.2026.11484318

Rana, R., & Bhambri, P. (2026). Zero-Trust for Immersive Systems and Spatial Computing: In O. Almomani, A. Almomani, & M. Alauthman (Eds.), *Advances in Computational Intelligence and Robotics* (pp. 279–340). IGI Global Scientific Publishing. https://doi.org/10.4018/979-8-3373-7882-4.ch010

Sable, N. M., Sharma, V. K., & Manjre, B. (2025). Design of an Iterative Method for Zero-Day Attack Detection Using Adversarial Graph Temporal Convolutional Networks and Federated Learning. In A. Bagwari, J. L. V. Barbosa, E. Babulak, & D. S. Chauhan (Eds.), *Emerging Wireless Technologies and Sciences* (Vol. 2399, pp. 49–63). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-87886-2_4

Shin, D., Kim, J., Pawana, I. W. A. J., & You, I. (2025). Enhancing cloud-native DevSecOps: A Zero Trust approach for the financial sector. *Computer Standards & Interfaces*, *93*, 103975. https://doi.org/10.1016/j.csi.2025.103975

Singh, A. (2025). Enhancing 5G Security Through Zero Trust Architecture: A Threat-Centric Analysis of Critical Interfaces. *2025 IEEE 4th International Conference on Intelligent Reality (ICIR)*, 1–6. https://doi.org/10.1109/ICIR68135.2025.11361566

Singh, S. (2025). Zero-Trust Interconnect Framework for Secure Multi-Cloud Connectivity. *2025 International Conference on Computational Engineering, Sensing Technology and Management (ICCETM)*, 1–6. https://doi.org/10.1109/ICCETM66557.2025.11557841

Somayajula, R., Pai, R. R., Sajanraj, N., & Shah, K. (2025). Zero-ETL Architectures for AI Workloads Direct ML Model Access on Cloud-Native OLAP Systems. *2025 International Conference on Computing Technologies (ICOCT)*, 1–7. https://doi.org/10.1109/ICOCT64433.2025.11118928

Sultana, N., Miah, M. K., Ahmed, K. R., Goyal, S. K., Ali, A. H. M. M., & Karim, M. R. (2026). Zero-Trust Access Control Via Blockchain-Backed Identity and AI-Based Anomaly Detection. *2026 International Conference on Signal Analysis for Smart Systems (SIGNASS)*, 1–6. https://doi.org/10.1109/SIGNASS67826.2026.11480024

Vikas, Wahi, C., Sagar, B. B., & Manjul, M. (2025). Zero Trust for Secure Wireless Sensor Networks. *2025 International Conference on Networks and Cryptology (NETCRYPT)*, 42–46. https://doi.org/10.1109/NETCRYPT65877.2025.11102154

Viswanathan, J., N, D. Kumar., & Kumar, S. U. (2024). Zero Trust Security for Web Applications in Microservice-Based Environments. *2024 First International Conference on Data, Computation and Communication (ICDCC)*, 488–494. https://doi.org/10.1109/ICDCC62744.2024.10960955

Vusumuzi, M., & Godwin, M. (2025). AI-Driven Zero-Trust Models for Blockchain-Supported Healthcare Ecosystems. *Proceedings of the 2025 International Conference on Artificial Intelligence and Its Applications*, 1–11. https://doi.org/10.1145/3774791.3774801

Yashu, F., Rochester, S. M., & Saqib, M. (2025). Optimizing Intra-Container Communication with Memory Protection Keys: A Novel Approach to Secure and Efficient Microservice Interaction. *2025 International Conference on Business Intelligence for Technology Innovation (ICBITI)*, 1–7. https://doi.org/10.1109/ICBITI65527.2025.11501025

Yue, P., Wang, K., Wu, H., & Liu, R. (2025). Recent progress and technical practices of GIServices. *Chinese Science Bulletin*, *70*(30), 5163–5179. https://doi.org/10.1360/TB-2025-0107

Zeydan, E., Arslan, S., & Turk, Y. (2026). A Cloud Native Journey for Telecommunication Networks: Components, Applications and Open Challenges. *ACM Computing Surveys*, *58*(10), 1–38. https://doi.org/10.1145/3801492