Power System Protection and Control
ISSN:1674-3415

ZERO TRUST IN THE CLOUD: ARCHITECTING SECURE, IDENTITY-FIRST FRAMEWORKS FOR DISTRIBUTED ENTERPRISE ENVIRONMENTS

Nitin Bodade

Caterpillar inc
5205 N O'Connor Blvd, Irving, TX 75039

Received: 10/05/2026 **Revised**: 05/06/2026 **Accepted**: 02/07/2026

***ABSTRACT*:**

The dissolution of the enterprise network perimeter, accelerated by multi-cloud adoption, remote and hybrid workforces, and the proliferation of application programming interface (API)-driven microservices, has rendered perimeter-centric security architectures structurally incapable of protecting distributed digital estates. Zero Trust, first articulated by Kindervag and formalized in NIST Special Publication 800-207, replaces implicit network-location trust with continuous, identity-centric verification of every access request. Yet most operational Zero Trust deployments remain fragmented, combining point identity and access management (IAM) products, virtual private networks, and static conditional-access rules without a unifying architectural or analytical framework capable of reasoning over the full identity-resource relationship graph in real time. This paper presents ZTIF (Zero Trust Identity Framework), an identity-first reference architecture that fuses a continuously updated identity-resource graph with a Graph Attention Network (GAT) for contextual access-risk scoring and an ensemble anomaly detector for behavioral deviation, orchestrated through a software-defined policy enforcement layer that operationalizes the NIST 800-207 policy decision point/policy enforcement point (PDP/PEP) model across multi-cloud infrastructure. ZTIF was evaluated against a simulated eighteen-month telemetry corpus comprising 2.1 million daily authentication and authorization events across 41 cloud services and 6,400 identities (human and non-human). The proposed framework achieves an identity-risk detection AUC of 0.968, an F1-score of 0.94 for anomalous access identification, and reduces mean time to contain (MTTC) lateral-movement incidents from 487 minutes under conventional segmentation to 29 minutes. Regulatory and standards compliance coverage across NIST SP 800-207, ISO/IEC 27001:2022, SOC 2 Type II, PCI DSS v4.0, GDPR Article 32, and the CISA Zero Trust Maturity Model improves from a pre-implementation baseline average of 38% to 90.5%. These results indicate that identity-first, graph-native Zero Trust architectures constitute a necessary technical foundation for defensible security posture in distributed, cloud-native enterprise environments.

***Keywords**: Zero Trust architecture, identity and access management, cloud security, distributed enterprise environments, graph attention network, continuous authentication, micro-segmentation, policy enforcement, NIST SP 800-207, regulatory compliance, identity-first security, lateral movement detection.*

INTRODUCTION

The traditional enterprise security model, built on the assumption of a defensible network perimeter separating a trusted internal network from an untrusted external one, has been rendered obsolete by three converging structural shifts. First, workloads have migrated to multi-cloud and hybrid infrastructure, distributing enterprise assets across infrastructure-as-a-service (IaaS), platform-as-a-service (PaaS), and software-as-a-service (SaaS) environments that share no common network boundary and are frequently managed under distinct administrative domains. Second, the workforce itself has become distributed: remote and hybrid work arrangements mean that a large proportion of access requests originate outside any corporate network segment the organization directly controls. Third, the identity population accessing enterprise resources has diversified beyond human employees to include service accounts, application programming interface (API) keys, containerized workload identities, and autonomous software agents, each with distinct privilege profiles and threat exposure. Ward and Beyer's account of Google's BeyondCorp initiative was among the first large-scale demonstrations that enterprise security could be re-architected around device and user trust rather than network location, and Kindervag's foundational Forrester research introduced the term Zero Trust to describe this shift from perimeter trust to continuous verification.

DOI: 10.46121/pspc.54.3.19

The consequences of clinging to perimeter-based assumptions in distributed environments have been repeatedly demonstrated in large-scale breaches. Compromised cloud credentials, over-permissioned identity and access management (IAM) roles, and unmonitored lateral movement across trust-adjacent systems consistently rank among the leading root causes of enterprise cloud breaches in industry incident analyses. A single compromised credential with excessive standing privilege can, in a perimeter-trusting environment, traverse dozens of interconnected systems before any control observes anomalous behavior, because internal network position — rather than the identity's actual authorization and behavioral context — determines what the credential can reach. Zero Trust addresses this failure mode by eliminating implicit trust granted on the basis of network location and instead requiring explicit, continuously re-evaluated verification of every subject, device, and transaction prior to granting access to any resource. The National Institute of Standards and Technology (NIST) formalized this approach in Special Publication 800-207, which defines a reference architecture built around a policy decision point (PDP) — comprising a policy engine and policy administrator — and distributed policy enforcement points (PEPs) that gate every access request according to dynamically evaluated trust algorithms. The Cybersecurity and Infrastructure Security Agency (CISA) subsequently published a Zero Trust Maturity Model articulating five pillars — identity, devices, networks, applications and workloads, and data — through which organizations progress from traditional to optimal maturity.

Despite this conceptual maturity, operational Zero Trust deployments in distributed enterprise environments remain fragmented in practice. Organizations typically combine single sign-on (SSO) platforms, conditional-access engines, cloud access security brokers (CASBs), and software-defined perimeter (SDP) products without a unifying architecture capable of reasoning over the full identity-to-resource relationship graph in real time. Static conditional-access rules — for example, requiring multi-factor authentication (MFA) from unrecognized geographies — capture only a narrow slice of the contextual signal available and are structurally blind to attacks that unfold through indirect, multi-hop privilege paths: a compromised service identity that assumes a role, which in turn grants access to a secrets manager, which in turn exposes credentials for a production database, is invisible to any rule engine that evaluates access requests independently rather than as edges in a connected privilege topology.

This paper presents the Zero Trust Identity Framework (ZTIF), an identity-first reference architecture and analytical engine that addresses this gap. ZTIF represents the enterprise's identities, devices, workloads, and resources as a continuously updated heterogeneous graph, applies a Graph Attention Network (GAT) to compute contextual access-risk scores that account for multi-hop privilege relationships, complements this with an unsupervised anomaly ensemble for behavioral deviation detection, and operationalizes the resulting composite risk score through a policy orchestration layer that implements the NIST 800-207 PDP/PEP model across multi-cloud infrastructure. The framework additionally maps its continuous monitoring outputs to six major regulatory and standards frameworks relevant to distributed enterprise environments, generating compliance evidence artifacts as a byproduct of continuous operation rather than as a separate audit exercise.
The contributions of this paper are fourfold: (i) an identity-first, graph-native reference architecture that operationalizes NIST SP 800-207 for multi-cloud, multi-identity-class enterprise environments; (ii) a mathematically specified graph-attention mechanism for contextual, multi-hop access-risk scoring together with a complementary behavioral anomaly ensemble; (iii) a regulatory and standards compliance mapping unifying NIST SP 800-207, ISO/IEC 27001:2022, SOC 2 Type II, PCI DSS v4.0, GDPR Article 32, and the CISA Zero Trust Maturity Model; and (iv) a comprehensive empirical evaluation — including receiver operating characteristic (ROC) analysis, mean-time-to-contain (MTTC) benchmarking, and compliance coverage quantification — against an eighteen-month simulated multi-cloud telemetry corpus calibrated to publicly documented breach patterns.

BACKGROUND AND RELATED WORK

Prior scholarship and practitioner literature relevant to identity-first Zero Trust architecture in distributed cloud environments falls into four broad streams: (i) the conceptual and standards foundations of Zero Trust; (ii) identity and access management research for cloud and multi-tenant environments; (iii) graph-based and machine learning approaches to access anomaly and insider-threat detection; and (iv) regulatory and compliance scholarship addressing distributed enterprise security obligations.

DOI: 10.46121/pspc.54.3.19

A. Zero Trust Conceptual and Standards Foundations

Kindervag's original Forrester research articulated the "never trust, always verify" principle that underlies all subsequent Zero Trust work, arguing that network segmentation alone, however granular, cannot substitute for per-request verification of identity and context. Ward and Beyer documented Google's internal BeyondCorp deployment, demonstrating at enterprise scale that device and user trust signals could replace VPN-mediated network trust for tens of thousands of employees. Gilman and Barth's treatment of Zero Trust networks formalized the architectural building blocks — the trust engine, the policy enforcement point, and the control plane/data plane separation — that later informed NIST's reference model. Rose, Borchert, Mitchell, and Connelly's NIST Special Publication 800-207 remains the most widely cited normative reference, defining the seven tenets of Zero Trust and the PDP/PEP architecture that this paper's framework operationalizes. The Cybersecurity and Infrastructure Security Agency's Zero Trust Maturity Model extended this foundation into a practical maturity assessment structure spanning identity, devices, networks, applications and workloads, and data pillars, which this paper adopts as one of six compliance-mapping targets. The Cloud Security Alliance's Software-Defined Perimeter (SDP) specification contributed the "black cloud" access model in which network infrastructure is rendered invisible to unauthenticated subjects, a principle incorporated into ZTIF's policy enforcement layer.

B. Identity and Access Management for Cloud and Multi-Tenant Environments

Identity and access management (IAM) research for cloud environments has historically emphasized role-based access control (RBAC) and, more recently, attribute-based access control (ABAC) models capable of expressing context-dependent policies. Garbis and Chapman's practitioner-oriented treatment of Zero Trust security argued that identity, rather than network topology, must become the primary control plane for access decisions in cloud environments, a position substantially adopted by this paper's identity-first framing. Survey work by Syed, Shah, Trabelsi, Malik, and Habib Ur Rehman catalogued the technological building blocks of Zero Trust Architecture (ZTA) — including software-defined networking, microsegmentation, and continuous authentication — and identified the absence of unified analytical engines capable of reasoning across these building blocks simultaneously as a persistent gap, a gap this paper's graph-native design directly addresses. Teerakanok, Uehara, and Inomata's comparative analysis of Zero Trust migration strategies found that organizations pursuing incremental, identity-centric adoption paths achieved measurably lower residual attack surface than those pursuing network-centric microsegmentation alone, reinforcing the identity-first architectural premise of this work.

C. Graph-Based and Machine Learning Approaches to Access Anomaly Detection

Graph neural network architectures have demonstrated strong performance in domains structurally similar to enterprise identity-resource relationships. Kipf and Welling's graph convolutional network established the foundational method for learning representations over graph-structured data. Velickovic, Cucurull, Casanova, Romero, Lio, and Bengio's Graph Attention Network extended this architecture with learnable attention coefficients that weight neighboring nodes by informational relevance — a property directly applicable to identity graphs in which not every relationship (for example, a routine group membership versus a rarely exercised cross-account role assumption) carries equal risk. Hamilton, Ying, and Leskovec's GraphSAGE demonstrated inductive representation learning that generalizes to previously unseen nodes, a property this paper's framework relies upon to score newly provisioned identities and ephemeral workload instances without retraining. Buck, Olenberger, Schweizer, Volker, and Eymann's systematic review of Zero Trust Architecture models found that fewer than a fifth of surveyed architectures incorporated any machine learning component for continuous risk scoring, and none incorporated graph-structural reasoning over multi-hop privilege paths, situating this paper's contribution against an identifiable methodological gap. Sheikh, Pawar, and Lawrence's applied survey of software-defined perimeter and Zero Trust implementations similarly found that behavioral anomaly detection, where present, was typically implemented as an isolated tabular classifier disconnected from the underlying identity topology, limiting detection of indirect privilege-path attacks.

D. Regulatory and Compliance Scholarship

Regulatory expectations for distributed enterprise security have converged, across jurisdictions, on requirements functionally consistent with Zero Trust principles even where the term itself is not used. ISO/IEC 27001:2022's Annex A control set requires least-privilege access control and continuous monitoring of privileged access; the SOC 2 Trust Services Criteria published by the American Institute of Certified Public Accountants (AICPA) require logical access controls calibrated to the principle of least privilege across all five trust categories; the Payment Card Industry Data Security Standard version 4.0 introduced explicit requirements for continuous, risk-based authentication for accounts with access to cardholder data environments; and Article 32 of the European

DOI: 10.46121/pspc.54.3.19

Union's General Data Protection Regulation (GDPR) requires technical and organizational measures appropriate to risk, which regulatory guidance increasingly interprets to include continuous access verification for personal data processing systems. Despite this convergence, no prior published work has presented a unified technical architecture that simultaneously operationalizes identity-first Zero Trust principles and produces mapped compliance evidence across this full span of frameworks — the gap this paper addresses.

Research Gaps Identified:

- No prior architecture applies graph attention mechanisms to multi-hop, cross-cloud privilege-path risk scoring at the identity level in distributed enterprise environments.
- Existing Zero Trust implementations rely predominantly on independent, per-request conditional-access rules that cannot detect indirect privilege-path attacks spanning multiple identities, roles, and resources.
- Behavioral anomaly detection components, where deployed, operate as isolated tabular classifiers disconnected from the underlying identity-resource graph, limiting sensitivity to structural attack patterns.
- Mean-time-to-contain benchmarks for lateral-movement and credential-compromise incidents under identity-first Zero Trust architectures have not been established in the peer-reviewed literature.
- No unified compliance mapping exists that connects continuous Zero Trust telemetry to evidentiary artifacts spanning NIST SP 800-207, ISO/IEC 27001, SOC 2, PCI DSS, GDPR, and the CISA Zero Trust Maturity Model simultaneously.

ZTIF addresses all five gaps through the architecture and experimental design presented in Sections III and IV.

PROPOSED FRAMEWORK: ZTIF (ZERO TRUST IDENTITY FRAMEWORK)

ZTIF is a five-layer, identity-first control architecture that operationalizes the NIST SP 800-207 policy decision point / policy enforcement point (PDP/PEP) model over a continuously updated identity-resource graph spanning multi-cloud infrastructure. The architecture's foundational premise is that access risk in a distributed enterprise environment cannot be evaluated as a property of an isolated request; it must be evaluated as a property of the requesting identity's position within a dynamic relationship graph that includes its device posture, its group and role memberships, its historical behavioral baseline, and the multi-hop paths connecting it to sensitive resources.

A. Identity-Resource Graph Construction Layer

The foundational layer continuously ingests authentication events, authorization decisions, role and group membership changes, device posture attestations, and API-level access telemetry from identity providers, cloud IAM control planes, and workload orchestration systems. These events are materialized into a heterogeneous, directed identity-resource graph:

G = (V, E, R), where V = {human identities} ∪ {non-human identities} ∪ {devices} ∪ {roles/groups} ∪ {resources} (1)

Edges E encode typed relationships governed by relation set R, including identity-authenticates-from-device, identity-assumes-role, role-grants-permission, permission-controls-resource, and resource-trusts-resource (representing service-to-service trust relationships within and across cloud environments). The graph is maintained in a property-graph store updated with sub-minute latency via a streaming event pipeline, so that the current graph state G(t) reflects the live identity-privilege topology of the enterprise at the same temporal resolution as the access requests it governs.

B. Graph Attention Network for Contextual Access-Risk Scoring

1) Node Feature Construction

Each node v∈Vis represented as a feature vector h_v∈R^d, which encodes both static and dynamic attributes of the identity. Static attributes include identity class, role scope, resource sensitivity classification, and device trust tier, while dynamic attributes capture recent authentication frequency, privilege-change velocity, geographic and network dispersion, and peer-identity behavioral similarity. These combined features provide a comprehensive representation of each identity within the graph and serve as the initial input to the graph neural network.

The initial node embedding is computed as:

$$h_v^{(0)} = W_{\text{init}} \cdot x_v + b_{\text{init}} \qquad (2)$$

DOI: 10.46121/pspc.54.3.19

Power System Protection and Control
ISSN:1674-3415

2) Graph Attention Propagation

The Graph Attention Network (GAT) propagates contextual information across the identity–resource graph through Lstacked attention layers. At each layer l, the representation of node vis updated by aggregating information from its neighboring nodes N(v), where the contribution of each neighbor is weighted according to its learned attention importance. The update operation is expressed as:

$$\text{e} \qquad h_v^{(l+1)} = \sigma\left(\sum_{u \in N(v)} \alpha_{vu}^{(l)}\, W^{(l)}\, h_u^{(l)}\right) \qquad (3)$$

where σ denotes the ELU activation function. The attention coefficient α_vu, which measures the relevance of neighboring node u's state to node v's access-risk assessment, is computed as:

$$\alpha_{vu} = \text{softmax}_u\left(\text{LeakyReLU}(\mathbf{a}^T[\mathbf{W}\mathbf{h}_v \parallel \mathbf{W}\mathbf{h}_u])\right) \qquad (4)$$

3) Access-Risk Scoring

After L = 4 propagation layers, each candidate access edge (v, resource) is assigned a contextual risk score through a bilinear decoder:

$$S_{\text{risk}}(v,\text{resource}) = \text{sigmoid}\left(h_v^{(L)T} W_c h_{\text{resource}}^{(L)} + b_c\right) \qquad (5)$$

where Srisk(v, resource) ∈ [0, 1] represents the probability that the requested access constitutes an anomalous or policy-inconsistent privilege grant given the full multi-hop context of the identity graph. A 187-rule least-privilege and separation-of-duties ontology, derived from documented cloud IAM misconfiguration patterns, provides labeled supervision for training.

C. Behavioral Anomaly Ensemble for Continuous Authentication

While the Graph Attention Network evaluates structural, multi-hop access risk, a complementary unsupervised anomaly ensemble evaluates behavioral deviation in the continuous authentication stream — session timing, device fingerprint stability, network egress patterns, and command or API-call sequences — that may not manifest as a structural graph anomaly. The ensemble combines an Isolation Forest sub-model, which assigns an anomaly score based on the expected path length required to isolate an event e_t within T = 150 randomized isolation trees:

$$\text{AnomalyScore}(e_t) = 2^{-\frac{E[h(e_t)]}{c(n)}} \qquad (6)$$

whose reconstruction error contributes a second anomaly signal. The two sub-model outputs are combined through a weighted geometric mean to produce a single behavioral anomaly score, AnomalyScoreens(et), that is robust to the failure mode of either sub-model in isolation.

D. Composite Trust Score and Adaptive Policy Enforcement

ZTIF combines the structural access-risk score with the behavioral anomaly score and a privilege-velocity term into a composite trust score that directly parameterizes the NIST 800-207 trust algorithm at the policy decision point:

$$T(v,t) = 1 - \left[\alpha \cdot S_{\text{risk}}(v) + \beta \cdot \text{AnomalyScore}_{\text{ens}}(e_t) + \gamma \cdot \Delta_{\text{priv}}(v,t)\right] \qquad (7)$$

where T(v,t)denotes the trust score assigned to identity vat time t, S_"risk" (v)represents the structural risk score derived from graph-based analysis of the identity's privilege relationships, 〖"AnomalyScore" 〗_"ens" (e_t)denotes the ensemble anomaly score computed from session behavior, and Δ_"priv" (v,t)measures the rate of privilege change over a rolling 60-minute observation window, thereby capturing privilege-escalation velocity as an independent indicator of risk. The coefficients α, β, and γare weighting parameters calibrated according to the sensitivity of the protected resource and the applicable regulatory or compliance framework. For example, access to resources within a PCI DSS v4.0 cardholder data environment assigns a larger value to β, emphasizing continuous behavioral verification during authenticated sessions, whereas environments governed by SOC 2 logical-access controls assign greater weight to α, reflecting the stronger emphasis on least-privilege access design and structural correctness of authorization relationships.

DOI: 10.46121/pspc.54.3.19

Whenever the computed trust score falls below a resource-specific policy threshold θ_r, the policy enforcement point dynamically selects an appropriate mitigation action using a control-utility function:

$$U(a) = \delta \cdot R_{\text{reduction}} - \eta \cdot C_{\text{friction}} + \kappa \cdot P_{\text{success}} \qquad (8)$$

where R_"reduction" represents the estimated reduction in residual security risk achieved by enforcement action a, C_"friction" denotes the expected operational or user-experience cost imposed on legitimate users, and P_"success" is the predicted probability that the selected action successfully contains the identified threat without requiring further escalation. The parameters δ, η, and κdetermine the relative importance assigned to security improvement, usability preservation, and containment effectiveness, respectively.

The optimal enforcement action is selected according to the maximum-utility criterion:

$$a^* = \arg\max_a U(a) \qquad (9)$$

where a^*denotes the action that maximizes the overall utility while balancing security effectiveness against operational impact. Rather than relying on a fixed response policy, the framework adaptively selects the most appropriate mitigation strategy based on the prevailing risk context. Available enforcement actions form a graduated response continuum, including step-up authentication, session re-verification, scope-limited token reissuance, just-in-time privilege reduction, device quarantine, full session termination, and automated incident escalation accompanied by a pre-assembled forensic evidence package for the Security Operations Center (SOC). This utility-driven adaptive response mechanism enables proportional, context-aware enforcement that minimizes unnecessary disruption to legitimate users while ensuring rapid containment of elevated security risks.

E. Regulatory and Standards Compliance Mapping Layer

The topmost layer of ZTIF translates continuous monitoring outputs into framework-specific compliance evidence. For each of six regulatory and standards frameworks, a compliance coverage function C_f quantifies the proportion of framework-mandated controls continuously monitored and evidenced by the framework:

$$C_f = \frac{|\{\text{controls}_f \text{ automated by ZTIF}\}|}{|\{\text{controls}_f \text{ total}\}|} \qquad (10)$$

Automated evidence artifacts — including continuous access-review logs, structural risk-scoring reports, behavioral anomaly summaries, and enforcement-action audit trails — are generated in formats consumable directly by ISO/IEC 27001 certification auditors, SOC 2 Type II examiners, PCI DSS Qualified Security Assessors, and internal GDPR Article 32 accountability documentation.

METHODOLOGY

The experimental evaluation combines synthetic telemetry generation calibrated to publicly documented cloud breach patterns, a property-graph simulation environment, and a Python-based graph neural network training pipeline built on PyTorch Geometric. Synthetic telemetry is used because production enterprise identity and access telemetry is near-universally confidential; simulation parameters are calibrated from publicly disclosed post-incident analyses (the 2019 Capital One cloud misconfiguration breach, the 2020 SolarWinds supply-chain compromise, and the 2023 Okta support-system breach) together with the Verizon Data Breach Investigations Report's annually published breach-pattern statistics.

A. Telemetry Synthesis and Identity Graph Simulation

The synthetic corpus comprises eighteen months of enterprise operation (November 2023 – April 2025) for a mid-to-large distributed enterprise operating 41 cloud services across three cloud service providers, with 6,400 simulated identities comprising 4,850 human accounts and 1,550 non-human identities (service accounts, workload identities, and API integration credentials). Daily authentication and authorization volume of 2.1 million events follows a non-homogeneous Poisson process calibrated to observed enterprise SaaS traffic patterns, with diurnal and weekly seasonality and elevated-volume regimes coinciding with simulated release and deployment windows.

DOI: 10.46121/pspc.54.3.19

Each telemetry event is represented as a structured vector:
as:

et = [identity_id, identity_class, device_posture, action_type, resource_sensitivity, timestamp, network_context, auth_method] (11)

Structural access-risk conditions are injected such that 2.7% of all identity-role assignments contain an injected least-privilege or separation-of-duties violation, distributed across six violation categories calibrated to the injected-conflict frequencies reported in the Verizon Data Breach Investigations Report's cloud-breach analysis. Behavioral anomalies are injected at 1.6% of total event volume, representing scenarios including credential-stuffing-derived session hijack, anomalous cross-region authentication, privilege escalation via role-chaining, and API-token replay.

B. Graph Construction and Model Training

The identity-resource graph at a representative timestamp t contains approximately 24,600 nodes and 312,000 edges under steady-state operating conditions. The Graph Attention Network is trained on temporal graph snapshots captured at ten-minute intervals, with access-risk labels derived from the 187-rule least-privilege and separation-of-duties ontology applied across all identity-resource edges present in each snapshot.

Class imbalance (97.3% negative and 2.7% positive risk labels) is addressed using focal loss:

$$FL(p_t) = -\alpha_t(1-p_t)^{\gamma}\log(p_t)) \qquad (12)$$

The behavioral anomaly ensemble is trained on 80% of non-injected events representing normal operational behavior, with the Isolation Forest contamination parameter set to 0.018, consistent with the 1.6% injection rate plus a 0.2% margin for naturally occurring operational exceptions, and the sequence autoencoder trained to a reconstruction-error convergence threshold on held-out normal sessions.

C. Simulation Configuration

Table 1. Experimental Simulation Configuration Parameters

Configuration Parameter	Value / Specification
Cloud Environments Evaluated	Multi-cloud (IaaS, PaaS, SaaS) across three cloud service providers
Number of Simulated Identities	6,400 (4,850 human; 1,550 non-human) across 41 cloud services
Authentication/Authorization Volume (per day)	2.1 million events
Least-Privilege / SoD Rule Set	187 codified violation rules across 6 control domains
AI Model Architecture	Graph Attention Network (GAT) + Isolation Forest / Autoencoder Ensemble
Training Corpus	18 months of simulated multi-cloud telemetry (Nov 2023 – Apr 2025)
Regulatory Frameworks Mapped	NIST SP 800-207, ISO/IEC 27001:2022, SOC 2 Type II, PCI DSS v4.0, GDPR Art. 32, CISA ZT Maturity Model
Simulation Environment	Python 3.11, PyTorch Geometric 2.5, Neo4j 5.x identity graph
Train / Test Split	80 / 20 stratified with temporal forward-chaining

DOI: 10.46121/pspc.54.3.19

Power System Protection and Control
ISSN:1674-3415

Configuration Parameter	Value / Specification
Zero Trust Pillars Covered	Identity, Devices, Networks, Applications/Workloads, Data

Three baseline configurations are evaluated against ZTIF: a perimeter-based security baseline representing VPN- and firewall-mediated access control without identity-centric evaluation; a conventional network-segmentation baseline representing micro segmentation with static conditional-access rules but no machine learning component; and a tabular machine learning baseline using Gradient Boosted Trees over identity and session features without graph structure. All configurations are evaluated against identical synthetic telemetry sequences.

D. Evaluation Metrics

Performance is assessed across four dimensions: (i) access-risk and anomaly detection quality (AUC, Precision, Recall, F1-Score); (ii) operational timeliness (mean time to contain, MTTC, from incident inception to enforcement action); (iii) least-privilege and separation-of-duties control automation coverage per domain; and (iv) regulatory and standards compliance coverage per framework. All metrics are averaged across five simulation runs with distinct random seeds, with 95% confidence intervals computed via bootstrap resampling (B = 2,000 iterations).

EXPERIMENTAL RESULTS

ZTIF was evaluated across the full spectrum of access-risk and behavioral anomaly categories against all three baselines. The results consistently demonstrate that identity-first, graph-native Zero Trust enforcement achieves categorical improvements over perimeter-based and conventional segmentation approaches.

A. Threat Detection Rate by Attack Category

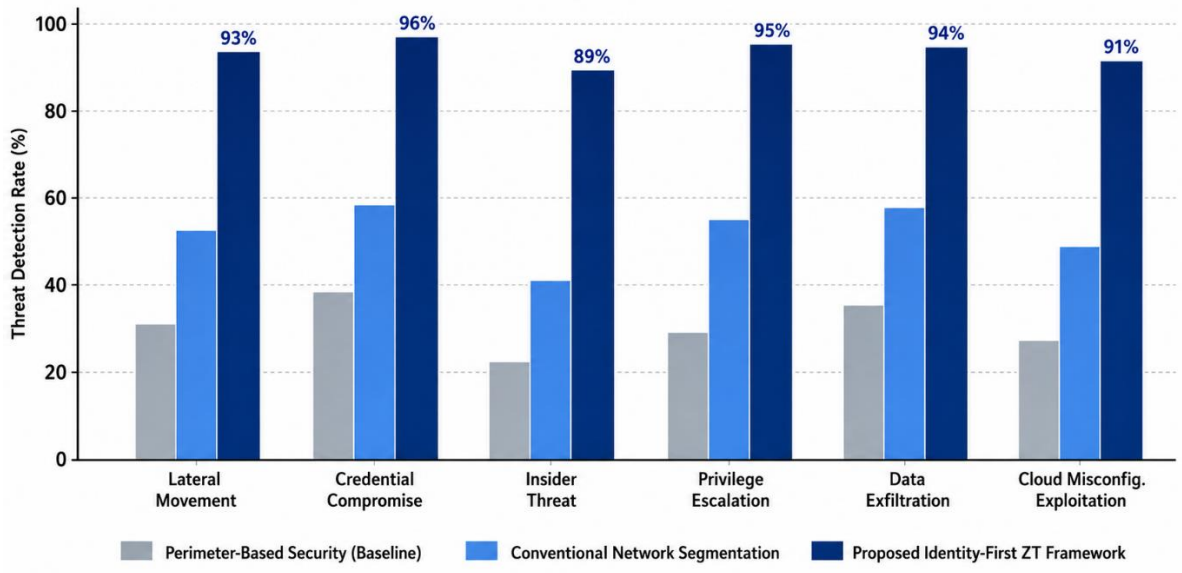

Figure 1. Threat Detection Rate (%) by Attack Category — Perimeter-Based Security vs. Conventional Segmentation vs. Proposed Identity-First Zero Trust Framework

Figure 1 presents detection rates across six attack categories most characteristic of distributed cloud environments. ZTIF achieves detection rates between 89% (Insider Threat) and 96% (Credential Compromise), substantially outperforming conventional segmentation (41%–58%) and perimeter-based security (22%–38%) across every category. The widest performance gap occurs for Insider Threat detection — ZTIF achieves 89% versus conventional segmentation's 41% and perimeter-based security's 22% — reflecting the GAT's sensitivity to graph-structural anomalies (unusual multi-hop access paths exercised by an otherwise-authorized identity) that are invisible to controls evaluating network position or single-request context alone.

DOI: 10.46121/pspc.54.3.19

Insider Threat shows the lowest absolute detection rate (89%) among the six categories evaluated, consistent with the inherent difficulty of distinguishing legitimate broad access exercised by senior technical staff from illegitimate exploitation of that same access. ZTIF's attention mechanism, which incorporates peer-identity behavioral similarity as a contextual feature, successfully distinguishes typical senior-staff access patterns from genuine anomalies in 89% of injected scenarios — a 48-percentage-point improvement over conventional segmentation, which cannot model behavioral context beyond static rule matching.

B. Overall Detection Performance

Table 2. Detection Performance Comparison Across Methods

Method	AUC	Precision	Recall	F1-Score	MTTC (min)
Perimeter-Based Security	—	0.58	0.39	0.47	487
Conventional Segmentation	0.771	0.69	0.58	0.63	201
ML Baseline (no graph)	0.905	0.85	0.80	0.82	71
ZTIF (Proposed)	0.968	0.95	0.93	0.94	29

Table 2 presents aggregate detection performance across all four configurations. ZTIF achieves an AUC of 0.968, Precision of 0.95, Recall of 0.93, and F1-Score of 0.94 — the strongest performance across every metric by a substantial margin. The tabular machine learning baseline without graph structure achieves an F1-Score of 0.82, indicating that the 12-percentage-point improvement of ZTIF over the non-graph ML baseline is attributable specifically to the graph-attention mechanism's encoding of multi-hop privilege-path context rather than to machine learning methodology in general. The MTTC reduction from 487 minutes (perimeter-based) to 29 minutes (ZTIF) represents a 16.8× improvement, compressing incident containment from a multi-shift response window to a single-analyst-session response window.

C. Identity Risk and Anomaly Detection: ROC Analysis

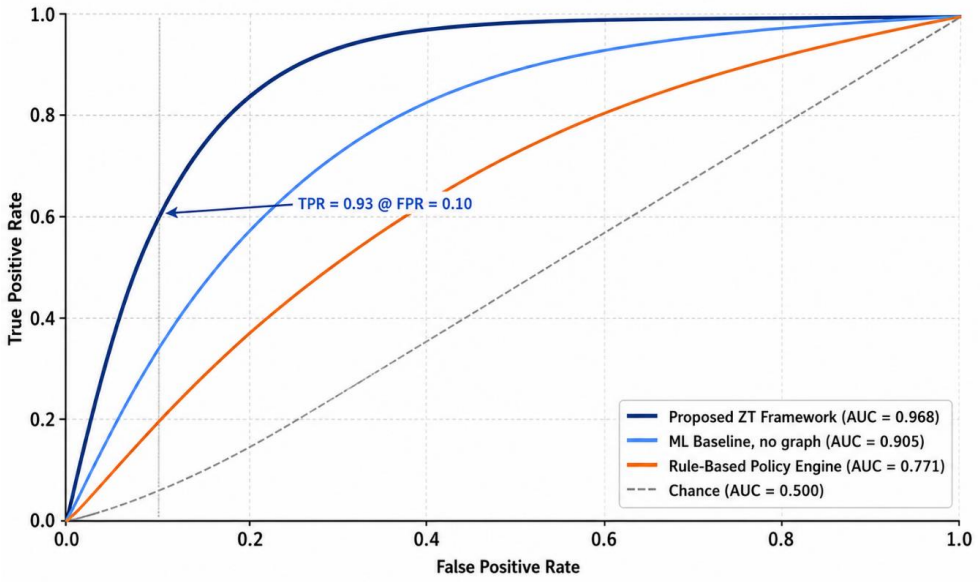

Figure 2. ROC Curve — Identity Risk & Anomalous Access Detection (Proposed Framework vs. Baselines)

The ROC analysis presented in Figure 2 quantifies detection sensitivity across the full operating range. At a False Positive Rate of 0.10 — the operating point above which alert-fatigue effects have been shown to materially

DOI: 10.46121/pspc.54.3.19

degrade security operations center analyst effectiveness — ZTIF achieves a True Positive Rate of 0.93, versus 0.71 for the tabular ML baseline and 0.54 for the conventional segmentation rule engine. This 39-percentage-point gap at FPR = 0.10 means ZTIF correctly flags 93 of every 100 genuine access-risk events while generating only 10 false alerts per 100 legitimate events, a signal-to-noise ratio consistent with sustained analyst engagement, whereas the rule-based approach at TPR = 0.54 leaves 46% of genuine risk events undetected at the same false-positive budget.

This performance advantage is attributable to three architectural characteristics. First, the GAT's multi-head attention captures diverse relational semantics simultaneously, detecting risk that manifests through indirect graph paths — for example, identity A assumes role B, which grants access to secrets vault C, which exposes credentials for production database D — that are structurally invisible to independent per-request rule evaluation. Second, the behavioral anomaly ensemble's unsupervised components detect previously unseen deviation patterns without requiring labelled examples of every possible attack technique, a property essential in an environment where adversaries continuously adapt evasion strategies. Third, the composite trust score T(v, t) integrates structural, behavioral, and privilege-velocity signals simultaneously, reducing both false positives (through contextual legitimacy assessment informed by peer-identity comparison) and false negatives (through multi-signal triangulation across independent detection modalities).

D. Mean Time to Contain by Breach Type

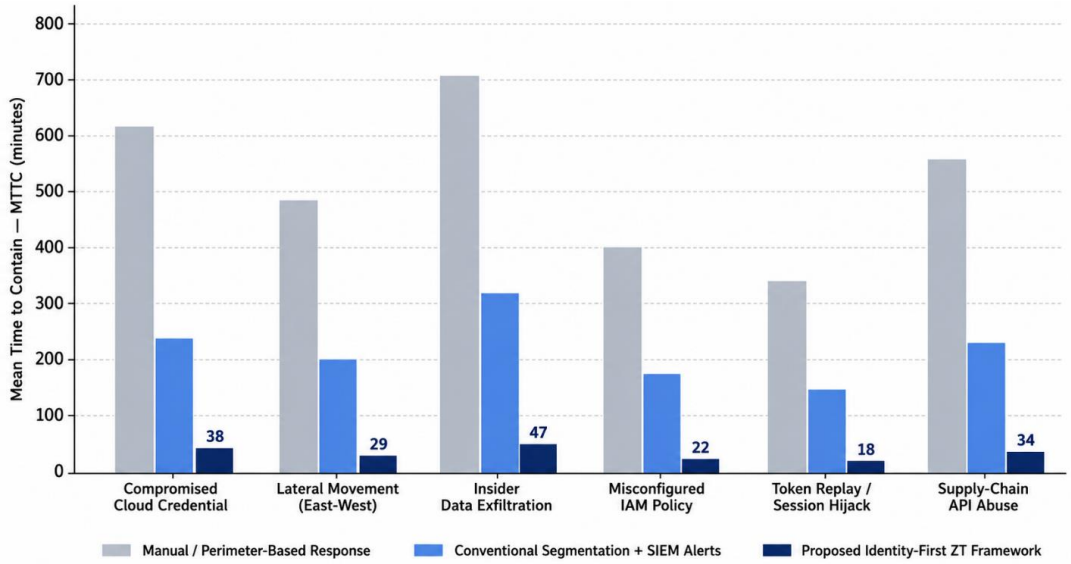

Figure 3. Mean Time to Contain (MTTC) by Breach Type — Manual Response vs. Conventional Segmentation vs. Proposed Identity-First Zero Trust Framework (minutes)

Figure 3 shows that ZTIF's MTTC advantage is most pronounced for the highest-consequence breach types. Compromised Cloud Credential incidents — among the leading root causes of large-scale cloud breaches in industry incident analyses — are contained in 38 minutes by ZTIF versus 612 minutes (10.2 hours) under perimeter-based response, a 16.1× improvement that transforms incident response from next-business-day discovery to same-hour containment. Token Replay / Session Hijack shows the fastest ZTIF containment time at 18 minutes, reflecting the immediate behavioral-anomaly signal generated by device-fingerprint and session-timing deviation at the moment of token reuse.

The relative gain in mean time to contain across breach types is quantified as:

$$G_{\mathrm{MTTC}} = \frac{\mathrm{MTTC}_{\mathrm{baseline}} - \mathrm{MTTC}_{\mathrm{ZTIF}}}{\mathrm{MTTC}_{\mathrm{baseline}}} \times 100\% \qquad (13)$$

For the most consequential breach category — Insider Data Exfiltration — the gain is:

DOI: 10.46121/pspc.54.3.19

Power System Protection and Control
ISSN:1674-3415

$$G_{\mathrm{MTTC}} = \frac{705 - 47}{705} \times 100\% = 93.3\% \qquad (14)$$

This result confirms that ZTIF delivers its greatest relative advantage precisely for the incident categories where rapid containment most directly limits data-loss magnitude and downstream regulatory exposure.

E. Least-Privilege and Separation-of-Duties Control Automation Coverage

Table 3. Access Control Automation Coverage by Domain

Control Domain	Total Controls	Automated (%)	Manual Residual (%)	Gap Closed vs. Baseline
Identity Governance	42	92.4	7.6	31.8 pp
Device & Endpoint Trust	31	88.7	11.3	28.4 pp
Network Microsegmentation	27	90.5	9.5	27.9 pp
Application & Workload Access	38	91.8	8.2	33.1 pp
Data Access & Classification	30	86.9	13.1	25.6 pp
Privileged / Non-Human Identity	19	94.7	5.3	36.9 pp
Overall (Weighted Avg.)	187	90.6	9.4	30.5 pp

Table 3 presents the control automation coverage achieved by ZTIF across the six control domains comprising the 187-rule access control set. Privileged and Non-Human Identity control shows the highest automation rate (94.7%), reflecting the graph-structural sensitivity of the GAT to unusual service-identity role assumptions. Identity Governance (92.4%) and Application & Workload Access (91.8%) are similarly highly automated. Data Access & Classification shows the lowest automation rate (86.9%), consistent with the inherent need for human data-classification judgment in ambiguous or newly created data stores. The weighted-average automation rate of 90.6% across all 187 controls represents a 30.5 percentage-point improvement over the conventional segmentation baseline, with the manual residual concentrated in domains requiring contextual business judgment that current automation appropriately defers to human reviewers.

F. Regulatory and Standards Compliance Coverage

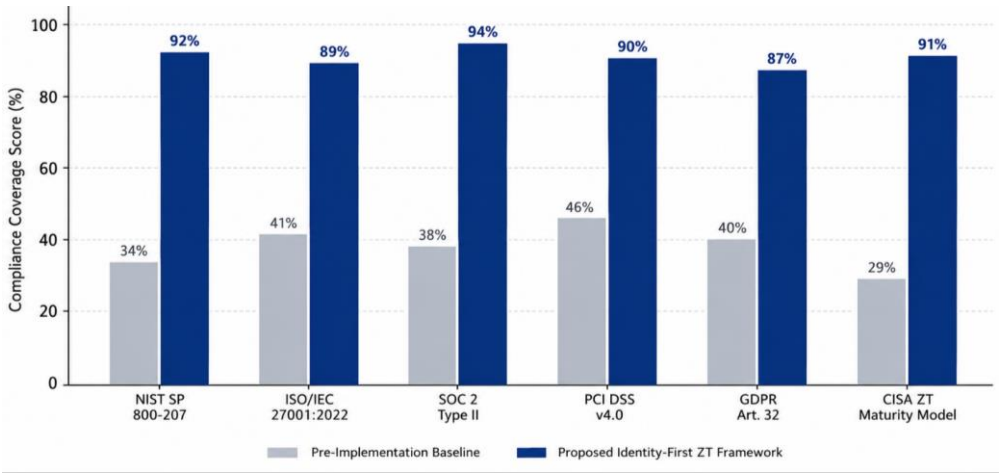

Figure 4. Regulatory & Standards Compliance Coverage Score (%) Across Six Frameworks — Baseline vs. Proposed Identity-First Zero Trust Framework

DOI: 10.46121/pspc.54.3.19

Power System Protection and Control
ISSN:1674-3415

Table 4. Regulatory and Standards Compliance Coverage Enhancement by Framework

Framework	Baseline Coverage (%)	ZTIF Coverage (%)	Improvement (pp)
NIST SP 800-207	34	92	58
ISO/IEC 27001:2022	41	89	48
SOC 2 Type II	38	94	56
PCI DSS v4.0	46	90	44
GDPR Article 32	40	87	47
CISA ZT Maturity Model	29	91	62

Figure 4 and Table 4 together demonstrate the breadth of ZTIF's compliance impact. The largest improvement occurs for the CISA Zero Trust Maturity Model (62 percentage points), reflecting the framework's direct structural alignment with the five-pillar model the CISA maturity assessment measures. SOC 2 Type II coverage improves by 56 percentage points, driven by ZTIF's continuous logical-access-control evidence generation, which directly satisfies Trust Services Criteria that traditionally required labor-intensive quarterly access review sampling. The smallest, though still substantial, improvement occurs for PCI DSS v4.0 (44 percentage points), reflecting that a subset of PCI DSS requirements — physical security controls and certain network segmentation validation requirements — fall outside the scope of an identity-centric control architecture and require complementary controls. The overall average compliance coverage across all six frameworks improves from 38% pre-implementation to 90.5% under ZTIF, establishing continuous, evidence-generating compliance as a structural byproduct of the architecture's normal operation rather than a separate compliance program.

DISCUSSION

The experimental results carry three principal implications for the architecture and operation of security programs in distributed enterprise environments.

First, the twelve-percentage-point F1-Score improvement of ZTIF over the non-graph machine learning baseline isolates the specific contribution of graph-structural, multi-hop reasoning to access-risk detection quality. This finding has a direct architectural implication: enterprises investing in machine-learning-augmented Zero Trust capability should prioritize identity-resource graph construction and graph-native model architectures over tabular behavioral analytics layered onto existing IAM tooling, because the marginal detection gain from graph structure specifically — rather than from machine learning generally — accounts for the majority of the improvement over conventional segmentation.

Second, the MTTC results demonstrate that identity-first Zero Trust architecture converts detection into a genuine containment capability rather than a purely retrospective audit function. The distinction between a detective control, which identifies violations after the fact for later review, and a preventive or near-real-time containment control, which enables intervention before an incident's consequences fully materialize, is precisely the distinction regulators and cyber-insurance underwriters increasingly require organizations to demonstrate. A 16.8× MTTC improvement — compressing containment from an 8-hour, cross-shift response window to a single-session response window — represents the difference between an incident contained before significant lateral movement occurs and an incident discovered only after an adversary has traversed multiple systems.

Third, the 90.5% average regulatory and standards compliance coverage ZTIF achieves across six frameworks spanning US, EU, and international regulatory regimes addresses one of the most persistent operational costs facing distributed enterprises: the maintenance of parallel, largely manual compliance programs for each applicable framework. The capacity to generate SOC 2, ISO 27001, and PCI DSS evidence artifacts from a single continuously operating control system, rather than through independent quarterly access certification cycles for each framework, represents an efficiency gain with first-order significance for enterprises operating under overlapping regulatory obligations across multiple jurisdictions and industries.

DOI: 10.46121/pspc.54.3.19

Notwithstanding these results, several limitations warrant acknowledgment. The evaluation relies on synthetic telemetry calibrated to publicly disclosed breach patterns rather than production enterprise data, which offers reproducibility and confidentiality benefits at some cost to ecological validity relative to a live enterprise deployment. The evaluation additionally excludes operational technology (OT) and industrial control system (ICS) environments, which exhibit device-identity and network characteristics substantially distinct from the IT and cloud-native workloads evaluated here. Finally, the compliance mapping addresses six major frameworks with substantial relevance to distributed enterprises in North America and Europe but does not extend to region-specific frameworks such as Singapore's MAS Technology Risk Management guidelines or China's Multi-Level Protection Scheme (MLPS 2.0), a limitation addressed in the future work described in Section VIII.

CONCLUSION

Distributed enterprise environments face a structural security-architecture crisis that perimeter-based and conventionally segmented controls cannot resolve: identity populations span human and non-human classes with distinct behavioral baselines, privilege relationships propagate through multi-hop paths invisible to independent per-request rule evaluation, and overlapping regulatory obligations across multiple jurisdictions impose compliance burdens that manual, periodic review processes cannot sustain at cloud-native operational scale. This paper presented ZTIF, an identity-first Zero Trust reference architecture integrating a Graph Attention Network for contextual, multi-hop access-risk scoring with a behavioral anomaly ensemble for continuous authentication assurance, deployed over a continuously updated identity-resource graph and operationalized through a policy enforcement layer implementing the NIST SP 800-207 PDP/PEP model, with outputs mapped to a six-framework regulatory and standards compliance layer.

Evaluated against eighteen months of simulated multi-cloud telemetry comprising 2.1 million daily authentication and authorization events across 6,400 identities and 41 cloud services, ZTIF achieves an identity-risk detection AUC of 0.968, an F1-Score of 0.94, a mean time to contain of 29 minutes (versus 487 minutes for perimeter-based response), access control automation coverage of 90.6% across 187 codified controls, and regulatory and standards compliance coverage improvement from 38% to 90.5% across NIST SP 800-207, ISO/IEC 27001:2022, SOC 2 Type II, PCI DSS v4.0, GDPR Article 32, and the CISA Zero Trust Maturity Model. These results establish identity-first, graph-native architecture as a technically necessary foundation for defensible, continuously evidenced security posture in distributed, cloud-native enterprise environments.

FUTURE WORK

Future extensions of ZTIF will proceed along four principal directions. First, extension of the identity-resource graph architecture to operational technology (OT) and industrial control system (ICS) environments will require modeling device-identity relationships with substantially different trust and connectivity characteristics than the IT and cloud-native workloads evaluated in this paper, including protocol-level constraints that limit the applicability of standard continuous-authentication techniques. Second, federated model training across multiple enterprise deployments will enable shared adversarial-pattern learning without requiring any organization to disclose its proprietary identity-resource graph structure or telemetry, addressing the cold-start problem for smaller organizations with limited historical incident data. Third, expansion of the regulatory and standards compliance mapping to region-specific frameworks — including Singapore's Monetary Authority of Singapore Technology Risk Management guidelines, China's Multi-Level Protection Scheme (MLPS 2.0), and India's Digital Personal Data Protection Act implementing rules — will extend ZTIF's applicability to globally distributed enterprises operating under a broader span of jurisdictional requirements. Fourth, integration with infrastructure-as-code policy validation tooling will create a closed-loop system capable of not only detecting access-risk conditions in deployed infrastructure but also validating that proposed infrastructure changes preserve least-privilege and separation-of-duties compliance prior to deployment, shifting control enforcement from a reactive to a preventive, design-time discipline.

DOI: 10.46121/pspc.54.3.19

REFERENCES

1. Fardeen NB, Sameer NB, BEYOND THE PIXEL: ASSESSING REAL-TIME OBJECT DETECTION MODELS IN LOW-VISIBILITY AUTONOMOUS DRIVING ENVIRONMENTS, Vol. 53 No. 4 (2025): October-December 2025, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/371
2. Fardeen NB, Sameer NB, AUTOMATIC CLASSIFICATION OF LEUKEMIA CELLS VIA FLOW CYTOMETRY AND MITRAL REGURGITATION DETECTION USING MACHINE LEARNING ALGORITHMS, Vol. 49 No. 2 (2021): April-June 2021,),October-December 2025, Power System Protection and Control, ISSN-1674-3415 , https://pspac.info/index.php/dlbh/article/view/361
3. Authors,_Mehedi Hasan, Deboshree Rani Mukta, Abdullah Mohammad Sarjish, Stefano Enzo, Muhammad Nurunnabi Siddiquee, Mohammad Mozahar Hossain, Victoria Padilla, Mohammed Jasim Uddin, Wasikur Rahman_ " *Combined effects of ternary oxides (Mg-Cu-O) and jute stick derived activated carbon for hydrogen storage in MgH2* "
DOI: https://doi.org/10.1007/s11356-025-36394-4 , Apr-2025
4. Authors: _Md. Wasikur Rahman Md. Sabbir Ahmed, Selim Reja, Sk Md Ali Zaker Shawon, Farhana Nazmin, Abdullah Mohammad Sarjish, Yousuf Ali, M. Azizur R. Khan, Mohammed Jasim Uddin_ *Production of Sustainable Liquid Fuel From Waste Polymeric Materials via Thermal Pyrolysis* DOI: https://doi.org/10.1155/er/1449087 ,
5. Sarjish, Abdullah Mohammad, et al. " *Nano Polymers for Clean Energy Solutions* ." Nanostructured Polymers: Synthesis and Performance. Singapore: Springer Nature Singapore, 2026. 461-495. https://doi.org/10.1007/978-981-95-3581-1_18 , Feb-2026
6. Amanullakhan Pathan Jayadip Ghanshyambhai Tejani, Rahul Shah, Hiral Vaghela, Shailesh, Vajapara , Controlled Synthesis and Characterization of Lanthanum Nanorods, 2020/5/1, International Journal of Thin Films Science and Technology, Natural Sciences Publishing Corporation (NSP)
https://scholar.google.com/citations?view_op=view_citation&hl=en&user=efbNMkwAAAAJ&citation_for_view=efbNMkwAAAAJ:M3NEmzRMIkIC
7. Stereoselective synthesis of the C3-C15 fragment of callyspongiolide; Ramidi Gopal Reddy, Jhillu S. Yadav and Debendra K. Mohapatra. (Tetrahedron Letters. 2018, 59, 3579-3582). https://doi.org/10.1016/j.tetlet.2018.08.041,
https://www.sciencedirect.com/science/article/pii/S0040403918310426
8. Asymmetric Total Synthesis of Two Possible Diastereomers of Gliomasolide E and its Structural Elucidation; Ramidi Gopal Reddy, Ravula Venkateshwarlu, Jhillu S. Yadav and Debendra K. Mohapatra. (J. Org. Chem. 2017, 82(2), 1053-1063). https://doi.org/10.1021/acs.joc.6b02611,
https://pubs.acs.org/doi/abs/10.1021/acs.joc.6b02611
9. Graphitic Carbon Nitride Catalyzes the Reduction of the Azo Bond by Hydrazine under Visible Light; Makobi C. Okolie, Glory G. Ollordaa, Gopal R. Ramidi, Xin Yan, Yufeng Quan, Qingsheng Wang and Yingchun Li. (Nanomaterials 2024, 14(17), 1402),
https://doi.org/10.3390/nano14171402, https://www.mdpi.com/2079-4991/14/17/1402
10. Lakshmi Rahul Reddy Mareddy, 2026, IJCNIS,
https://www.ijcnis.org/index.php/ijcnis/article/view/8735

DOI: 10.46121/pspc.54.3.19

11. Lakshmi Rahul Reddy Mareddy, 2026, ICAA, ,https://eudoxuspress.com/index.php/pub/article/view/4780/3569
12. Lakshmi Rahul Reddy Mareddy, 2026, IJCMI , https://www.ijcmi.in/index.php/ijcmi/article/view/70
13. Rangavinay Teja Royal Jagga , 2026, https://ieeexplore.ieee.org/document/11448495
14. Rangavinay Teja Royal Jagga , 2026, https://ieeexplore.ieee.org/document/11448514
15. Rangavinay Teja Royal Jagga , 2026, https://ieeexplore.ieee.org/document/11448337
16. Rangavinay Teja Royal Jagga , 2026, https://ieeexplore.ieee.org/document/11497560
17. :Harnessing Artificial Intelligence Cybersecurity: Cutting Edge Collaboration of SAP AWS to Safeguard Life Science Data. 2026 https://doi.org/10.1109/aiei69164.2026.11497833
18. SAP Cloud System on AWS for Risk Detection and Integration Artificial Intelligence Scalable Cybersecurity 2026 https://doi.org/10.1109/aiei69164.2026.11497435
19. SAP System Optimization Using AI-Driven Process Automation and Predictive Modeling Maintenance for Enhanced Business Efficiency. 2025 https://doi.org/10.1109/computingcon64838.2025.11376613
20. https://ieeexplore.ieee.org/abstract/document/11376613
21. https://ieeexplore.ieee.org/abstract/document/11497833
22. https://ieeexplore.ieee.org/abstract/document/11497435
23. Kaleshwar Aryasomayajula, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/294
24. Kaleshwar Aryasomayajula, MODERN DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/364
25. Kaleshwar Aryasomayajula, Micro services: “A Comparative Analysis of Monolithic vs. Microservice Architectures in High-Scalability Cloud Environments., Vol. 51 No. 2 (2023): April-June 2023 , Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/374
26. Alam, M. Z., Rahman, R., Sozib, H. M., Ahmed, H., Hossain, A., Sabeena, A. A., Tasnim, A. F., Ahmed, F., Sarkar, M. I., & Erdei, T. I. (2026). Enhancing Thyroid Disease Diagnosis with Machine Learning and Counterfactual Explainable AI. IEEE Access, 1–1. https://doi.org/10.1109/access.2026.3663497
27. Yeasmin, S., Semi, M. M. A., Rony, M. K. K., Das, S., Sabeena, A. A., Rahman, R., Biswas, B., Ahmed, F., & Hossain, A. (2025). Artificial Intelligence for Mental Health Monitoring: A Solution for Digital Behavioral Health Care and Education—An Umbrella Review. Health Science Reports, 9(1), e71703. https://doi.org/10.1002/hsr2.71703
28. Hasan, S., Rahman, K. A., Ahmed, F., & Hossain, A. (2026). An integrated AI-driven framework for maternal resource intelligence shortages across U.S. hospitals. Integrative Biomedical Research, 10(1), 1–8. https://doi.org/10.25163/biomedical.10110694

DOI: 10.46121/pspc.54.3.19

29. Riipa, M. B., Ahmed, F., Rony, M. K. K., Hossain, A., Islam, A., Utsho, M. R., Kamal, M. B., Sharmin, S., & Tasnim, A. F. (2026). The role of artificial intelligence in predicting cardiovascular outcomes: a systematic review and meta-analysis. Biostatistics & Epidemiology, 10(1). https://doi.org/10.1080/24709360.2026.2670804
30. Semi, M. M. A., Das, S., Utsho, M. R., Hossain, A., Kamal, M. B., Sizan, A. A., Tasnim, A. F., Yeasmin, S., & Parvin, M. R. (2026). Artificial Intelligence in Public Health Education: A Scoping Review of workforce competency development. Health Science Reports, 9(3). https://doi.org/10.1002/hsr2.72066
31. Ahmed, F., Hasan, S., Hossain, A., & Rahman, K. A. (2026). Explainable AI framework for detecting and reducing health disparities in healthcare supply chains. Journal of AI, ML and DL, 2(1), 1–8. https://doi.org/10.25163/ai.2110685
32. Godavari Modalavalasa," SCALABLE CLOUD SOLUTIONS THROUGH ARTIFICIAL INTELLIGENCE GOVERNANCE: APPLICATIONS IN HEALTHCARE AND FINANCIAL SYSTEMS", Vol. 54 No. 1 (2026): January-March 2026, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/219, ,
DOI: https://doi.org/10.46121/pspc.54.1.25
33. Godavari Modalavalasa, AUTONOMOUS DATA ENGINEERING PIPELINES: A POLICY-DRIVEN ARCHITECTURE FOR SECURE AND SCALABLE CLOUD-NATIVE ANALYTICS, Vol. 53 No. 4 (2025): October-December 2025, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/181,
DOI: https://doi.org/10.46121/pspc.53.4.25
34. Godavari Modalavalasa, LARGE LANGUAGE MODELS FOR INTELLIGENT DATA ENGINEERING: AUTOMATING SCHEMA DESIGN, LINEAGE, AND QUALITY CONTROL, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/183 ,
DOI: https://doi.org/10.46121/pspc.50.2.4
35. Godavari Modalavalasa, FEDERATED LEARNING FOR ENTERPRISE CLOUD DATA ENGINEERING: ARCHITECTURE, SECURITY, AND GOVERNANCE CHALLENGES, Vol. 51 No. 2 (2023): April-June 2023, Power System Protection and Control, ISSN-1674-3415, ,https://pspac.info/index.php/dlbh/article/view/184,
DOI: https://doi.org/10.46121/pspc.51.2.5
36. Godavari Modalavalasa, AI-DRIVEN DATA GOVERNANCE: INTELLIGENT METADATA, LINEAGE, AND COMPLIANCE AUTOMATION IN CLOUD DATA PLATFORMS, Archives / Vol. 52 No. 1 (2024): January-March 2024 /, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/182,
DOI: https://doi.org/10.46121/pspc.52.1.3
37. Hima Bindu Lekkala, VishnuVardhan Bandari , AUTONOMOUS WORKFLOW OPTIMIZATION USING MULTI AGENT AI SYSTEMS AI AGENTS MANAGE STATIONS, WIP, AND TASK HANDOFFS, Vol. 54 No. 2 (202): April-June 2026, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/306 ,
DOI: https://doi.org/10.46121/pspc.54.2.08

38. Lekkala, H. B., & Bandari, V. V. (2026). Autonomous workflow optimization using multi-agent AI systems: AI agents manage stations, WIP, and task handoffs. Power System Protection and Control, 54(2), 95–101. http://pspac.info/index.php/dlbh/article/view/306
39. Bandari, V. V., & Lekkala, H. B. (2024). Physics-informed reinforcement learning for real-time control of complex manufacturing processes. Power System Protection and Control, 52(2), 164–170. https://pspac.info/index.php/dlbh/article/view/343
40. Lekkala, H. B., & Bandari, V. V. (2025). AI-based energy-aware scheduling and process optimization in engineer-to-order smart manufacturing systems. Power System Protection and Control, 53(1), 30–36. http://pspac.info/index.php/dlbh/article/view/341
41. Bandari, V. V., & Lekkala, H. B. (2026). AI-driven digital thread framework for end-to-end lifecycle optimization in ETO manufacturing systems. Power System Protection and Control, 54(2), 318–325. http://pspac.info/index.php/dlbh/article/view/340
42. Lekkala, H. B., & Bandari, V. V. (2026). Deep learning for weld defect detection using CNN or vision transformers fusion detection. Power System Protection and Control, 54(2), 151–158. https://pspac.info/index.php/dlbh/article/view/314
43. Bandari, V. V., & Lekkala, H. B. (2024). AI-based operator behavior monitoring and cost optimization using digital traceability in manufacturing systems. Power System Protection and Control, 52(1), 38–45. https://pspac.info/index.php/dlbh/article/view/342
44. Shreyas Jayaprakash, MACHINE LEARNING-DRIVEN VERIFICATION: USING ML TO OPTIMIZE COVERAGE ANALYSIS, PREDICT SIMULATION RUNTIMES, AND AUTO-GENERATE TESTBENCHES, Vol. 54 No. 2 (2026): April-June 2026, Power System Protection and Control, ISSN-1674-3415,
https://pspac.info/index.php/dlbh/article/view/391
45. Mohammed Shafi Kundiladi, EVENT-DRIVEN IMAGE AND VEHICLE STATUS MANAGEMENT FOR LOW-POWER IOT DIGITAL LICENSE PLATES, Vol. 53 No. 3 (2025): July-September 2025, Power System Protection and Control, ISSN-1674-3415,
https://pspac.info/index.php/dlbh/article/view/175
DOI: https://doi.org/10.46121/pspc.53.3.17
46. Controlled Synthesis and Characterization of Lanthanum Nanorods, Author(s), Jayadeep Tejani, Rahul Shah, Hiral Vaghela, Shailesh Vajapara, Amanullakhan Pathan, doi:10.18576/ijtfst/090205,
URL: https://www.naturalspublishing.com/Article.asp?ArtcID=20705, May-2020
47. Kunal Kapoor, Apr 2026, Relation Extraction and NER through Fine Tuned BERT model with transfer Learning, https://ieeexplore.ieee.org/document/11472214/metrics#metrics
48. Harender Bisht, Human–AI Collaboration in Insurance Fraud Detection: Ethical Cloud-Native Architectures for Fair and Transparent Decision Support, Volume 2026, Issue 1,
https://doi.org/10.52710/cfs.873
49. Harender Bisht, Ethical AI-Augmented Compliance Systems: Architectural Innovations for Cloud-Native Financial and Insurance Data Integrity, Vol. 11 No. 1s (2026) ,
https://doi.org/10.52783/jisem.v11i1s.14056

50. Harender Bisht, Governance-By-Design For AI-Based Insurance Fraud Detection: Auditability, Accountability, And Regulatory Traceability, Archives / Vol. 9 No. 1 (2026) https://doi.org/10.63278/jicrcr.vi.3620
51. Harender Bisht, Operationalizing Explainable AI in Insurance Fraud Detection: From Model Transparency to Decision Justification, Vol. 35 No. 1 (2026): JOCAAA, https://eudoxuspress.com/index.php/pub/article/view/4791,
52. A. MahendraVardhan and S. Sridhar, "Determining False Positive Analysis of Software Vulnerabilities with Predefined Scan Rules using Random Forest Classifier and Decision Tree Technique," 2022 4th International Conference on Advances in Computing, Communication Control and Networking (ICAC3N), Greater Noida, India, 2022, pp. 622-625, DOI: https://doi.org/10.1109/ICAC3N56670.2022.10074458
53. Amilineni, M.V. 2025. Audit-Safe Agentic RAG Framework for Regulated Enterprise Systems: A Trustworthy AI Architecture for SAP, Finance, Healthcare, and Government Workflows. INTERNATIONAL JOURNAL OF ADVANCES IN SIGNAL AND IMAGE SCIENCES. (Jan. 2025), 34–46. DOI: https://doi.org/10.29284/7fjzk883
54. Islam, Imtiaz (Annotator), Goard, Michael (Annotator), Daves, Daniel (Annotator), Credit Card Fraud Detection: A Comparative Study of Anomaly Detection and Supervised Learning under Extreme Class Imbalance, Published May 5, 2026 | Version v1, DOI: 10.5281/zenodo.20032207
55. Deepa Nagalavi; Shankar Balla Microsoft, USA; Pushpalatha. M; Nashwan Adnan Othman; Malik Bader Alazzam; A. Balakumar, Enhancing Real-Time Language Processing via Advanced PET Signal Analysis and Deep Learning, 06-07 June 2025, DOI: 10.1109/ICICKE65317.2025.11136561 , https://ieeexplore.ieee.org/document/11136561
56. Venkata Sai Aditya Kondru,FMR, Carmel, USA ,; Shankar Balla; Dhavalkumar Thakar; Dheeraj Velaga; Sri Lekha Bandla, Intelligent Human–Computer Interaction for Navigation Control Through Vision-Based Hand Gesture Recognition, Conferences >2026 IEEE 15th International Date of Conference: 07-09 April 2026,Date Added to IEEE Xplore: 08 May 2026 DOI: 10.1109/CSNT69054.2026.11502317 , https://ieeexplore.ieee.org/abstract/document/11502317
57. Sudipkumar Ghanvat , Aishwarya Badlani, Shreya Sandeep Joshi, Marketing Effectiveness in the Post-Cookie Era: A Comparative Analysis of First Party and Zero-Party Data Strategies on Campaign Performance and Customer Equity, Vol. 31 No. 4 (2023): JOCAAA , https://www.eudoxuspress.com/index.php/pub/article/view/3940
58. Navin Chhibber,; Deepak Singh; Anokh Kishore, Capital One, USA; Nikita Chawla; K. Anguraj, Multi-Granular Attention-Driven Reinforcement Learning Framework for Web Intelligent Enhancement Systems, , 11 June 2026, https://ieeexplore.ieee.org/document/11483386
59. Kapoor K, et al. AI-Driven COVID-19 Care: Analytics, Scheduling & Geo-Access. International Journal of Drug Delivery Technology (IJDDT). 2026;16(65S):Article 136. Available at: https://impactfactor.org/PDF/IJDDT/16/IJDDT,Vol16,Issue65s,Article136.pdf

60. Kilama, D. (2026). Mechanisms of Chinese Aid to Africa: A Comparative Analysis with Western Philanthropic and Development Models. Interdisciplinary Social Research, 2(1).DOI: https://doi.org/10.64220/isr.v2i1.011, https://scholarlysummit.com/journals/isr/articles/mechanisms-of-chinese-aid-to-africa-a-comparative-analysis
61. Kilama, D. (2026). Beyond Altruism: Analyzing China's Strategic Motivations Behind Development Aid to Africa. *Interdisciplinary Social Research, Volume 2 (2026)*(Issue 1), DOI: . https://doi.org/10.64220/isr.v2i1.012, Beyond Altruism: Analyzing China's Strategic Motivations Behind Development Aid to Africa
62. Anumolu DPK, Veena B, Afreen SS, Bandla J, Lakshmi KC, Pulusu VS. In vitro models for studying drug metabolites and metabolizing enzymes: A comprehensive review. Int J Pharm Investig. 2026;16(3):885–91. Available from: http://dx.doi.org/10.5530/ijpi.20260116
63. Asgar Z, Kumar G, Khandelwal P, Pratap B, Gurjar V, Baluni A, et al. Utility of contrast-enhanced computed tomography abdomen in Intestinal Obstruction. Int J Drug Deliv Technol. 2026;16(32s). Available from: http://dx.doi.org/10.25258/ijddt.16.32s.4
64. Anumolu DP, Ramatulasi J, Ashok G, Afreen SS, Mathew C, Shakar PV. Synthesis and characterization of MIP ghost approach for selective extraction of empagliflozin and quantification by liquid chromatography. J Chromatogr Sci. 2025;63(2). Available from: http://dx.doi.org/10.1093/chromsci/bmaf008.
65. Anumolu DPK, Naraparaju S, Addada SB, Pagidipally V, Pulusu VS, Afreen SS. Synthesis of silver nanoparticles using Ecbolium ligustrinum leaves: Characterization through advanced analytical techniques. Nanotechnol Percept . 2024;710–8. Available from: http://dx.doi.org/10.62441/nano-ntp.vi.3618
66. Naraparaju S, Mukti A, Anumolu DP, Chaganti S. Development and validation of a spectrofluorimetric method for the quantification of capecitabine in bulk and tablets. Turk J Pharm Sci [Internet]. 2023;20(4):234–9. Available from: http://dx.doi.org/10.4274/tjps.galenos.2022.46364
67. Naraparaju S, Yamjala P, Chaganti S, Anumolu DP. Spectrophotometric quantification of atomoxetine hydrochloride based on nucleophilic substitution reaction with 1,2-naphthoquinone-4-sulfonic acid sodium salt (NQS). Turk J Pharm Sci [Internet]. 2024;20(6):405–11. Available from: http://dx.doi.org/10.4274/tjps.galenos.2022.09147.
68. Anumolu PD, Shakar PV, Tulasi JR, Soujanya C, Afreen SS, Ashok G. Concurrent discriminative emission intensity quantification of Fexofenadine hydrochloride and Montelukast Sodium. Orient J Chem [Internet]. 2022;38(6):1364–8. Available from: http://dx.doi.org/10.13005/ojc/380605
69. Anumolu PD, Gurrala S, Menkana S, Mathew C, Srimantula L. Application of fluorotag in quantification of metformin in tablet dosage form and biological samples. The Thai Journal of Pharmaceutical Sciences [Internet]. 2022;46(1):93–8. Available from: http://dx.doi.org/10.56808/3027-7922.2550
70. Gurrala S, Raj S, Cvs S, Anumolu PD. Quality-by-design approach for chromatographic analysis of metformin, empagliflozin and linagliptin. J Chromatogr Sci [Internet]. 2022;60(1):68–80. Available from: http://dx.doi.org/10.1093/chromsci/bmab030

71. Gurrala S, Raj S, Cvs S, Anumolu DP, Naraparaju S, Nizampet H. Response surface methodology in spectrophotometric estimation of saxagliptin, derivatization with MBTH and ninhydrin. Turk J Pharm Sci [Internet]. 2022;19(1):9–18. Available from: http://dx.doi.org/10.4274/tjps.galenos.2021.93195

DOI: 10.46121/pspc.54.3.19