Excel International Journal of
Technology, Engineering & Management
Journal homepage: exceljournals.org.in
ISSN ONLINE: https://portal.issn.org/resource/ISSN/2349-8455
ISSN PRINT: https://portal.issn.org/resource/ISSN/2348-0378

AI-Driven Continuous Governance for Machine Identities in Cloud-Native Zero Trust Environments

Kaushik Reddy Muppa
Jawahar Lal Technological University Hyderabad

Published online: April 2026
DOI Link: https://doi.org/10.64971/j.cph.eijtem.v13.i2.11.2026

Abstract
Cloud-native systems and Zero Trust architectures have fundamentally restructured enterprise security by elevating identity to the primary enforcement boundary. As distributed infrastructures scale, Non-Human Identities (NHIs), including workload credentials, service accounts, APIs, containers, serverless functions, and autonomous agents, now dominate authentication flows and execute most privileged interactions. However, identity governance mechanisms have not evolved at the same pace. Existing Identity Governance and Administration (IGA) models remain largely human-centric, review-driven, and static, resulting in a structural gap between the scale of identity and governance capability.
This gap manifests as privilege drift, entitlement sprawl, prolonged credential exposure, and expanded lateral-movement surfaces across service communication graphs. We formalize this challenge as the Machine Identity Governance Problem (MIGP): the need to continuously assess, constrain, and adapt privileges for dynamically provisioned machine identities operating at automation scale.
To address this problem, we propose an AI-driven Continuous Governance Framework (ACGF) that reconceptualizes identity governance as a closed-loop, risk-adaptive control system. ACGF integrates behavioral telemetry, privilege relationship modeling, credential lifecycle intelligence, and automated policy enforcement to minimize machine identity risk while preserving operational continuity continuously. We present a cloud-native reference architecture, describe its practical implementation, and evaluate its effectiveness in reducing privilege accumulation and containment latency in high-churn environments.
Our findings indicate that continuous, telemetry-driven governance significantly improves privilege hygiene and limits lateral exposure without introducing instability. More broadly, this work establishes adaptive machine identity governance as a foundational security discipline for Zero Trust ecosystems, providing a conceptual and architectural framework for future research in machine-scale identity control.
Keywords: AI-Driven, Machine Identities, Cloud-Native Zero Trust Environments

1. Introduction

Cloud-native computing has transformed the structure of modern enterprise systems. Applications are increasingly decomposed into microservices, deployed through container orchestration platforms, and interconnected through service meshes and APIs. In parallel, Zero Trust architectures have shifted the enforcement boundary from network perimeters to identity-driven access control. Within this paradigm, identity is no longer an attribute of human users alone it is the primary security primitive governing every interaction across the distributed infrastructure.
As automation expands, Non-Human Identities (NHIs) including workload certificates, service accounts, API credentials, containers, serverless functions, and autonomous agents now execute the majority of privileged operations in production environments. In large-scale systems, machine identities can outnumber human identities by orders of magnitude. These identities are dynamically provisioned, short-lived, replicated across clusters, and embedded in CI/CD workflows, infrastructure-as-code pipelines, and service-to-service communication paths.
Despite this fundamental shift, identity governance mechanisms have not evolved at the same pace. Traditional Identity Governance and Administration (IGA) frameworks were designed for relatively stable human identities characterized by predictable lifecycle events, clear ownership attribution, and periodic

access review cycles. These models assume bounded growth and manual oversight. Such assumptions do not hold in machine-dominated systems.

This structural mismatch introduces systemic risk. Machine identities frequently accumulate excessive permissions through configuration inheritance, role reuse, deployment artifacts, and temporary operational exceptions. Privilege drift persists because review-based governance mechanisms cannot operate at machine scale. Furthermore, service communication graphs create complex lateral-movement surfaces, where compromise of a single workload identity may expose downstream services through pre-authorized trust relationships. Long-lived credentials exacerbate this exposure, increasing attacker dwell time and blast radius.

We refer to this emerging challenge as the **Machine Identity Governance Problem (MIGP)**: the need to continuously assess, constrain, and adapt privileges for dynamically provisioned non-human identities operating at automation scale. Unlike traditional IAM problems centered on authentication or misconfiguration detection, MIGP concerns the continuous alignment between identity behavior, privilege assignment, and system topology in high-churn environments.

Addressing MIGP requires rethinking governance as a continuous control discipline rather than a periodic compliance exercise. Detection, policy modeling, and enforcement must operate as an integrated feedback loop. Behavioral telemetry must inform privilege adaptation. Credential lifecycle management must be risk aware. Governance decisions must be automated yet bounded by operational guardrails.

In this paper, we propose an AI-driven Continuous Governance Framework (ACGF) that reconceptualizes machine identity governance as a closed-loop, risk-adaptive control system aligned with Zero Trust principles. The framework integrates privilege relationship modeling, behavioral analytics, credential lifecycle intelligence, and automated enforcement to continuously minimize machine identity risk without destabilizing production workloads.

This work makes three primary contributions:

1. It formalizes the Machine Identity Governance Problem as a distinct and emerging security discipline within Zero Trust ecosystems.
2. It presents a cloud-native architectural model for continuous, adaptive governance of non-human identities.
3. It demonstrates, through implementation and evaluation, that telemetry-driven privilege adaptation can significantly improve privilege hygiene and containment efficiency in high-scale environments.

By establishing adaptive machine identity governance as a foundational security requirement, this work aims to provide a conceptual and architectural basis for future research in machine-scale identity control, risk-adaptive access enforcement, and autonomous security operations.

2. Zero Trust Foundations

Zero Trust Architecture (ZTA), formally defined in NIST SP 800-207, reorients enterprise security by eliminating implicit trust derived from network location or perimeter boundaries. Rather than assuming internal traffic is trustworthy, Zero Trust mandates continuous verification of identity, device posture, contextual signals, and policy compliance prior to granting access to protected resources.

At its core, Zero Trust reframes identity as the primary enforcement boundary of modern systems.

The foundational tenets of Zero Trust include:

- All data sources and services are treated as protected resources
- Communication is secured regardless of network location
- Access is granted on a per-session basis
- Policy decisions are dynamically evaluated based on context
- Continuous monitoring of asset posture and integrity is required
- Authentication and authorization are strictly enforced before access

In cloud-native environments, these principles extend beyond human users to encompass workloads, APIs, microservices, containers, serverless functions, and autonomous agents. Identity is no longer merely an authentication mechanism; it becomes the control plane governing service-to-service interaction, infrastructure orchestration, and automated execution flows.

However, while Zero Trust establishes continuous verification as a policy requirement, it does not prescribe how identity privileges should evolve in response to behavioral change, topology shifts, or automation-driven scale. The Zero Trust model focuses primarily on enforcing access decision enforcement rather than managing the ongoing privilege lifecycle governance.

This distinction becomes critical in machine-dominated systems. In large-scale cloud-native environments, non-human identities frequently outnumber human users by orders of magnitude. These identities are dynamically created, replicated, and terminated through orchestration platforms and CI/CD pipelines. Their privilege sets evolve through deployment artifacts, inherited roles, and operational adjustments.

Zero Trust ensures that each access request is evaluated, but it does not inherently guarantee that the underlying privilege assignments remain minimal, adaptive, or behaviorally aligned.

As a result, an architectural gap emerges between Zero Trust policy enforcement and machine identity governance at scale. This gap motivates the formalization of the Machine Identity Governance Problem (MIGP), which extends Zero Trust principles into continuous, telemetry-driven privilege adaptation for non-human identities.

NIST Tenets of Zero Trust Architecture

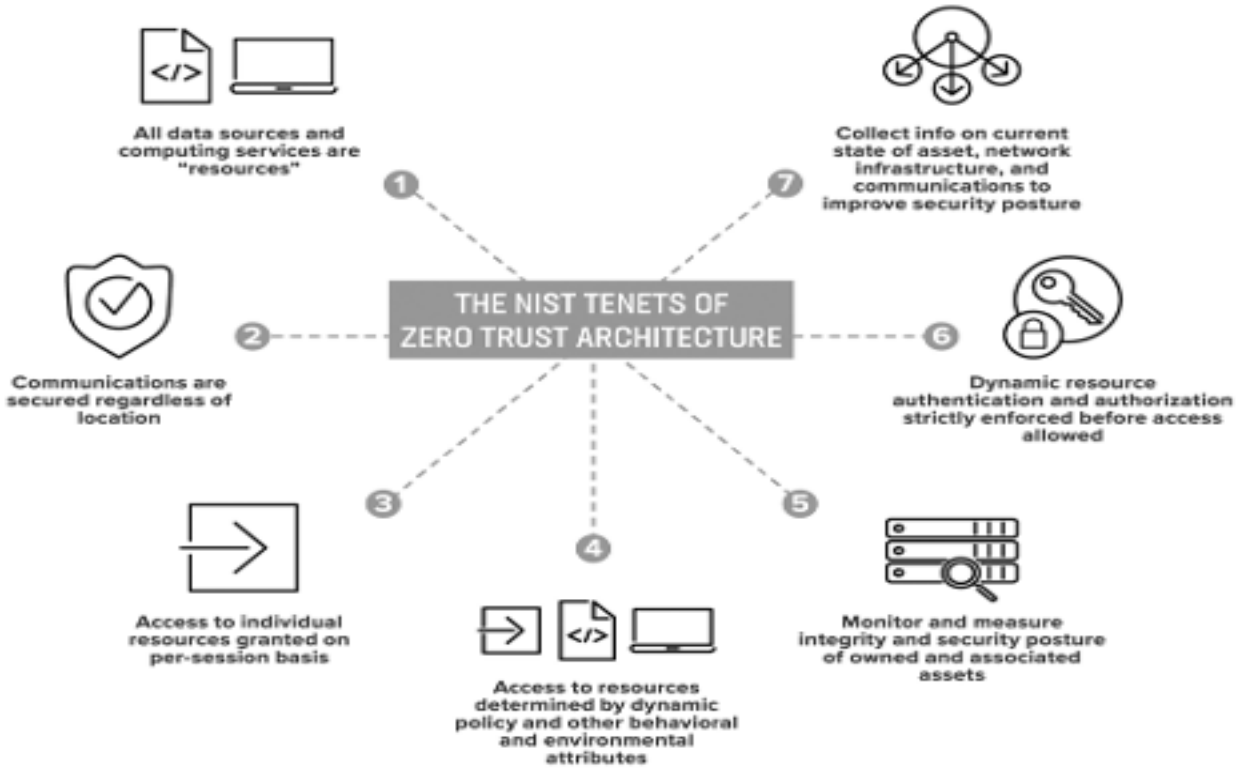

While ZTA provides foundational guidance, it does not fully address how governance mechanisms must evolve when machine identities dominate access decisions.

3. Machine Identity Governance Problem

In modern cloud-native environments, machine identities are provisioned automatically through orchestration platforms, infrastructure-as-code pipelines, and dynamic scaling systems. These identities often lack explicit ownership attribution and are frequently short-lived yet highly privileged.

Over time, privileges assigned to machine identities may expand due to:

- Configuration inheritance
- Role reuse across deployments
- Temporary debugging permissions
- Pipeline misconfigurations
- Operational shortcuts

Without continuous evaluation, these privileges persist and accumulate, creating privilege drift.

Additionally, service-to-service communication graphs introduce lateral-movement pathways. If a workload identity is compromised, downstream services and sensitive resources may become reachable through pre-authorized trust relationships.

The Machine Identity Governance Problem therefore involves designing mechanisms that:

- Continuously evaluate machine identity risk posture
- Detect excessive or unused entitlements
- Reduce reachable service exposure
- Minimize credential exposure windows
- Automatically enforce least-privilege policies
- Maintain operational continuity

Traditional review-based governance processes cannot operate at the scale and velocity of machine-driven systems. Continuous, automated governance is required.

4. Cloud-Native Identity Governance Model

Effective machine identity governance requires an integrated lifecycle model that operates across discovery, validation, control, intelligence, and response.

A cloud-native identity governance platform must incorporate:

1. **Posture Management** – Continuous discovery and inventory of machine identities and their associated privileges
2. **Validation and Policy Enforcement** – Context-aware authentication and dynamic policy evaluation
3. **Lifecycle Management** – Automated provisioning, deprovisioning, and access review for machine identities
4. **Intelligence and Analytics** – Behavioral analysis and risk-based identity scoring
5. **Detection and Response** – Segmentation, privilege restriction, and remediation workflows

These domains form a continuous feedback loop in which identity behavior informs policy adaptation and enforcement.

Cloud-Native Identity Management Platform

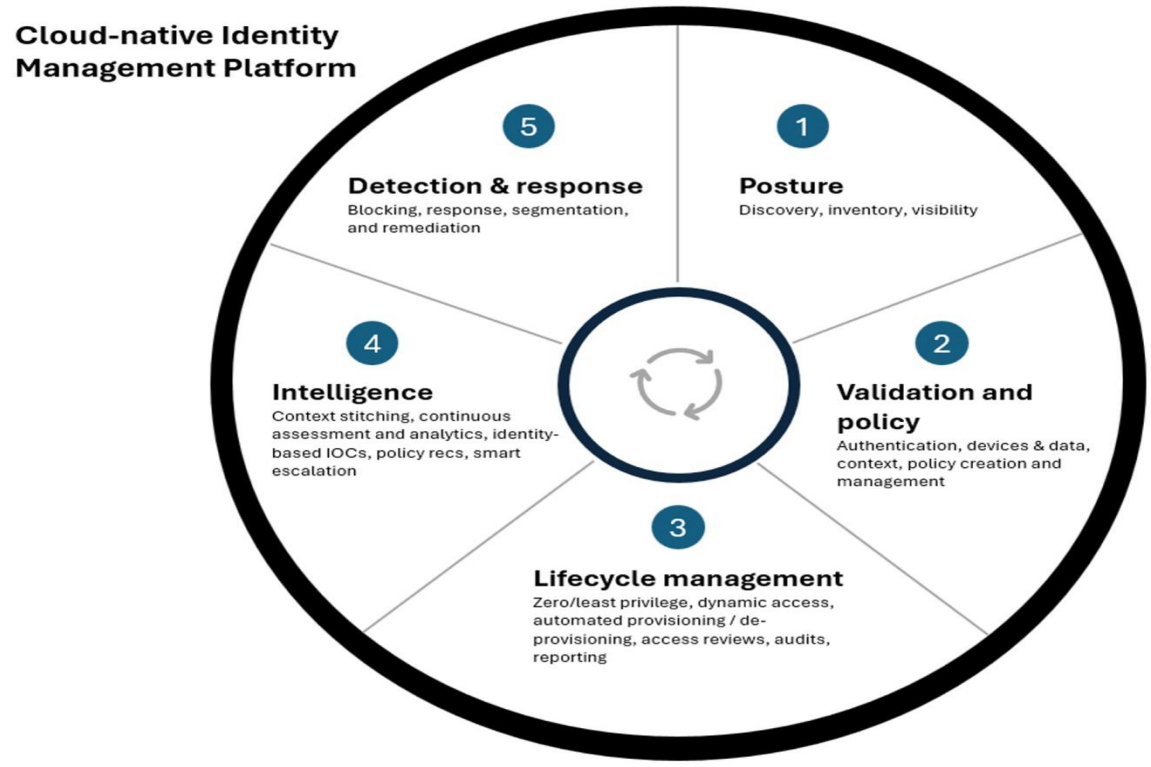

This model transitions governance from periodic human review to dynamic machine-driven control

5. **Literature Review**

Zero Trust Foundations and Operationalization: Zero Trust Architecture (ZTA) has become the dominant security paradigm for modern enterprises, shifting enforcement from perimeter controls to identity-, context-, and policy-driven access decisions. NIST SP 800-207 formally defines the canonical

ZTA model and emphasizes continuous verification and least privilege as foundational tenets. NIST SP 800-207A further extends these principles into access control modeling and policy evaluation mechanisms.

Government and public-sector frameworks, including CISA's Zero Trust Maturity Model and the U.S. Department of Defense Zero Trust Reference Architecture, operationalize these principles across identity, device, network, application, and data domains.

In industry, Google's Beyond Corp model demonstrated a large-scale implementation of identity-centric enterprise security, eliminating implicit trust based on network location.

Workload Identity Standards for Cloud-Native Systems: Workload identity standards such as SPIFFE and SPIRE provide mechanisms for issuing cryptographically verifiable identities to services and workloads. These frameworks enable mutual authentication without relying on long-lived shared secrets.

However, workload identity standards primarily address identity establishment and attestation. They do not inherently provide mechanisms for continuous privilege governance, entitlement minimization, or risk-adaptive enforcement across dynamic service graphs.

Cloud IAM Misconfiguration and Multi-Step Privilege Escalation: Research has extensively examined cloud IAM misconfiguration and privilege escalation. Recent USENIX Security work models AWS IAM policies using formal verification techniques to identify multi-step privilege escalation paths. Other work proposes automated repair of IAM privilege escalation vulnerabilities through graph analysis and optimization approaches.

These contributions significantly improve detection and repair of static policy vulnerabilities. However, they largely operate on static configurations and do not incorporate continuous behavioral telemetry or automated runtime governance.

Kubernetes RBAC Over-Permission and the NHI Attack Surface: Kubernetes RBAC has been shown to frequently grant excessive privileges to service accounts and third-party applications. CCS research has demonstrated that excessive permission assignments can lead to cluster compromise, establishing over-permission as a systemic cloud-native risk. More recent work proposes automated detection techniques for excessive RBAC privileges in Kubernetes environments.

These findings reinforce the prevalence of privilege sprawl in machine identities. However, most approaches focus on detection rather than continuous, telemetry-driven privilege adaptation.

Policy Mining and Data-Driven Least Privilege: Policy mining research seeks to derive least-privilege policies from observed access logs. Work on mining attribute-based access control (ABAC) policies demonstrates the feasibility of data-driven least-privilege generation.

While promising, traditional policy mining assumes relatively stable identity populations. Extending these approaches to ephemeral machine identities requires handling identity churn, evolving service graphs, and continuous telemetry ingestion

Risk-Adaptive Access Control and Continuous Decision-Making: Risk-Adaptive Access Control (RAdAC) proposes incorporating real-time risk into authorization decisions. This model aligns conceptually with continuous governance by framing access as a function of contextual risk rather than static policy.

However, modern cloud-native systems require extending risk-adaptive principles beyond decision-time access control to include privilege lifecycle management, credential TTL adaptation, and automated enforcement within distributed service architectures.

Telemetry-Driven Security Analytics and ML Caution: Machine learning approaches have been widely explored for anomaly detection in security systems. However, foundational work cautions against naïve deployment of anomaly detection without operational constraints, explainability, and integration into governance processes.

This motivates a hybrid approach that combines telemetry-driven analytics with bounded automation and enforceable guardrails.

Summary: Why Existing Literature Is Insufficient : Across Zero Trust frameworks, workload identity standards, IAM misconfiguration analysis, Kubernetes privilege research, policy mining, and risk-adaptive access control, the literature provides essential building blocks. However, a unified, continuous governance loop for non-human identities remains underdeveloped.

This paper addresses that gap by integrating behavioral telemetry, privilege modeling, credential lifecycle intelligence, and automated enforcement into a cohesive adaptive governance framework aligned with Zero Trust principles.

6. AI-Driven Continuous Governance Architecture

The AI-driven Continuous Governance Framework (ACGF) operationalizes machine identity governance as a closed-loop control system.
The architecture spans:

- Identity verification across human and non-human entities
- Device and workload posture validation
- Context-aware Zero Trust policy enforcement
- Service-to-service segmentation
- Runtime infrastructure protection
- Continuous threat intelligence integration
- Policy optimization through AI-based analytics

AI enhances governance by:

- Detecting anomalous identity behavior
- Identifying unused or excessive permissions
- Scoring identity risk based on behavioral and contextual signals
- Recommending privilege adjustments
- Automating remediation under controlled guardrails

AI-Driven Zero Trust Governance Architecture

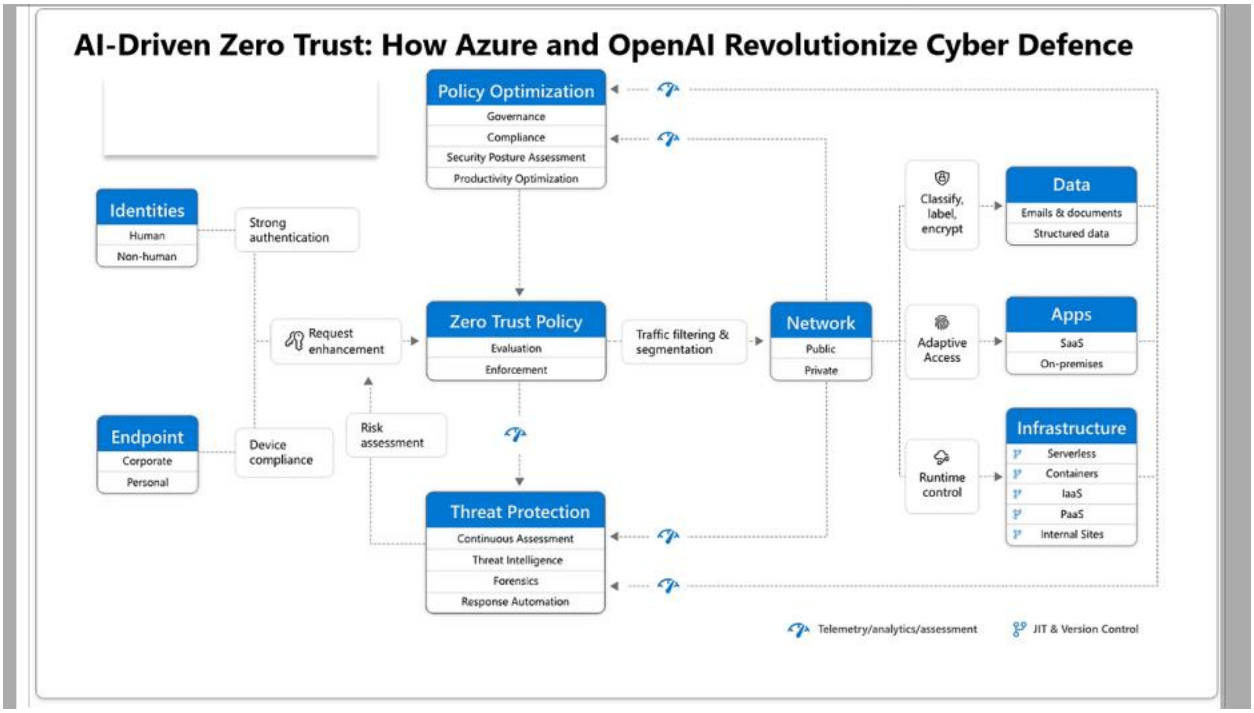

Unlike static IAM systems, ACGF continuously learns from telemetry and adapts policies dynamically.

7. Conclusion

Machine-dominated cloud-native systems require governance models that extend beyond periodic review and static IAM policies. While Zero Trust establishes identity as the primary control boundary, it does not prescribe continuous governance mechanisms for non-human identities operating at machine scale.
This paper presented an AI-driven Continuous Governance Framework that integrates behavioral intelligence, privilege modeling, and automated enforcement into a unified architecture. By transforming identity governance into a closed-loop control system, organizations can materially reduce privilege drift, lateral-movement exposure, and credential risk while maintaining operational stability.
As automation, microservices, and AI agents continue to expand, continuous machine identity governance will become foundational to secure Zero Trust deployments.

8. Reference:

1. National Institute of Standards and Technology, "SP 800-207A: A Zero Trust Architecture Model for Access Control," 2023.
2. Cybersecurity and Infrastructure Security Agency, "Zero Trust Maturity Model v2.0," 2023.
3. U.S. Department of Defense CIO, "Zero Trust Reference Architecture v2.0," 2022.
4. R. Ward and B. Beyer, "BeyondCorp: A New Approach to Enterprise Security," USENIX ;login:, 2014.
5. Shevrin, et al., "Detecting Multi-Step IAM Attacks in AWS Environments via Model Checking," USENIX Security Symposium, 2023.
6. Y. Hu, et al., "Fixing Privilege Escalations in Cloud Access Control with MaxSAT and Graph Neural Networks," Proceedings of the 38th IEEE/ACM International Conference on Automated Software Engineering (ASE), 2023.
7. S. Yang, et al., "Attacking Kubernetes via Excessive Permissions of Third-Party Applications," Proceedings of the 30th ACM Conference on Computer and Communications Security (CCS), 2023.
8. Z. Gu, et al., "EPScan: Automated Detection of Excessive RBAC Permissions in Kubernetes Applications," IEEE Symposium on Security and Privacy, 2025.
9. W. Sanders and M. Yue, "Mining Least Privilege Attribute-Based Access Control Policies," Annual Computer Security Applications Conference (ACSAC), 2019.
10. National Institute of Standards and Technology, "Risk-Adaptive Access Control (RAdAC) Concept Paper," [Online]. Available: https://csrc.nist.gov/publications/detail/white-paper/2010/07/01/risk-adaptive-access-control-radac-concept-paper/final
11. R. Sommer and V. Paxson, "Outside the Closed World: On Using Machine Learning for Network Intrusion Detection," IEEE Symposium on Security and Privacy, 2010.
12. SPIFFE Project, "SPIFFE Identity Specifications," [Online]. Available: https://spiffe.io/docs/latest/spiffe-about/overview/
13. SPIRE Project, "SPIRE Runtime Identity Framework," [Online]. Available: https://spiffe.io/spire/
14. K. R. Muppa, "Study on cloud-based identity and access management in cyber security," *International Journal of Data Analytics Research and Development (IJDARD)*, vol. 2, no. 1, pp. 40–49, 2024.
15. K. R. Muppa, "Analysis on the role of artificial intelligence and identity and access management (IAM) in cyber security," *International Journal of Artificial Intelligence Research and Development (IJAIRD)*, vol. 2, no. 1, pp. 113–122, 2024.
16. K. R. Muppa, "Analysis on cyber risk exposures and an evaluation of the elements that go into being ready to deal with cyber threats," *International Journal of Computer Engineering and Technology (IJCET)*, vol. 15, no. 3, pp. 12–20, 2024.
17. K. R. Muppa, "Enhanced identity and access management with artificial intelligence: A strategic overview," *International Journal of Information Security and Cybercrime (IJISC)*, vol. 13, no. 2, pp. 9–17, 2024.
18. K. R. Muppa, "Optimizing security in the cloud: Strengthening protection through single sign-on implementation," *International Research Journal of Engineering & Applied Sciences (IRJEAS)*, vol. 11, no. 2, 2023.
19. K. R. Muppa, "Advancing cloud security with AI-enhanced AWS identity and access management," *International Research Journal of Engineering & Applied Sciences (IRJEAS)*, vol. 10, no. 1, p. 4, 2022.

How do I cite this article?
Kaushik Reddy Muppa, AI-Driven Continuous Governance for Machine Identities in Cloud-Native Zero Trust Environments, Excel International Journal of Technology, Engineering and Management, 2026; Volume -13, Issue-2_Page_72-78. DOI Link: https://doi.org/10.64971/j.cph.eijtem.v12.i3.11.2025