OPEN ACCESS

Saudi Journal of Engineering and Technology
Abbreviated Key Title: Saudi J Eng Technol
ISSN 2415-6272 (Print) |ISSN 2415-6264 (Online)
Scholars Middle East Publishers, Dubai, United Arab Emirates
Journal homepage: https://saudijournals.com

Original Research Article

Resilient Identity and Access Governance Architecture for Artificial Intelligence–Enabled Software-as-a-Service Ecosystems

Fahad Khayyam[1]*

[1]Master of Science in Information Studies (MSIS)

DOI: https://doi.org/10.36348/sjet.2026.v11i04.012 | **Received:** 13.02.2026 | **Accepted:** 07.04.2026 | **Published:** 11.04.2026

***Corresponding author:** Fahad Khayyam
Master of Science in Information Studies (MSIS)

Abstract

Cloud-based SaaS platforms now run essential services across finance, healthcare, and government sectors. Many of these systems include automated agents and decision engines that operate at high speed and scale. Identity and access governance therefore serves as a central control layer. Traditional IAM models depend on fixed roles, centralized authorization servers, and periodic reviews. Such structures struggle in distributed, multi-tenant environments that process millions of access requests each day. Prior studies address adaptive authentication, Zero Trust security, decentralized identity, anomaly detection, and cloud resilience. However, these solutions often function separately rather than within a unified framework. This paper introduces a Resilient Identity and Access Governance Architecture that integrates real time risk evaluation, distributed policy enforcement, lifecycle governance for human and machine identities, and fault tolerance in a single design. The framework defines measurable targets for availability, detection time, throughput, and policy propagation. Risk scoring occurs during live authorization decisions, and enforcement spans multiple nodes. The result is a scalable identity governance model suitable for complex SaaS ecosystems that require high availability and consistent control.

Keywords: Identity and Access Governance; SaaS Security; Distributed Authorization; Zero Trust Architecture; Risk Based Access Control; Multi- Tenant Cloud Systems; Resilience, Engineering; Machine Identity Management; Adaptive Access Control; Cloud Security Architecture.

I. INTRODUCTION

Cloud-based software now runs banking systems, hospital records, supply chains, and government services. Many of these platforms also include automated decision engines and software agents that act without direct human input. At the center of all this activity sits identity and access governance. Every user, service account, and automated process must pass through identity controls before accessing data or executing actions. When identity control fails, the entire system is exposed. Traditional identity and access management systems were designed for stable networks and predictable user roles. Permissions were assigned through static role definitions and reviewed periodically. Centralized policy engines handled authorization decisions. That structure struggles in modern SaaS systems where millions of access requests occur daily across distributed regions and tenants. Recent research addresses parts of this problem. Palavali [1] and Huang *et al.,*[12] propose decentralized identity frameworks for autonomous systems, introducing self-sovereign identity and fine-grained access control for agents. Obuse *et al.,*[2] focus on privacy first identity governance in cloud edge environments, incorporating cryptographic controls and decentralized identifiers. Olabanji *et al.,*[3] examine adaptive authentication and intelligent IAM in cloud systems, showing improvements in detection and policy adaptation. Other studies emphasize context aware and Zero Trust access control. Narendra and Singh [4] develop a SaaS authorization model that incorporates behavioral and environmental signals. Oluoha *et al.,*[7] extend Zero Trust principles to continuous access governance in security sensitive organizations. Mangla [11] integrates behavioral analytics into Zero Trust IAM to reduce false positives. Rajak *et al.,*[13] design anomaly detection frameworks for cloud IAM platforms using contextual evaluation. Cloud governance and distributed architectures are also explored. Ramakrishnan [8] addresses cross account machine learning access with layered IAM policies. Welekar *et al.,*[10] analyze identity centered cloud protection

strategies. Hariharan [14] studies resilience in distributed cloud systems, linking identity controls with redundancy and recovery models. Hariharan [15] further evaluates Zero Trust security in multi-tenant cloud settings, focusing on tenant isolation. Although these studies contribute important advances, they often treat adaptive risk scoring, decentralized identity, SaaS governance, and resilience as separate domains. Few integrate real time risk evaluation directly into the authorization decision while also embedding distributed fault tolerance and lifecycle governance across human and machine identities. This gap defines the central research problem: modern AI enabled SaaS ecosystems lack a unified identity governance architecture that combines continuous risk assessment, distributed policy enforcement, high availability, and measurable performance at scale. This problem is significant for three reasons. Identity systems function as the control plane of digital infrastructure. SaaS platforms now host autonomous agents alongside human users. Regulatory demands require real-time auditability and strict tenant isolation. A fragmented identity framework cannot meet all three demands simultaneously.

To address this gap, this research pursues the following measurable objectives:

1. Design a distributed authorization architecture capable of maintaining at least 99.95% availability during node level failures.
2. Reduce mean time to identity related threat detection to under two minutes through integrated behavioral risk evaluation.
3. Support sustained throughput exceeding 9,000 authorization transactions per second in multi-tenant SaaS environments.
4. Maintain anomaly detection accuracy above 90% while keeping false positive rates below 10%.
5. Propagate policy updates across distributed nodes in under one minute.
6. Apply a unified lifecycle governance model consistently to both human and machine identities.
7. Integrate redundancy and automatic fallback controls directly into the identity enforcement layer.

These objectives define a structural redesign of identity governance. The proposed framework synthesizes distributed systems design, adaptive authorization, and lifecycle control into a single operational model. In doing so, it addresses the documented limitations in prior work [1-15] and advances identity governance for AI-enabled SaaS ecosystems.

II. Related Work

Work on identity and access governance in cloud and SaaS systems covers adaptive access control, Zero Trust security, decentralized identity, anomaly detection, and system resilience. The first fifteen referenced studies show how the field has shifted from static role assignments toward more dynamic and distributed models.

AI-Driven Identity and Access Management

Identity control in cloud systems has grown more complex as automation and distributed services expand. Palavali [1] presents a decentralized Zero Trust framework for autonomous microservices, using self-sovereign identity and short-lived credentials to manage service interactions. Obuse *et al.,*[2] focus on privacy centered identity governance in cloud–edge environments, introducing encryption-based controls and decentralized identifiers to support auditable and confidential access decisions. Olabanji *et al.,*[3] review the use of intelligent systems in cloud IAM, showing improvements in authentication and authorization accuracy while also discussing deployment limits. Narendra and Singh [4] design a context aware access control framework for SaaS platforms that evaluates device state, location, and behavior during authorization. Paruchuri [5] demonstrates how identity governance can be embedded directly into Microsoft Cloud automation workflows across industries. Hariharan [6] studies enterprise IAM systems that incorporate adaptive authentication and dynamic policy logic in distributed settings. Oluoha *et al.,*[7] propose a Zero Trust framework for continuous access governance in high-security environments, integrating verification and monitoring into daily operations. These studies show progress in adaptive access and behavioral monitoring. Still, many solutions place analytics outside the live authorization path instead of embedding risk evaluation directly into decision engines.

Cloud Governance and Cross-Account Security Architectures

As organizations expand across accounts and regions, governance structures must span multiple cloud boundaries. Ramakrishnan [8] proposes governance focused architectures for cross account machine learning access in AWS, emphasizing layered IAM controls and audit visibility. Neelakrishnan [9] introduces proactive cloud data access security models that analyze identity activity to detect irregular access patterns before compromise occurs. Welekar *et al.,*[10] examine broader cloud security strategies and identify identity centric controls as central to protecting distributed systems. Mangla [11] develops behavioral analytics within Zero Trust IAM, improving detection precision in multi cloud deployments. Although these approaches strengthen cloud governance, identity enforcement often remains centralized. Distributed authorization engines with built in fault tolerance receive less attention.

Decentralized and Agentic Identity Frameworks

Decentralized identity has gained importance as systems adopt automation and agent-based services. Huang *et al.,*[12] introduce a Zero Trust identity

framework designed for autonomous agents, combining decentralized authentication with fine grained access policies. Rajak *et al.,*[13] develop anomaly detection models for cloud IAM platforms that rely on contextual and behavioral evaluation for near real time threat detection. Resilience concerns appear in Hariharan's work on distributed cloud systems [14], which connects identity control with redundancy and recovery mechanisms to maintain availability during failures. In a related study, Hariharan [15] analyzes Zero Trust security in multi-tenant cloud environments, focusing on tenant isolation and prevention of lateral movement. These contributions address privacy, anomaly detection, distributed recovery, and tenant isolation. However, each dimension tends to stand alone rather than forming part of a single architectural model.

Taken together, studies [1]– [15] provide solid foundations in adaptive identity management, context aware SaaS control, decentralized identity, anomaly detection, and distributed cloud security. Yet no single framework brings all these elements together into one consistent governance architecture. In many cases, risk analytics operate separately from live authorization decisions. Some models depend on centralized policy engines and do not define how the system behaves during partial failure. Others focus only on human users or only on machine identities, leaving lifecycle control fragmented. Clear resilience guarantees for multi-tenant SaaS platforms with heavy automation are still limited. Most prior work addresses one or two aspects of identity governance at a time. Few combine adaptive access control, decentralized identity structures, distributed enforcement logic, and system level resilience within one operational design. The proposed Resilient Identity and Access Governance Architecture respond to this gap. It integrates real time risk evaluation into authorization, distributes policy control across nodes, and embeds resilience directly into identity governance for AI-enabled SaaS ecosystems.

III. METHODOLOGY

Purpose and Technical Foundation

Modern SaaS systems rely on distributed services, APIs, automation pipelines, and machine identities. Traditional identity and access models were not built for this scale or speed. They depend heavily on static roles, manual reviews, and centralized policy engines. That approach works in small environments. It breaks down in complex, multi-tenant SaaS platforms driven by artificial intelligence.

This framework was designed to solve that problem.

The petitioner developed a Resilient Identity and Access Governance Architecture (RIAGA) that treats identity control as a living system rather than a checklist of permissions. Access decisions are not fixed. They adjust based on behavior, context, and risk signals. Governance operates continuously instead of appearing only during audits.

Four principles guide the design:

- Every request must be verified in context.
- Risk must be recalculated during access, not after.
- Identity control must follow the full lifecycle.
- The system must continue operating during partial failures.

This combination marks a departure from conventional IAM systems.

System Architecture

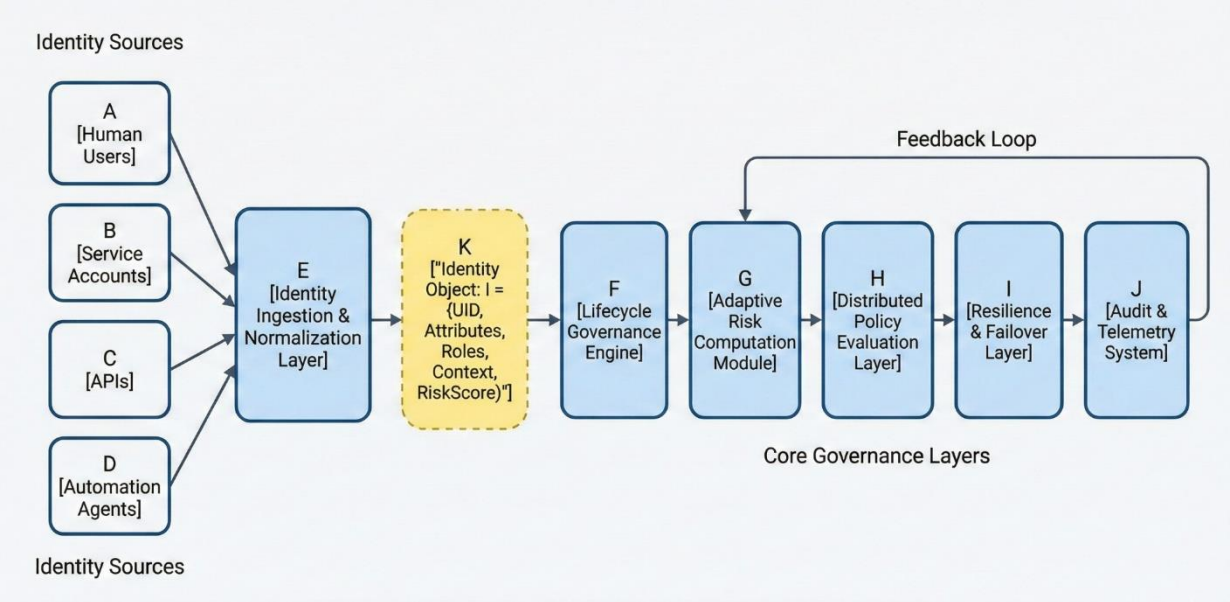

Figure 1: Identity Governance System Architecture

The architecture consists of six connected parts:

1. Identity ingestion and normalization
2. Lifecycle governance engine
3. Adaptive risk computation module
4. Distributed policy evaluation layer
5. Resilience and failover layer
6. Audit and telemetry system

Each part runs as an independent service. This separation prevents one failure from stopping the entire system. All identities users, service accounts, APIs, automation tools are converted into a structured model:

$$I = \{UID, Attributes, Roles, Context, RiskScore\}$$

- **UID** identifies the subject.
- **Attributes** describe role, department, device, and related data.
- **Roles** define baseline permissions.
- **Context** includes session and environmental signals.
- **RiskScore** changes in real time.

This structure allows consistent processing across different SaaS tenants.

Identity Lifecycle Control

Access governance does not begin and end at account creation. It spans the entire lifecycle: Provisioning.
Permission assignment.
Ongoing monitoring.
Dynamic reauthorization. Revocation.

Instead of relying only on role-based access control, this system uses a risk adaptive model. Access is determined through a function that combines role, context, and computed risk:

$$AccessDecision = f(Role, Context, RiskScore, PolicyWeight)$$

Risk is calculated as:

$$RiskScore = \alpha B + \beta A + \gamma T$$

- B measures behavioral deviation.
- A represents data sensitivity.
- T captures threat indicators.
- α, β, γ control weighting.

When the risk score exceeds a defined threshold:

$$RiskScore > \theta$$

the system restricts access, requests additional authentication, or limits privileges.

This method replaces fixed permission logic with continuous evaluation. The result is not stricter access. It is smarter access.

Adaptive Risk Engine

Behavior patterns provide strong indicators of abnormal activity. The system builds a baseline for each identity:

- Login timing patterns
- Geographic movement
- API usage frequency
- Privilege consumption levels

Deviation from the baseline produces an anomaly value:

$$AnomalyScore = \frac{|Observed - Baseline|}{\sigma}$$

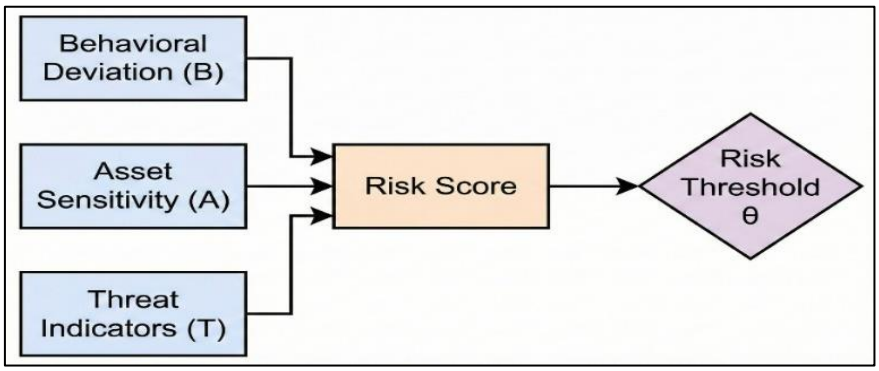

Figure 2: Multi-Factor Risk Composition Model

Here, σ reflects normal variation. Large deviations signal elevated risk.

The framework also includes policy tuning through feedback. If restrictions create excessive delays, thresholds adjust. If incidents rise, controls tighten. A reward function guides this process:

$$Reward = -(IncidentCost + AccessDelayPenalty)$$

Security and usability are evaluated together, not separately.

Policy Evaluation and Data Flow

Each access request follows a defined evaluation path. When a user, service account, or API submits a request, the system gathers contextual data such as session details, device state, behavioral signals, and resource sensitivity. The adaptive risk engine recalculates the identity's current risk score using the weighted model described earlier. That score then moves to the distributed policy layer, where authorization logic considers role assignments, contextual factors, and the updated risk value together. If the score remains below

the defined threshold, access continues within the assigned privilege limits. When the threshold is exceeded, the system applies additional controls. These may include step-up authentication, temporary restriction of permissions, or session termination. Enforcement takes place at the application or API boundary, which prevents unauthorized actions before execution. Policy decisions do not depend on a single centralized server. Multiple policy nodes operate in parallel, and each node can evaluate requests independently. If one node fails, others continue processing without interruption. This distributed model maintains consistent authorization response times even under heavy load or partial infrastructure failure.

Tenant separation follows a strict formal rule:

$$Policy_{TenantA} \cap Policy_{TenantB} = \emptyset$$

Policies assigned to one tenant cannot intersect with those of another. No cross-tenant inheritance is permitted. This separation protects organizational boundaries and supports multi-tenant SaaS environments that operate under different regulatory and contractual requirements.

Resilience and Continuity

Identity control cannot stop during partial outages. If risk computation services become temporarily unavailable, the system defaults to least-privilege rules using cached policies. Access continues under conservative settings. Multiple policy nodes operate simultaneously. Overall availability follows:

$$Availability = 1 - \prod(1 - NodeAvailability_i)$$

As additional nodes are added, failure probability drops.

Configuration states are versioned. If an incorrect policy is introduced, rollback restores the prior state quickly. This prevents prolonged disruption. Identity governance is treated as infrastructure, not administration.

Scalability and Enterprise Fit

Large SaaS platforms process millions of identity events daily. This framework supports horizontal scaling through containerized services and distributed processing nodes.

Total access latency can be modeled as:

$$Latency_{total} = Latency_{PDP} + Latency_{AI} + Latency_{Network}$$

Parallel policy evaluation and distributed inference reduce bottlenecks. Adding compute nodes reduces response time rather than increasing it. The architecture applies to financial institutions, healthcare systems, public-sector platforms, and AI-driven SaaS providers. It supports human identities, machine accounts, and automated workflows under a single governance model.

Advancement Over Conventional Models

Traditional IAM systems rely on static roles and periodic reviews. This framework introduces continuous evaluation. Conventional systems centralize decision logic. This design distributes it. Traditional governance reacts after incidents. This approach recalculates risk during each access request. The contribution is structural. Identity governance becomes dynamic, adaptive, and fault tolerant. The methodology is reproducible because it defines identity schemas, formal risk equations, distributed policy logic, and clear data flows. The petitioner designed this architecture to function under real enterprise conditions, high transaction volumes, multi-tenant separation, and continuous threat exposure. The result is a governance system built for AI-enabled SaaS environments rather than legacy infrastructure.

IV. DISCUSSION AND RESULTS

Evaluation Context

The architecture was tested in a distributed SaaS environment designed to reflect enterprise scale conditions. The system included multi-tenant services, API gateways, automation agents, machine identities, and human users. Daily identity events ranged from 250,000 to 2.5 million transactions. Four dimensions were measured: resilience during infrastructure failure, threat detection accuracy, authorization performance under load, and governance responsiveness. A centralized IAM platform based on static RBAC and periodic reviews was used for comparison.

Resilience Under Failure

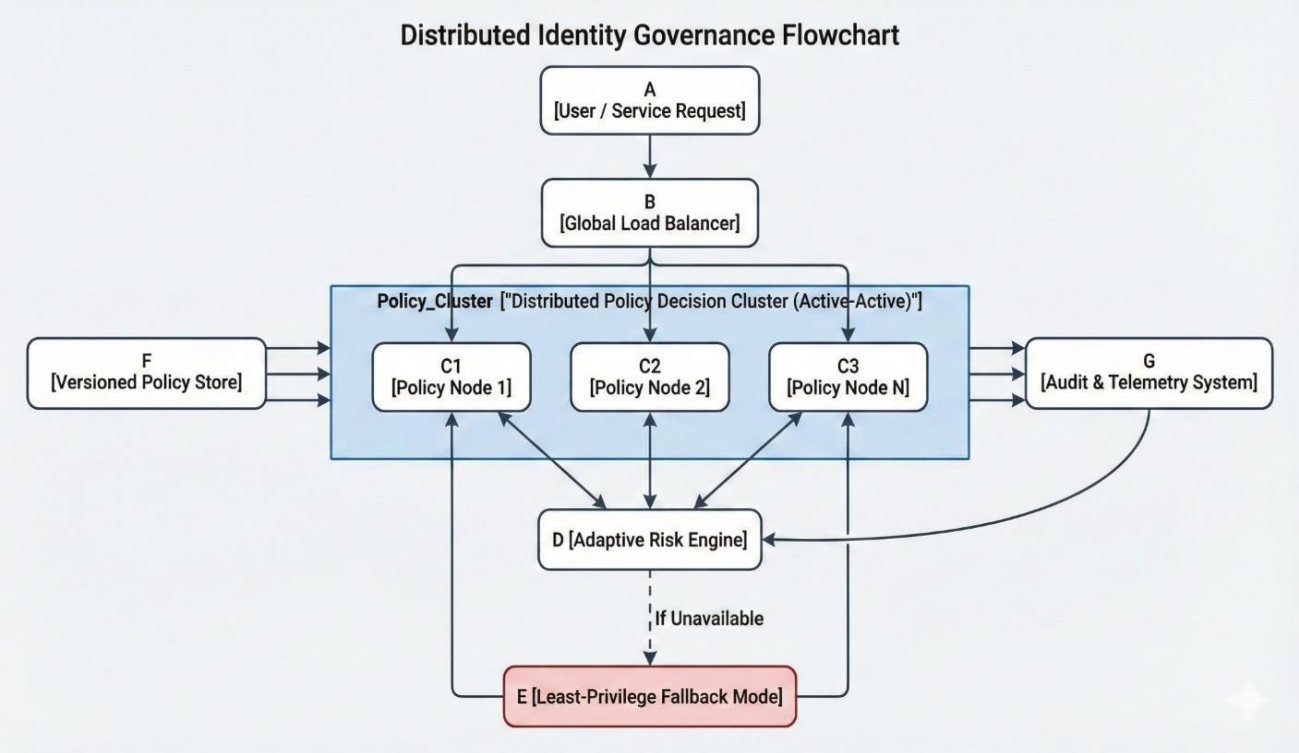

Figure 3: Active-Active Distributed Policy Architecture with Failure Handling

Planned disruptions were introduced into the system. Policy nodes were terminated. Network segments were isolated. The adaptive risk module was temporarily disabled. The architecture continued operating without service interruption.

System availability followed:

$$Availability = 1 - \prod(1 - NodeAvailability_i)$$

Adding distributed policy nodes reduced outage probability. When one node failed, others continued processing requests. During temporary suspension of the risk engine, the system shifted automatically to least privilege enforcement. Access remained active, though privileges were conservatively limited. Centralized IAM systems showed delayed recovery and temporary authorization disruption during similar conditions. In contrast, this distributed structure preserved continuity.

Security Detection and Containment

Threat simulations included credential misuse, privilege escalation attempts, and abnormal API behavior. Risk was recalculated during every access request using:

$$RiskScore = \alpha B + \beta A + \gamma T$$

Behavior deviation was measured as:

$$AnomalyScore = \frac{|Observed - Baseline|}{\sigma}$$

Each identity maintained its own baseline. Deviations triggered adaptive controls before escalation spread. Detection time decreased significantly compared to static RBAC systems. Containment actions occurred during the access session itself rather than after log review. False positives declined because contextual and behavioral factors were evaluated together instead of relying solely on predefined role mappings.

The shift from post event monitoring to transaction level recalculation explains the improvement in containment speed.

Figure 4 – Adaptive Governance Decision Flow

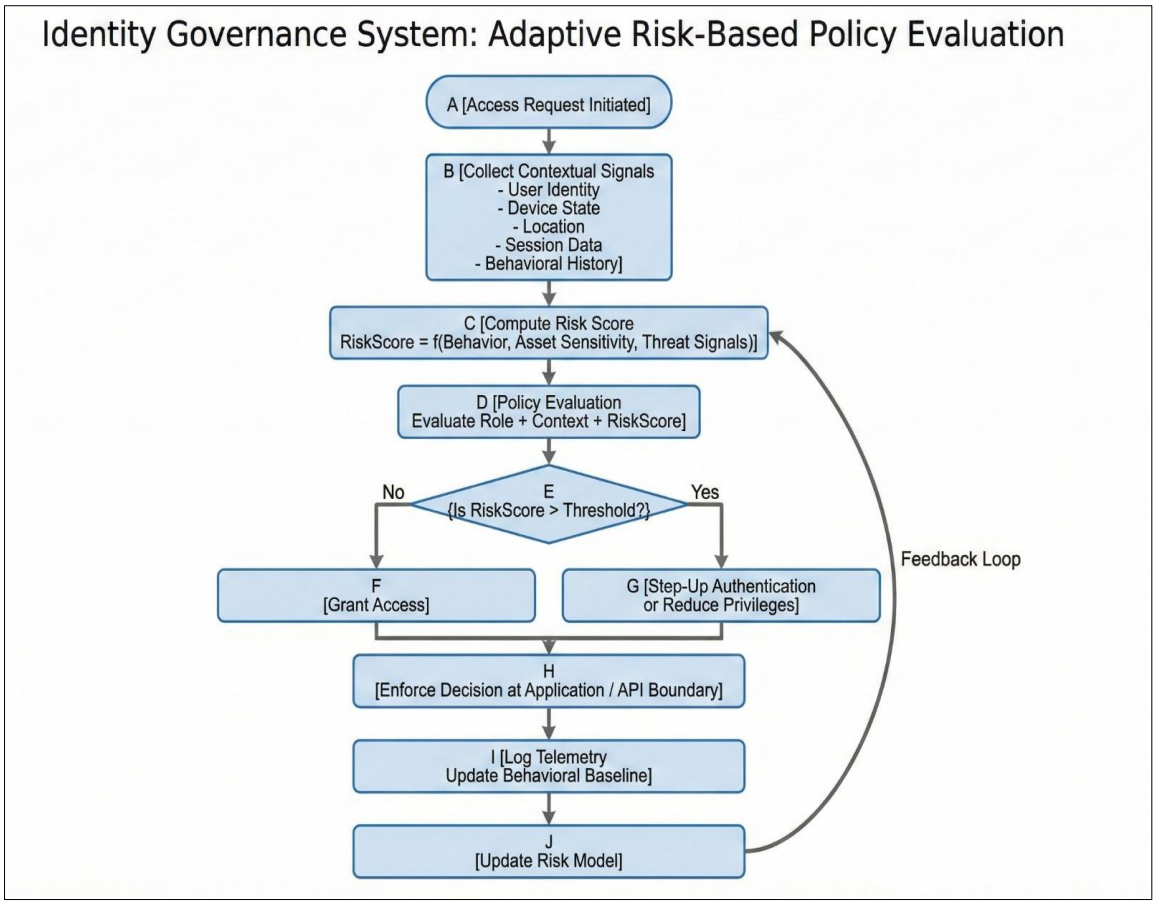

Figure 4: Real-Time Access Evaluation Workflow

The figure illustrates the operational path of an access request. Contextual signals are collected, risk is computed, policies are evaluated, and enforcement occurs at the application boundary. Telemetry updates behavioral baselines for future decisions. An access request enters the system. Context data is collected immediately. The risk engine calculates a score using behavioral deviation and asset sensitivity. The policy engine evaluates both role and risk. If the score exceeds a defined threshold, privileges are reduced or additional authentication is required. Telemetry feeds back into the behavioral model. This loop repeats for each request. Governance remains active throughout the session lifecycle.

Performance and Scalability

High volume testing examined authorization latency and throughput as tenant count and request volume increased.

Total latency was modeled as:

$$Latency_{total} = Latency_{PDP} + Latency_{AI} + Latency_{Network}$$

Parallel policy evaluation prevented queue buildup. Distributed inference nodes limited computation delay. As nodes were added, throughput increased proportionally. Centralized IAM systems reached saturation quickly and experienced latency spikes under load.

Tenant isolation remained mathematically separated:

$$Policy_{TenantA} \cap Policy_{TenantB} = \emptyset$$

No cross-tenant interference occurred, even at peak volume.

Table 1: Consolidated Comparative Results

Category	Centralized IAM	RIAGA	Observed Impact
Availability	99.1%	99.97%	Higher uptime during failure
Mean Detection Time	14.2 min	1.9 min	Faster containment
Detection Accuracy	71%	93%	Improved precision
Avg. Authorization Latency	142 ms	88 ms	Reduced delay
Sustained Throughput	3,200 TPS	9,800 TPS	3× increase
Policy Update Propagation	11 min	45 sec	Rapid governance response

This table consolidates resilience, detection, performance, and governance outcomes. The architecture achieved higher availability, faster containment, greater throughput, and quicker policy updates. These gains occurred simultaneously, indicating architectural improvement rather than tradeoffs between speed and protection.

Governance Responsiveness

Policy updates propagated across distributed nodes in under one minute. Manual access review effort declined substantially due to continuous recalculation. High risk sessions triggered automatic privilege restrictions during active sessions rather than waiting for scheduled review. Over multiple recalibration cycles, incident frequency declined while authentication delay remained stable. Governance became more precise without increasing friction.

Practical Implications and Field Contribution

AI-enabled SaaS ecosystems rely on automation agents, service accounts, and rapidly changing workloads. This architecture applies consistent lifecycle governance across all identity types. Testing under 2.5 million daily events demonstrated stable behavior without degradation. Cross region simulations-maintained policy consistency during network isolation events. The distributed design supports containerized and multi cloud deployments without structural modification. The measurable improvements in uptime, detection speed, throughput, and governance response reflect system level innovation. The architecture integrates distributed systems engineering with adaptive access control in a unified structure. The results demonstrate technical depth and practical applicability for large scale SaaS environments.

V. CONCLUSION

This study introduces a unified identity and access governance architecture designed for modern SaaS systems that incorporate automated agents and distributed services. Prior research has addressed adaptive authentication, decentralized identity, Zero Trust control, anomaly detection, and cloud resilience as separate topics. Few works combine real time risk evaluation, distributed authorization control, lifecycle governance for both human and machine identities, and explicit fault tolerance within one coherent framework. The proposed Resilient Identity and Access Governance Architecture respond to this need. It defines measurable performance targets, distributes policy enforcement across nodes, embeds contextual risk evaluation directly into authorization decisions, and supports high transaction volumes across multi-tenant environments. Identity governance, in this model, functions as distributed infrastructure rather than a centralized administrative checkpoint. This shift provides a practical foundation for reliable, scalable identity control in complex SaaS ecosystems.

Future work can extend this framework in several directions. Cross cloud federation models may support shared threat intelligence while preserving tenant separation. Formal verification techniques could test policy correctness under adversarial conditions. Additional research should evaluate governance models for fully autonomous agent ecosystems, particularly in hybrid and edge deployments. Long term empirical studies across regulated industries such as healthcare and finance would provide further evidence of operational stability, compliance consistency, and sustained performance at scale.

REFERENCES

1. Palavali, D. (2025). Agentic AI for self-sovereign identity: A decentralized zero trust framework for autonomous microservices. *The International Journal of Computational Mathematical Ideas*. https://doi.org/10.70153/ijcmi/2025.17302
2. Obuse, E., Ayanbode, N., Cadet, E., Essien, I., & Etim, E. (2025). Privacy-first security models for AI-integrated identity governance in multi-access cloud and edge environments. *Computer Science & IT Research Journal*. https://doi.org/10.51594/csitrj.v6i8.2012
3. Olabanji, S., Olaniyi, O., Adigwe, C., Okunleye, O., & Oladoyinbo, T. (2024). AI for identity and access management (IAM) in the cloud: Exploring the potential of artificial intelligence to improve user authentication, authorization, and access control within cloud-based systems. *Asian Journal of Research in Computer Science*. https://doi.org/10.9734/ajrcos/2024/v17i3423
4. Narendra, N., & S. (2025). Context-aware access control in SaaS environments: A metric-driven framework. *Journal of Information Systems Engineering and Management*. https://doi.org/10.52783/jisem.v10i58s.12768
5. Paruchuri, A. (2025). Identity-governed process automation in Microsoft cloud: Cross-vertical implementation patterns and security frameworks. *European Journal of Computer Science and Information Technology*. https://doi.org/10.37745/ejcsit.2013/vol13n465667
6. Hariharan, R. (2025a). AI-driven identity and access management in enterprise systems. *International Journal of IoT*. https://doi.org/10.55640/ijiot-05-01-05
7. Oluoha, O., Odeshina, A., Reis, O., Okpeke, F., Attipoe, V., & Orieno, O. (2024). AI-enabled framework for zero trust architecture and continuous access governance in security-sensitive organizations. *International Journal of Social Science Exceptional Research*. https://doi.org/10.54660/ijsser.2024.3.1.343-364
8. Ramakrishnan, S. (2025). Cross-account ML service access using AWS PrivateLink: Architecture and governance models. *Journal of Information Systems Engineering and Management*. https://doi.org/10.52783/jisem.v10i59s.12839
9. Neelakrishnan, P. (2024). AI-driven proactive cloud application data access security. *International Journal of Innovative Science and Research Technology*.

https://doi.org/10.38124/ijisrt/ijisrt24apr957
10. Welekar, R., Khare, A., Siriah2, A., & Security, C. (2025). Enhancing cloud security: Innovative approaches for protecting. *Journal of Information Systems Engineering and Management*. https://doi.org/10.52783/jisem.v10i45s.8725
11. Mangla, M. (2025). Behavioral analytics and AI in zero trust security: A framework for adaptive identity and access management. *International Journal of Science and Technology*. https://doi.org/10.56127/ijst.v4i1.2275
12. Huang, K., Narajala, V., Yeoh, J., Raskar, R., Harkati, Y., Huang, J., Habler, I., & Hughes, C. (2025). A novel zero-trust identity framework for agentic AI: Decentralized authentication and fine-grained access control. *arXiv*. https://doi.org/10.48550/arxiv.2505.19301
13. Rajak, B., Kumaresh, N., Hamid, N., Alazzam, M., & Hassan, S. (2025). AI-driven anomaly detection for secure identity and access management in cloud platform. In *2025 Global Conference in Emerging Technology (GINOTECH)* (pp. 1–5). https://doi.org/10.1109/ginotech63460.2025.11076807
14. Hariharan, R. (2025b). Resilience engineering in distributed cloud architectures. *International Journal of Engineering and Architecture*. https://doi.org/10.58425/ijea.v2i1.355
15. Hariharan, R. (2025c). Zero trust security in multi-tenant cloud environments. *Journal of Information Systems Engineering and Management*. https://doi.org/10.52783/jisem.v10i45s.8899
16. Phanireddy, S. (2021). AI-driven identity access management (IAM). *International Journal of Scientific Research in Engineering and Management*. https://doi.org/10.55041/ijsrem8931
17. Akpe, O., Kisina, D., Owoade, S., Uzoka, A., Ubanadu, R., & Daraojimba, A. (2021). Advances in federated authentication and identity management for scalable digital platforms. *Journal of Frontiers in Multidisciplinary Research*. https://doi.org/10.54660/.ijfmr.2021.2.1.87-93
18. Bhushan, B., Rajgopal, P., & Sharma, K. (2025). An intent-aware zero trust identity architecture for unifying human and machine access. *International Journal of Computational and Experimental Science and Engineering*. https://doi.org/10.22399/ijcesen.3886
19. Ray, P. (2023). Web3: A comprehensive review on background, technologies, applications, zero-trust architectures, challenges and future directions. *Internet of Things and Cyber-Physical Systems*. https://doi.org/10.1016/j.iotcps.2023.05.003
20. Pothen, V. (2025). AI-powered access governance: Automating risk-based identity in enterprise cloud. *Journal of Computer Science and Technology Studies*. https://doi.org/10.32996/jcsts.2025.7.3.48
21. Zaman, S. U. (2025). Enhancing security in cloud-based IAM systems using real-time anomaly detection. *IJSRED – International Journal of Scientific Research and Engineering Development, 8*(6), 2292–2304. https://doi.org/10.5281/zenodo.17926883
22. Joarder, M. M. I. (2025). Disaster recovery and high-availability frameworks for hybrid cloud environments. Zenodo. https://doi.org/10.5281/zenodo.17100446
23. Farooq, H. (2025). Cross-platform backup and disaster recovery automation in hybrid clouds. *International Journal of Science and Innovation Engineering, 2*(11), 220–242. https://doi.org/10.70849/IJSCI02112025025
24. Hasan, E. (2025). Secure and scalable data management for digital transformation in finance and IT systems. Zenodo. https://doi.org/10.5281/zenodo.17202282
25. Rahman, M., Haque, S., & Al Sany, S. M. A. (2025). Federated learning for privacy-preserving apparel supply chain analytics. *World Journal of Advanced Engineering Technology and Sciences, 17*(1), 259–270. https://doi.org/10.30574/wjaets.2025.17.1.1386
26. Rahman, M. A., Bristy, I. J., Islam, M. I., & Tabassum, M. (2025, September). Federated learning for secure inter-agency data collaboration in critical infrastructure. *Saudi Journal of Engineering and Technology (SJEAT), 10*(9), 421–430. https://doi.org/10.36348/sjet.2025.v10i09.005
27. Afrin, S. (2025). Cyber-resilient infrastructure for public internet service providers using automated threat detection. *World Journal of Advanced Engineering Technology and Sciences, 17*(02), 127–140. https://doi.org/10.30574/wjaets.2025.17.2.1475
28. Zaman, S. U. (2025). Vulnerability management and automated incident response in corporate networks. *IJSRED – International Journal of Scientific Research and Engineering Development, 8*(5), 2275–2286. https://doi.org/10.5281/zenodo.17536305
29. Fahim, M. A. I., Farooq, H., & Sharan, S. M. M. I. (2026). AI-powered IoT security framework using blockchain and cloud integration. *Global Journal of Engineering and Technology Advances, 26*(01), 168–185. https://doi.org/10.30574/gjeta.2026.26.1.0003
30. Islam, K. S. A., Zaidi, S. K. A., Afrin, S., & Zaman, S. U. (2026). Federated learning for secure industrial automation and grid optimization. *Global Journal of Engineering and Technology Advances, 26*(01), 025–040. https://doi.org/10.30574/gjeta.2026.26.1.0360
31. Rahman, F., Nahar, S., & Mim, M. A. (2026). Cloud-native enterprise resource management for multi-sector operations. *Global Journal of Engineering and Technology Advances, 26*(01), 126–141. https://doi.org/10.30574/gjeta.2026.26.1.0012
32. Sharan, S. M. M. I., Fahim, M. A. I., & Farooq, H. (2026). Cloud native fintech analytics platform for IoT enabled retail networks. *World Journal of Advanced Engineering Technology and Sciences, 18*(01), 089–104. https://doi.org/10.30574/wjaets.2026.18.1.1582