IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

AI-Driven Zero-Trust Identity Risk and Adaptive Access Engine: A Literature Review

[1]Rahul Bembade, [2]Sanchit Agalave, [3]Pushkar Patade, [4]Sarang Bele, [5]Sarthak Harpal

[1]Guide, [2]Researcher, [3]Researcher, [4]Researcher, [5] Researcher

[1,2,3,4,5]School Of Computing,

[1,2,3,4,5]MIT ADT University, Pune, India

Abstract: The transition from traditional, perimeter-based network defense to Zero-Trust Architecture (ZTA) represents a fundamental paradigm shift in modern cybersecurity. Historically, enterprise security models relied heavily on a "castle-and-moat" topology, operating under the assumption of implicit trust for any user, device, or application situated within the corporate firewall. The proliferation of remote workforces, cloud-native infrastructures, and sophisticated adversary-in-the-middle (AitM) phishing campaigns has rendered this perimeter obsolete, exposing the systemic vulnerabilities of persistent, session-based trust. In response, modern cybersecurity frameworks mandate a "never trust, always verify" methodology, necessitating continuous authentication and dynamic authorization on a per-request basis. This literature review critically examines the theoretical foundations and real-world implementation of AI-driven continuous authentication mechanisms within modern Identity and Access Management (IAM) systems. Specifically, it explores the deployment of lightweight, hybrid machine learning models within the Zero-Trust Policy Engine. This review details the efficacy of combining an unsupervised Isolation Forest algorithm for behavioral baselining—capable of detecting novel anomalies such as massive data exfiltration or atypical access hours without prior threat labeling—with a supervised Random Forest classifier that contextualizes risk based on device posture, geographic velocity, and threat intelligence. These models synthesize a dynamic risk score (1-100) that triggers automated Adaptive Enforcement protocols, dynamically gating access, mandating Step-Up Multi-Factor Authentication (MFA), or executing session termination. Furthermore, this paper analyzes the critical integration of identity risk telemetry into next-generation Security Information and Event Management (SIEM) systems. By leveraging standardized schemas such as the Open Cybersecurity Schema Framework (OCSF), modern Security Operations Centers (SOCs) can achieve unified threat visualization and orchestrate automated incident response. Ultimately, this review establishes that a unified, lightweight AI engine is an operational imperative for overcoming the latency, false-positive alert fatigue, and interoperability limitations inherent in existing enterprise security architectures.

Index Terms - Zero-Trust Architecture, Continuous Authentication, Machine Learning, Adaptive Access, Isolation Forest,Random Forest

I. INTRODUCTION

The fundamental failure of legacy cybersecurity architectures lies in their reliance on binary, location-centric trust. For decades, enterprise networks were constructed utilizing a perimeter defense model, heavily fortified by firewalls, Virtual Private Networks (VPNs), and intrusion detection systems. Within this "castle-and-moat" paradigm, identity verification occurred as a discrete, singular event at the network boundary. Once a user provided valid credentials, the system granted broad, persistent access to internal resources, assuming that internal network traffic was inherently benign. This architectural assumption has proven catastrophic in the face of modern cyber threats.

According to recent threat intelligence, including the Verizon Data Breach Investigations Reports from 2023 and 2025, approximately 88% of basic web application breaches are facilitated by the use of stolen or compromised credentials, with third-party breaches now accounting for 30% of all cases. When identity verification is limited to the perimeter, a compromised session token or bypassed password instantly grants threat actors unfettered lateral movement, allowing them to escalate privileges and exfiltrate sensitive data without triggering subsequent authentication challenges.

Recognizing these systemic vulnerabilities, the cybersecurity industry, guided by federal mandates and enterprise imperatives, has systematically adopted Zero Trust Architecture (ZTA). ZTA operates on the foundational premise that trust must never be granted implicitly based on physical or network location. Instead, trust must be continuously evaluated and explicitly verified for every access request, utilizing a comprehensive array of contextual and behavioral signals. However, the operationalization of continuous verification introduces immense computational and logistical challenges. In a distributed enterprise environment, evaluating millions of daily network interactions and access requests requires an analytical engine capable of processing high-velocity data streams in real time.

Early iterations of continuous authentication relied on static, heuristic rule sets crafted by security engineers. These rule-based systems proved highly brittle; they suffered from rapid decay as adversary tactics evolved and generated overwhelming volumes of false-positive alerts, which severely degraded user productivity and exhausted Security Operations Center (SOC) personnel. Conversely, recent academic proposals have advocated for the deployment of complex deep learning architectures, such as Long Short-Term Memory (LSTM) networks and Convolutional Neural Networks (CNNs), to analyze user behavior. While these deep learning models demonstrate high detection accuracy in theoretical simulations, they present severe limitations in production enterprise environments. Deep learning architectures impose prohibitive inference latency—often failing to meet the sub-millisecond response times required for seamless authentication gating—and function as opaque "black boxes," depriving security analysts of the explainability necessary to confidently triage and respond to alerts.

To resolve the dichotomy between static brittleness and deep learning opacity, modern enterprise Identity and Access Management (IAM) systems are pivoting toward lightweight, hybrid machine learning methodologies. The focus of this literature review is the implementation of an AI-Driven Zero-Trust Identity Risk and Adaptive Access Engine that utilizes a dual-layered algorithmic approach. By deploying an unsupervised Isolation Forest for dynamic behavioral baselining and pairing it with a supervised Random Forest for contextual risk classification, security systems can generate a highly accurate, explainable identity risk score. This risk score dictates real-time Adaptive Enforcement, enforcing the principle of least privilege through dynamic security friction. Furthermore, to bridge the gap between automated identity enforcement and broader incident response, this review explores the critical necessity of forwarding structured identity telemetry to SIEM dashboards utilizing industry-standard schemas.

By analyzing prominent frameworks such as the National Institute of Standards and Technology (NIST) Special Publication 800-207 and the Cybersecurity and Infrastructure Security Agency (CISA) Zero Trust Maturity Model (ZTMM), this paper maps theoretical machine learning concepts directly to their practical applications in real-world SOCs and enterprise architectures. The subsequent sections will detail the evolution of these enterprise networks, the mathematical and operational mechanics of the chosen machine learning algorithms, the execution of risk-based access controls, and the overarching role of unified telemetry in modern threat mitigation.

II. Literature Review

The Evolution of Zero Trust in Enterprise Networks

The conceptual framework of Zero Trust has evolved from an abstract theoretical proposition into a rigorously standardized architectural mandate. While the term was initially coined by industry analysts at Forrester, its practical viability at an enterprise scale was first demonstrated by Google's BeyondCorp initiative, which began deployment in 2014. The impetus for BeyondCorp was "Operation Aurora," a series of sophisticated, state-sponsored cyberattacks that successfully breached the internal networks of numerous technology companies by exploiting the implicit trust granted to internal corporate resources.

In response, Google engineered an architecture that systematically dismantled its traditional VPNs and privileged intranets. The BeyondCorp model shifted access controls entirely away from the network perimeter, placing them directly on individual users and their specific devices, regardless of whether those devices were connecting from a corporate office or a public Wi-Fi network. At the core of the BeyondCorp architecture is the "Trust Inferer," a dynamic access control engine that continuously analyzes millions of data

points—including device encryption status, the health of endpoint management agents, and historical user behavior—to dynamically adjust a centralized risk posture. This architectural shift proved that a massive, distributed workforce could securely access highly sensitive internal applications without relying on a protected network segment. The success of BeyondCorp established the foundational blueprint that would heavily influence subsequent federal guidelines and commercial IAM platforms.

Following the success of pioneering enterprise deployments, federal agencies moved to formally codify the principles of ZTA. The definitive standard governing this architecture is the National Institute of Standards and Technology (NIST) Special Publication (SP) 800-207. NIST SP 800-207 establishes seven core tenets of Zero Trust, explicitly stating that all data sources and computing services are considered resources, all communication must be secured regardless of network location, and access to individual enterprise resources must be granted on a per-session basis. To execute these tenets, the NIST framework defines a specific logical architecture centered around the Policy Decision Point (PDP) and the Policy Enforcement Point (PEP). The PDP is further bifurcated into two distinct, highly critical components:

- The Policy Engine (PE): The analytical brain of the Zero Trust architecture. The PE utilizes a trust algorithm to evaluate a continuous stream of contextual signals—such as identity attributes, Continuous Diagnostics and Mitigation (CDM) data, threat intelligence feeds, and behavioral logs—to formulate a definitive access decision (grant, deny, or revoke).
- The Policy Administrator (PA): The control plane component responsible for executing the decisions made by the Policy Engine. The PA generates session credentials and configures communication paths to the requested resources.

All enterprise traffic must pass through the Policy Enforcement Point (PEP), which maintains or terminates the connection based strictly on the continuous directives issued by the PA. The operational challenge identified in the literature is that the Policy Engine cannot function effectively using static rules; it requires a highly responsive artificial intelligence engine to process the sheer volume of telemetry envisioned by the NIST standard.

Building upon the NIST definitions, the Cybersecurity and Infrastructure Security Agency (CISA) released the Zero Trust Maturity Model (ZTMM) Version 2.0 to provide organizations with a phased implementation roadmap. The CISA ZTMM divides the Zero Trust ecosystem into five foundational pillars: Identity, Devices, Networks, Applications and Workloads, and Data. These pillars are supported by three cross-cutting capabilities: Visibility and Analytics, Automation and Orchestration, and Governance. For the development of an AI-driven risk engine, the Identity pillar is paramount. CISA defines four progressive maturity levels within this pillar, guiding organizations away from legacy vulnerability toward dynamic, continuous verification:

Maturity Level	**Identity Controls & Authentication Mechanisms**	**Automation, Analytics & Policy Enforcement**
Traditional	Heavy reliance on static passwords and static credentials. Trust is primarily established at the initial provisioning phase and tied to network location.	Manual configurations for access assignments. Static security policies are applied broadly. Implicit trust persists post-authentication with no mid-session evaluation.
Initial	Agency-wide adoption of Multi-Factor Authentication (MFA), though often reliant on weaker, phishable factors (e.g., SMS, standard push notifications).	Early stages of automated lifecycle management. Basic integration of centralized identity stores. Policy enforcement remains largely binary and static.
Advanced / Optimal	Implementation of phishing-resistant MFA (e.g., FIDO2, smart cards). Dynamic attribute-based access control (ABAC) utilizing cross-pillar contextual data. Fully passwordless	Automated responses to common security events. Access decisions incorporate device compliance, location velocity, and network signals. Fully automated, just-in-time

Maturity Level	Identity Controls & Authentication Mechanisms	Automation, Analytics & Policy Enforcement
	authentication mechanisms. Continuous, real-time behavioral validation and dynamic risk assessment integrated directly into the access workflow.	privilege escalation. Machine learning-driven dynamic risk scoring determines access and enforcement down to the per-request level.

Achieving the "Optimal" maturity level described by CISA is technically unfeasible without the integration of advanced machine learning within the NIST Policy Engine. In a fully mature ZTA, the system must continuously map identity behavior across disparate environments, assign trust scores automatically, and proactively block malicious activity without human intervention. Consequently, the evolution of Zero Trust has transitioned from a networking and topology challenge into an applied data science imperative, wherein the overall efficacy of the security architecture is entirely dependent on the speed, accuracy, and interpretability of the underlying AI risk engine.

Machine Learning in Continuous Authentication

The mandate for continuous authentication requires a computational engine capable of processing high-velocity, high-dimensional telemetry streams to detect anomalous behavior while simultaneously assessing complex contextual risk. As established, purely rule-based systems are inadequate for this task. They require constant, manual tuning by security engineers, generate excessive false positives, and inherently fail to adapt to novel, zero-day adversary techniques. Conversely, a significant portion of academic literature focuses on the application of deep learning models—such as Artificial Neural Networks (ANNs), Convolutional Neural Networks (CNNs), and Long Short-Term Memory (LSTM) networks—for behavioral profiling and anomaly detection.

While these deep architectures often achieve superior detection rates on static benchmark datasets, they present severe operational limitations when deployed in real-time enterprise Security Operations Centers (SOCs). Deep learning models suffer from high inference latency, requiring substantial computational resources that delay the authentication gateway process. Furthermore, they demand massive volumes of labeled training data, which is rarely available for novel insider threats. Most critically, deep learning networks function as "black boxes." In a SOC environment, an analyst cannot justify an automated containment action if the system cannot explicitly explain why an identity was scored as high risk.

To address these limitations, modern enterprise architectures and applied cybersecurity frameworks are increasingly adopting lightweight, tree-based machine learning algorithms. These algorithms provide a superior balance of high predictive performance, scalability, rapid inference, and critical transparency. The most effective architecture identified in recent applied research relies on a layered, hybrid approach: utilizing an unsupervised Isolation Forest for baseline anomaly detection and a supervised Random Forest for contextual risk classification.

Unsupervised Behavioral Baselining: The Isolation Forest: The Isolation Forest (iForest) algorithm is exceptionally well-suited for behavioral baselining within continuous authentication systems. Traditional anomaly detection models (e.g., One-Class SVM or k-means clustering) attempt to profile "normal" behavior and flag instances that deviate from that dense baseline. This approach is computationally expensive in high-dimensional spaces. In contrast, the Isolation Forest operates on a fundamentally different principle: it explicitly isolates anomalies by building an ensemble of random decision trees. The algorithm recursively partitions the dataset by randomly selecting a feature and then randomly selecting a split value between the maximum and minimum values of that selected feature.

Because anomalies are statistically "few and different," they are far more susceptible to isolation. Consequently, anomalous data points are isolated closer to the root of the isolation trees, requiring significantly fewer partitioning steps (shorter path lengths) than normal data points. The mathematical efficiency of this approach is vital for real-time risk scoring. The anomaly score s(x, n) for an observation is calculated based on its expected path length E(h(x)) across the ensemble of isolation trees, normalized by the average path length of unsuccessful searches in a Binary Search Tree c(n):

$$\frac{E(h(x))}{}$$

$$s(x, n) = 2 \quad c(n)$$

In this formulation, if the anomaly score approaches 1, the observation is highly anomalous; if it approaches 0.5, the observation exhibits normal behavior. In the context of enterprise IAM, the Isolation Forest continuously monitors unstructured behavioral telemetry without requiring prior attack labels. It evaluates variables such as session duration, resource access intensity, unusual login hours, and massive data exfiltration volumes. Its unsupervised nature is critical because malicious insider behavior and novel credential stuffing patterns are rarely represented in historical training datasets. Real-world implementations demonstrate that the Isolation Forest effectively handles sparse, high-dimensional behavioral data streams with minimal CPU overhead, making it an ideal first layer for a Zero-Trust Policy Engine.

Supervised Contextual Risk Classification: The Random Forest: While the Isolation Forest efficiently flags statistical outliers, operational reality dictates that not all outliers represent malicious activity. For example, a Chief Executive Officer logging in at 3:00 AM from an unrecognized IP address in a foreign country will generate a highly anomalous score in the Isolation Forest. If the system executes a hard block based solely on this unsupervised score, it severely disrupts legitimate business operations and generates false-positive alert fatigue.

To mitigate this, enterprise architectures pair the unsupervised layer with a supervised Random Forest classifier. The Random Forest algorithm utilizes an ensemble of decision trees trained on historically labeled contextual data. It employs bagging (bootstrap aggregating) to train individual trees on random subsets of the data, and random feature selection at each node split. This ensemble approach dramatically reduces the risk of overfitting, handling non-linear relationships across diverse feature spaces. Within the Zero-Trust Policy Engine, the Random Forest evaluates contextual features that qualify the behavioral anomaly. These features include device posture (e.g., managed vs. unmanaged, disk encryption status), IP reputation (e.g., known Tor exit nodes, anonymous VPNs), location velocity (impossible travel calculations), and threat intelligence feeds. Through ensemble voting, the model categorizes the authentication event, generating a precise threat probability score.

Critically, the Random Forest provides intrinsic feature importance analysis. This explainability allows the system to output structured reasoning alongside its risk score, enabling SOC analysts to understand the exact mechanics of the decision (e.g., "75% risk contribution from unmanaged device status, 20% from anomalous behavioral baseline, 5% from abnormal geographic location"). Recent frameworks documented in scientific literature, such as the "Shadow Sentinal" intelligence monitoring system and the "ZenGuard" Zero-Trust framework, effectively validate this hybrid ML architecture in simulated and real-world enterprise environments. By sequentially passing authentication telemetry through an Isolation Forest to detect baseline deviations, and subsequently through a Random Forest for contextual validation, these frameworks have achieved robust detection accuracies (exceeding 92%) while documenting up to a 45% reduction in false-positive alerts. The hybridization of these lightweight models provides a computationally efficient, highly interpretable mechanism to continuously recalculate identity trust without imposing undue latency on the authentication gateway.

Adaptive Access and Risk-Based MFA in the Industry

The immediate operational output of the hybrid AI-driven risk engine is a dynamic, continuous risk score. To integrate seamlessly with enterprise IAM and comply with standardized actuarial tracking, this score is typically normalized on a continuous scale from 1 to 100. In a legacy IAM environment, the authentication process is highly static: a user provides a username, a password, and perhaps a static secondary factor (like an SMS code); if the credentials match the directory, access is granted for the duration of a long-lived session token.

In a modern ZTA leveraging the continuous 1-100 risk score, the Policy Engine dictates dynamic Adaptive Enforcement actions. This methodology ensures that security friction is applied proportionally to the evaluated risk at any given moment, fulfilling the CISA ZTMM "Optimal" criteria for automated, risk-responsive application profile assessments. Prominent real-world enterprise IAM solutions, such as Microsoft Entra ID Protection and Okta Adaptive MFA, utilize this continuous scoring mechanism to drive comprehensive Conditional Access policies. Within these enterprise platforms, the overall identity risk is typically bifurcated into two distinct, yet interacting, categories:

- User Risk (Long-Term Compromise Probability): This metric represents the probability that a given identity has been fundamentally compromised over time. It is heavily informed by offline, asynchronous threat intelligence (e.g., the detection of the user's specific credentials in a dark web

data breach dump) and long-term macro-behavioral deviations calculated by the machine learning models.

- Sign-in Risk (Real-Time Session Probability): This metric represents the real-time probability that a specific, active authentication request is malicious. It is calculated by evaluating the immediate context of the request, such as connections originating from anonymous IP addresses, malware-linked networks, or highly unfamiliar sign-in properties.

Based on the aggregated risk score derived from the Isolation Forest and Random Forest models, the Policy Administrator component of the ZTA automatically triggers predefined enforcement actions. These Adaptive Enforcement tiers function as follows:

Score Tier	Risk Level	Automated Action	Operational Scenario & Context
1-30	Low Risk	Grant Access (Frictionless)	A user accesses an enterprise application from a managed, highly trusted device within standard operating hours. The login pattern matches the historical behavioral baseline. The system grants seamless access, potentially leveraging passwordless Single Sign-On (SSO).
31-70	Medium Risk	Require Step-Up MFA	A user logs in from a new geographic location or an unmanaged personal device. The Isolation Forest flags a behavioral anomaly, but the Random Forest identifies valid contextual credentials. The system dynamically interrupts the workflow to force a step-up challenge (e.g., FIDO2 hardware key, biometric validation).
71-100	High Risk	Block Access & Force Remediation	The ML engine detects “impossible travel” combined with an anomalous data exfiltration pattern. The Policy Administrator immediately terminates the active session. Access is strictly blocked, and the identity is suspended until an automated remediation playbook forces a secure password reset and notifies the SOC.

The operational and strategic benefits of implementing Adaptive MFA driven by an AI risk score are profound. It provides a highly robust, real-time defense against modern attack vectors—such as automated credential stuffing, SIM-swapping, and adversary-in-the-middle (AitM) proxy attacks—that easily bypass static MFA. Simultaneously, it drastically improves the end-user experience. By reserving stringent security challenges exclusively for high-risk anomalies, the system minimizes unnecessary prompts during legitimate daily activity, thereby reducing alert fatigue and preventing users from developing “MFA fatigue,” a psychological vulnerability increasingly exploited by threat actors.

Role of SIEM and Automated Threat Visualization in SOCs

While the AI-driven risk engine autonomously handles real-time access decisions and Adaptive Enforcement at the identity gateway, comprehensive enterprise security relies heavily on Security Information and Event Management (SIEM) systems. In a mature Zero Trust framework, the SIEM functions as the central intelligence hub, aggregating disparate telemetry from Policy Enforcement Points (network proxies), Identity Providers (IdPs), Cloud Access Security Brokers (CASBs), and Endpoint Detection and Response (EDR) agents.

A persistent, critical challenge in modern SOC operations is the “telemetry gap” caused by massive data fragmentation. As noted in prominent industry threat reports, advanced attackers rarely constrain their operations to a single domain. A typical kill chain involves compromising an identity, moving laterally to an endpoint, escalating privileges, and finally exfiltrating data through a cloud application. If the telemetry generated by the EDR platforms, cloud APIs, and identity systems remains siloed in proprietary, heterogeneous formats, the SIEM struggles to correlate these disparate signals. Historically, security engineers

were forced to waste hundreds of hours building and maintaining fragile, vendor-specific application programming interfaces (APIs) and regular expression (RegEx) parsers to normalize these logs. When a proprietary log format changed, the parser failed silently, resulting in critical identity anomalies failing to trigger broader infrastructure alerts, severely delaying the Mean Time to Respond (MTTR).

To resolve this systemic integration failure, a collaborative consortium of cybersecurity vendors introduced the Open Cybersecurity Schema Framework (OCSF). OCSF is an open-source, vendor-neutral data taxonomy designed to standardize the structure and semantics of security event data at the point of ingestion. By establishing a common language for cybersecurity telemetry, OCSF eliminates the need for fragile translation layers, significantly streamlining the ingestion and correlation of identity risk data into next-generation SIEMs (such as CrowdStrike Falcon Next-Gen SIEM or Datadog Cloud SIEM). Within the OCSF schema, identity risk and continuous authentication events are structured under highly specific, predefined event classes. For example, when the AI-driven risk engine processes an authentication attempt and triggers an Adaptive Access policy, it natively maps the event to the OCSF Authentication event class (Type ID 300201 for Logon or 300202 for Logoff).

OCSF Schema Field	**Description & Application for Identity Risk Telemetry**
class uid	A numerical identifier representing the specific event class (e.g., 3002 strictly denotes Authentication events), allowing the SIEM to instantly route the log to the identity analysis pipeline.
user	A structured entity array that universally identifies the subject, role, or account attempting the authentication, standardizing identifiers across cloud and on-premises directories.
device	A structured entity containing deep contextual posture details, including IP reputation, operating system compliance, and managed/unmanaged status as evaluated by the Random Forest.
risk score	The dynamically calculated, continuous AI trust score (1-100) generated by the Policy Engine, allowing analysts to filter dashboards strictly by quantitative risk.
severity id	A categorical urgency indicator (Low, Medium, High) that dictates the prioritization queue for SOC analysts and triggers corresponding automated triage levels.

The standardized attributes within this schema ensure that every system in the SOC understands the exact context of the event. By outputting structured OCSF logs natively, the AI risk engine seamlessly bridges the gap between identity security and broader infrastructure defense. When a “High Risk” score (e.g., an OCSF event with severity id of High and a risk score of 95) is ingested by the SIEM, it instantly visualizes the threat on unified SOC dashboards without parsing delays.

Furthermore, this standardized telemetry acts as the immediate catalyst for Security Orchestration, Automation, and Response (SOAR) platforms. Leveraging the standardized OCSF data, SOAR playbooks can execute dynamic, cross-domain response actions across the enterprise at machine speed. For instance, an identity compromise detected by the AI engine can automatically trigger a Python-based SOAR playbook that instructs edge firewalls to block the malicious IP address, commands the EDR platform to cryptographically isolate the user’s physical endpoint, and universally revokes all active OAuth session tokens across integrated SaaS applications. This automated threat visualization and orchestrated response, fueled by OCSF-compliant machine learning telemetry, is the ultimate realization of a resilient Zero-Trust architecture.

III. Gap Analysis & Problem Statement

Despite the significant architectural advancements formalized by NIST and CISA, and the rapid commercial adoption of Zero Trust principles, critical gaps remain in the operationalization of continuous identity verification within complex enterprise environments. The fundamental challenge lies in balancing the computational demands of advanced threat detection with the operational realities of real-time authentication workflows and SOC triage capabilities.

First, there is a pronounced operational conflict between the theoretical accuracy of complex artificial intelligence and the stringent performance requirements of a Zero-Trust Policy Engine. A substantial portion of academic literature advocates for the implementation of deep learning architectures—such as deep autoencoders, LSTMs, and CNNs—for behavioral anomaly detection. However, these models introduce significant computational overhead and inference latency. In a production environment, authentication gateways must render decisions in sub-milliseconds; introducing high-latency deep learning models results in unacceptable user friction and frequent session timeouts. Furthermore, deep learning networks function as opaque "black boxes." In a highly regulated SOC environment, analysts cannot confidently execute severe containment protocols—such as revoking enterprise-wide access or isolating executive endpoints—if the system cannot explicitly explain the origin of the risk.

Second, many commercial identity solutions attempting to bypass the complexity of AI continue to rely on static, location-centric, and heuristic risk policies. These static rules generate an excessive volume of false positives and are inherently brittle. As sophisticated adversaries increasingly leverage stolen, long-lived session tokens and route their attacks through residential proxy networks, static rules evaluating simple geographic locations or basic browser fingerprints are easily and consistently bypassed.

Finally, the telemetry generated by independent, disparate identity platforms often lacks a standardized taxonomy. If identity telemetry is forwarded to a SIEM in proprietary, unstructured formats, it exacerbates alert fatigue and creates massive blind spots in correlation engines. The foundational Zero Trust assumption of "continuous verification" breaks down into disconnected, siloed assessments if identity anomalies cannot be instantaneously correlated with corresponding network and endpoint telemetry.

Problem Statement: There is a critical, unmet need for a unified, highly efficient AI risk engine that simultaneously addresses algorithmic performance limitations, model explainability, and telemetry standardization. Existing IAM solutions either rely on brittle static rules that fail against modern credential abuse, or theoretical deep learning models that are operationally unviable due to latency and opacity.

The proposed "AI-Driven Zero-Trust Identity Risk and Adaptive Access Engine" directly addresses these systemic gaps. By eschewing deep learning in favor of lightweight, tree-based machine learning algorithms, the system achieves the necessary inference speeds for real-time authentication gating. The engine utilizes an Isolation Forest for unsupervised behavioral baselining, efficiently identifying novel anomalies without requiring prior threat labeling. It pairs this with a Random Forest for supervised contextual risk classification, ensuring that risk scores are highly accurate and, crucially, structurally explainable through feature importance weighting. By translating this aggregated intelligence into a dynamic 1-100 risk score, the system scales security friction proportionately via Adaptive Access controls (Grant, Step-Up MFA, Block). Finally, by standardizing the output telemetry strictly through the Open Cybersecurity Schema Framework (OCSF), the engine ensures that SOC analysts and automated SOAR playbooks possess the normalized data required to execute instantaneous, cross-domain threat mitigation

IV. Conclusion

The rapid dissolution of the traditional network perimeter, driven by cloud migration and remote enterprise operations, mandates the adoption of Zero Trust Architectures that evaluate risk continuously at the identity level. As definitively established by the NIST SP 800-207 guidelines and the CISA Zero Trust Maturity Model, the theoretical constructs of dynamic Policy Engines and Policy Enforcement Points are critical for defending against modern credential-based attacks. However, the true efficacy of these systems in live enterprise environments relies entirely on the processing capabilities, accuracy, and interpretability of the underlying risk assessment engine.

By strategically transitioning away from brittle, static heuristic rules and computationally prohibitive deep learning networks, modern security architectures can leverage the distinct advantages of lightweight machine learning methodologies. The hybridization of Isolation Forests—to detect subtle behavioral anomalies such as data exfiltration velocity without the need for historical attack labels—combined with Random Forests—to classify contextual risk and provide transparent feature explainability—enables highly accurate, continuous 1-100 risk scoring. When this intelligence is coupled with dynamic Adaptive Access mechanisms, the system

achieves the dual mandate of modern IAM: it dramatically reduces user friction for legitimate daily traffic while imposing stringent, phishing-resistant MFA challenges exclusively on suspicious requests.

Furthermore, standardizing this complex identity telemetry through the vendor-neutral Open Cybersecurity Schema Framework (OCSF) ensures that SIEM platforms and SOC analysts possess the immediate, correlated context required to orchestrate automated incident response via SOAR playbooks. Ultimately, implementing this intelligent, highly integrated access engine bridges the critical gap between academic algorithmic theory and applied security operations, transforming Zero Trust from a conceptual framework into a proactive, resilient, and fully automated defense mechanism capable of mitigating the most sophisticated modern cyber threats.

REFERENCES

[1] Rose, S., Borchert, O., Mitchell, S., and Connelly, S. 2020. Zero Trust Architecture. NIST Special Publication 800-207.

[2] Cybersecurity and Infrastructure Security Agency (CISA). 2023. Zero Trust Maturity Model Version 2.0.

[3] Ward, B., and Beyer, B. 2014. BeyondCorp: A New Approach to Enterprise Security. login:, 39(6): 6-11.

[4] Hassan, A., et al. 2025. ZenGuard: a machine learning based zero trust framework for context-aware threat mitigation using SIEM, SOAR and UEBA. Scientific Reports, 15(35871).

[5] Microsoft. 2024. Microsoft Entra ID Protection - Concept: Identity Protection Policies.

[6] OCSF Community. 2022. Open Cybersecurity Schema Framework (OCSF) Documentation.

[7] Janarthanan, R., et al. 2026. Shadow Sentinal: A Semi-Autonomous AI-Based Intelligence Monitoring Framework. International Journal of Research and Scientific Innovation (IJRSI), 13(3): 1172-1182.

[8] Verizon. 2023. 2023 Data Breach Investigations Report (DBIR).

[9] Okta. 2024. Implementing Risk-Based Authentication With Okta

[10] Verizon. 2025. 2025 Data Breach Investigations Report (DBIR). [11] Anonymous et al. 2024. Insider Threat Detection Model using Anomaly-Based Isolation Forest Algorithm. IEEE.