International Journal of Computational and Experimental Science and ENgineering (IJCESEN)

Vol. 12-No.1 (2026) pp. 1211-1217
http://www.ijcesen.com

ISSN: 2149-9144

Research Article

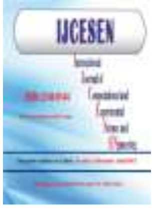

Research and Innovations in Identity and Access Management: Emerging Paradigms for Financial Services

Suneel Kumar Rawat*

Independent Researcher, USA
* **Corresponding Author Email:** rawatksuneel@gmail.com- **ORCID:** 0000-0002-5247-0772

Article Info:

DOI: 10.22399/ijcesen.4995
Received : 19 December 2025
Revised : 22 February 2026
Accepted : 01 March 2026

Keywords

Identity and Access Management,
Zero Trust Architecture,
Behavioral Biometrics,
Machine Identity Governance,
Decentralized Identity

Abstract:

Identity and Access Management has undergone a foundational architecture shift, based on the convergence of artificial intelligence, Zero Trust, and decentralized identity. Perimeter-based identity governance approaches, with static role provisioning and periodic access review, are structurally inadequate for the identity surfaces that modern financial enterprises operate today. Human users, service accounts, APIs, containerized workloads, and autonomous agents require context-sensitive, always-on access controls that run at a speed and granularity that existing access control and identity management technologies cannot accommodate. Seven innovation domains are changing the landscape of access control and identity management. Agentic artificial intelligence shifts access governance from reaction to prediction. Continuous behavioral authentication reduces implicit session trust. Just-in-time privilege reduces time of access to only what is needed to perform a task, while blockchain-anchored decentralized identity enables privacy-preserving credential verification. Machine identity governance extends least-privilege access to non-human identities. Platform convergence unifies fragmented access governance. Compliance frameworks require machine-speed policy enforcement. The financial services sector, with its heavily regulated environment, high-value targets, and desire for frictionless access experiences, serves as an unmatched proving ground for the approaches presented here. Some of the key open research questions are explainability for regulated decision trails, quantum-resistant crypto, privacy-preserving federated learning, and the measurement of return on investment.

1. Introduction

The modern financial enterprise identity landscape has vastly expanded beyond a small circle of human users and customary security infrastructure, accommodating contractors, service accounts, containerized workloads, APIs, and autonomous agents that fall into their own identity types, each governed by context-aware governance policies specific to their type. A bibliometric analysis of ten years of identity research in Scopus-indexed journals shows that access control, authentication, and privacy are the densest clusters, mirroring the growth of scholarly research on governance models for identity. This resulting surge in scholarship reflects the practical challenges to identity practitioners from the inability of static, role-based access control models in bounded perimeter environments to address the dynamic and multi-dimensional nature of access decisions in hybrid cloud environments.

The financial services sector is an acute example; credential theft is the highest attack vector for enterprises, and exposure to third parties increases the risk surface within businesses. In 2025, a study found that 35.5% of breaches were via third-party-based vectors; the abuse of credentials as a breach vector across all sectors was the highest of any confirmed breach vector. These findings point towards a structural issue whereby identity governance systems that stop at the enterprise boundary and do not monitor vendor, partner, or API access have been a source of opportunity for advanced attackers.

The modern IAM has evolved from a provisioning and authentication function to an all-inclusive control layer for all types of employees, contractors, service accounts, containerized workloads, bots, external partners, federated applications, and machine identities in hybrid and multi-cloud environments [2]. Three complementary technical forces drove this evolution. ZTNA's two primary innovations are (1) AI/ML and Zero Trust architecture that enables autonomous, contextualized decision-making and dynamic policy enforcement based on contextual parameters, behavior signals, device posture assessment, and threat intelligence while removing intrinsic trust and continuously verifying access to resources, and (2) decentralized identity, which uses cryptographically verifiable credentials and a distributed consensus to shift the control of identity from centralized identity providers and issuers to the individual while enabling privacy-preserving attribute assertions.

Regulatory requirements for explainability in strong domain applications can create friction for broad adoption of autonomous AI-driven access decisions. Financial regulators, for instance, require an auditable reasoning trail, which is not available at the same quality in current black-box models. At the moment, sensemaking for temporal access granularity is not fine-grained enough, which leads to standing privileges continuing to exist long after their operational need. Governance mechanisms for non-human identities (e.g., APIs, microservices, and autonomous agents) lack the maturity of those for human identities, despite non-human identities outnumbering human accounts in most enterprise environments [2].

The paper maps current research trajectories across seven innovation areas in IAM technologies and discusses the gaps between current capabilities and the operational and regulatory needs of the financial services industry. By drawing on empirical implementations and cross-domain bibliometrics of IAM research, we assess how emerging prototypes and frameworks currently being deployed in the financial domain meet existing research opportunities.

2. Related Work

Over time, IAM has transitioned from directory-based access control to a dynamic, intelligence-driven, risk-aware governance model, as previous models provided static identity governance based on roles and protection by a perimeter, implying trust once a user was previously authenticated. Subsequent research identified a number of shortcomings in those models (e.g., poor handling of credential-based attacks, privilege abuse, and lateral movement vectors), which led to the development of Zero Trust architectural principles, which require continuous verification regardless of where the network or session begins. Together, these and related contributions in behavioral biometrics establish the technical feasibility of continuous transparent identity verification during active user sessions across a variety of sensors and unimodal and multimodal systems. Parallel contributions in Just-in-Time privilege management formalize the temporal dimension of least-privilege enforcement and show how revocation of standing privileges to reduce exploit windows can be achieved efficiently. Thus, decentralized identity protocols, including blockchain and other distributed ledger technologies, are seen as mechanisms for cryptographically verifiable credentials, replacing central authorities and easing cross-boundary and privacy-optimized verification.

Table stakes like machine identity governance have emerged as a specialized contribution stream to address the challenges underlying the increasing lifecycle management of non-human identities. In aggregate, these contribution streams are transforming today's IAM from a collection of niche security controls into a flexible trust layer that manages access risk across the entire organizational identity surface.

3. Agentic Artificial Intelligence in IAM Operations

Agentic artificial intelligence fundamentally reorganizes the IAM automation landscape from rule-based automation to an agentic model, in which a clever agent does more than recommend or enforce policy but reasons over access behavior, detects privilege misuse, and remediates policy violations over a formally defined and population-level scoped access governance space, transforming IAM into a self-healing security control for the speed, scale, and complexity of modern financial infrastructure.

This problem is reflected in the numbers: the worldwide IAM market generated approximately USD 14.7 billion in 2022. Provisioning, or managing the lifecycle of user access and privileges, accounted for USD 4.3 billion of the IAM global revenue [3]. Market revenues for identity governance are projected to reach USD 22.2 billion in 2025 and USD 53.1 billion in 2032 at a 13.7% compound annual growth rate (CAGR), as organizations see a demand for advanced identity governance automation capabilities [3]. The most prominent operational model shift is agentic AI, with automation eliminating the need for human labor to continuously supervise the process.The short-term benefit of continuous access certification is that quarterly or annual access reviews may be insufficient in environments where privilege sprawl can happen between access reviews. Machine learning algorithms can be applied to entitlement usage data to determine inactive entitlements and anomalous access based on peer groups and to create automatic revocation of access when a specific condition is satisfied rather than at the next scheduled access review. Enterprise security research [4] estimates that it reduces the number of security incidents by 80%. This is consistent with the larger potential value of effective governance.

With agentic systems integration, transaction risk scoring engines in the financial services sector add an additional layer of precision to operational contexts. In high-value transactions, privilege elevation decisions are made at sub-second velocity based on behavioral baselines, a trust posture for a device, and real-time threat intelligence feeds. This architecture directly reduces the 74% of incidents involving human involvement (privilege abuse, stolen credentials, social engineering), of which three of four are directly identity-related [3]. Automated separation of duties violation detection, unlike retrospective audit, eliminates the window of time during which conflicting access combinations can be exploited by humans or human-controlled processes.

In regulated settings, audit-grade decision trails that express the reasoning behind a given action are required. Black box model outputs do not support auditability by design. This motivates research on human-in-the-loop hybrid systems, where agentic systems recommend remediation actions subject to requiring human endorsement for high-stakes actions, and constraint-based reasoning, where agentic systems are restricted to taking action on a procedurally verified set of procedures only. The experience of IAM as an institutional service shows that 92% of enterprises have adopted, or are planning to adopt, this technology. This reflects a demand that goes beyond any involvement of human supervisors. It is an open question how to reduce the risk of hallucinations in LLM components (by means of ensemble validation or other measures) before they are deployed in high-stakes operational environments governed by regulations.

4. Continuous Behavioral Authentication

In Zero Trust, continuous authentication is the verification of a user's identity during a session, as opposed to only at discrete points such as at each login. Behavior biometrics from user interactions may be used for continuous authentication. In reviewing over 140 approaches to continuous user authentication based on behavioral biometrics, six modalities can be identified: motion (28), gait (19), keystroke dynamics (20), touch gesture (29), voice (16), and multimodal (34) [5]. Taken together, this shows the maturity of this field as well as the variety of sensor input modern platforms have available.

By contrast, machine learning models based on behavioral signals can create a per-user model and operate transparently, without user initiation or reauthentication. For example, keystroke dynamics have been shown to be up to 99% accurate using Support Vector Machine classifiers. An EER of around 1.42% and a FAR of 2% has been recorded experimentally [5]. Touch gesture methods can attain a true positive of 99.9% and a 99.9% accuracy, as well as an Equal Error Rate of 0.01% when tested on a slide gesture dataset [5]. Multimodal systems improve robustness; for example, in a dataset including 35 users, the use of sensors embedded in the wearables and in the smartphone has caused an accuracy of 98.1% with an FRR of 0.9% and a FAR of 2.8%. On the other hand, a dataset including 1513 users with an SVM-based system reached an accuracy of 95.57% and a 3.2 s authentication time. An important factor is that in the European Union, the maximum FAR for commercial authentication systems is 0.01%, which current behavioral implementations are being optimized for.

Unlike knowledge-based authentication schemes, behavioral biometrics reduce the exploitable time window considerably. In a multimodal behavioral biometric system, the chance of performing any task without being detected is almost zero. An intruder will execute more than 1000 tasks if the security of a knowledge-based authentication scheme is broken, but in the multimodal behavioral biometrics case, the number is hardly one [5]. In a user perception experiment, 90% of the participants preferred transparent authentication based on behavioral biometrics to customary approaches, indicating a potential for common adoption [5].

Thus continuous authentication becomes a key requirement in a Zero Trust system. One approach to continuous authentication is to consider device identity as one of the factors in continuous authentication. The PUFDCA (Physical Unclonable Function-based Device Continuous Authentication) protocol can extend static PUF-based device authentication into continuous authentication using Channel State Information (CSI) to estimate device location at session start and during the session. Furthermore, the system has the property that the principal is authenticated only for the duration of the session, and reauthentication is required for each session (the principle of "never trust, always verify"). The system uses SHA-3 hash functions and a 128-bit Message Authentication Code, which has much less overhead as compared to the 128 to 256 bytes of overhead required for RSA signature schemes, thus making it suitable for resource-constrained devices [6]. The Scyther tool has been used to prove that the system is secure against man-in-the-middle, impersonation, replay attack, and physical attack [6].

5. Just-in-Time Trust and Decentralized Identity

5.1 Just-in-Time Privileged Access Management

"Always on" refers to privileged accounts, which are always typically in an active state. The 24/7 always-on state is typically used in enterprise identity management, describing the highly exploitable attack surface of always-on privileged accounts. For example, the always-on privileged account is privilege-active for 168 hours a week, regardless of whether legitimate activity requires that access level during any given period [7]. Just-in-Time Privileged Access Management (JIT PAM) is when privileges are only being granted when it is absolutely necessary to execute an authorized task. These privileges are revoked once the task is completed or the session expires. The goal of a JIT implementation is to achieve the Zero Standing Privileges (ZSP), where privileged access is never kept enabled within the enterprise accounts. [7]

The protection afforded by eliminating any privilege can be estimated. In a five-year analysis of vulnerabilities, 88% of critical vulnerabilities would have been removed if administrative rights had been removed from users, while 81% of vulnerabilities would have been eliminated had local administrator privileges been removed [7]. Time given to a privilege, in addition to its degree of privilege, is a major contributing factor to vulnerability. JIT PAM alleviates this problem by enforcing least privilege in both dimensions, by scoping the authentication to the minimum set of permissions required, and by constraining the duration to the minimum window of time required, unlike role-based PAM, which partially achieves this result.

Similar to Zero Trust, JIT PAM will take contextual information into consideration when privileges are activated, such as the authorization given for that task, the session context, and the assurance of identity. Because privileges are never persistent on an account and only checked out in a time-bound manner, JIT PAM eliminates the opportunity for lateral movement that on-demand access may provide, as there are no always-on privileged accounts for a threat actor to make use of to move around networked resources. [7] In financial services, for instance, using JIT PAM for trading system credentials, database administrative accounts, and payment processing interfaces limits the window of opportunity for credential abuse to the time required to perform authorized activities.

5.2 Decentralized Identity Frameworks

Blockchain-based decentralized identity systems supplement credential provenance and cross-boundary credential verification by addressing two of the most common failures of centralized identity architectures: a single point of failure and an organization's reliance on centralized identity providers whose security they cannot guarantee. In decentralized identity systems, cryptographically verifiable credentials are anchored in distributed ledger infrastructure to enable identity assertions from multiple sources not controlled by an authoritative issuer [8].

These approaches generally rely on the selective disclosure of verifiable credentials that allow the user to prove they possess some qualifying credential without revealing the underlying data. The financial services sector could adopt this for Know Your Customer (KYC) compliance, for example, allowing cryptographic attestation of eligibility to be done without full identity details being sent through the verification chain. Challenges exist to the adoption of blockchain-based identification systems, including scalability issues, energy consumed in the consensus mechanism, and interoperability across different distributed ledger technologies [8]. Few legal definitions exist for supporting on-chain identity claims. In the absence of legal clarity, intermediated designs are often proposed. They combine decentralized credential verification with required centralized oversight for access and audit needed to comply with laws and regulations. This practice has limited research into scalable decentralized identity in the context of financial services.

6. Machine Identity Governance and Platform Convergence

The true scale of the machine identity problem is structurally underappreciated, as machine identities (e.g., service accounts, application programming interface keys, secure shell keys, certificates, bots, and microservices) within enterprise environments far outnumber human identities (e.g., accounts for human users) by a ratio of 1 to 45 or even higher, but governance controls and processes are mainly human-centric [9]. The governance gap is reflected in adoption stats: 69% of organizations now have more machine identities than humans, while 72% say machines are harder to govern due to a gap in maturity for internal process and tooling [9]. This creates a structurally exploitable attack surface.

Overprivileged access is one of four threat vectors that together form the risk for machine identity management. Because machine identities are granted permissions that give them broad access based on the assumption that they always behave as desired, the blast radius becomes larger in case of a breach. The API key leak in 2024 stemmed from a lack of credential rotation and an overly permissive access policy [9]. In addition, if there is a lack of observability, there may be no detection of anomalous actions with a machine identity. For example, the 2023 Okta support system breach was enabled by the use of a machine identity [9]. Shadow machine identities are API keys, certificates, and service accounts invisible in IAM. Examples include the 2023 Microsoft SAS token leak, where a long-lived SAS token with no expiration or restrictions was leaked, giving access to sensitive production data, and the 2024 Internet Archive Zendesk exposure, where unrotated static tokens allowed attackers to gain access to hundreds of thousands of support records [9].

Automated lifecycle management reduces these attack vectors. Just-in-Time (JIT) and Just-Enough Privilege (JEP) grant fine-grained permissions to non-human identities for a limited time. These principles are based on the limited duration of the permissions, following which the permissions are revoked, thereby eliminating the vulnerability of standing privileges [9]. Automated and short expiry credential rotation limits static credential exposure that can be used for lateral movement to other hosts in the network. Continuous behavioral monitoring of requests, permissions, and resources accessed creates a baseline for every machine identity, enabling alerts based on unusual behavior, for example, when a CI/CD service account accesses production infrastructure during non-business hours [9].

Platform convergence helps to reduce governance fragmentation from administering cloud platforms, SaaS applications, databases, and CI/CD tooling as separate services with separate access control policies. A single centralized access control system with consistent least privilege policies for human and non-human identities helps reduce the risks from distributed, manual access control systems, like inconsistent policy enforcement and misconfigurations [9]. In financial services environments, where an API economy creates high volumes of programmatically generated machine identities through integrations between core banking systems and third-party fintech services, UGPs provide the cross-environment visibility and automated policy enforcement required to maintain security posture at cloud scale.

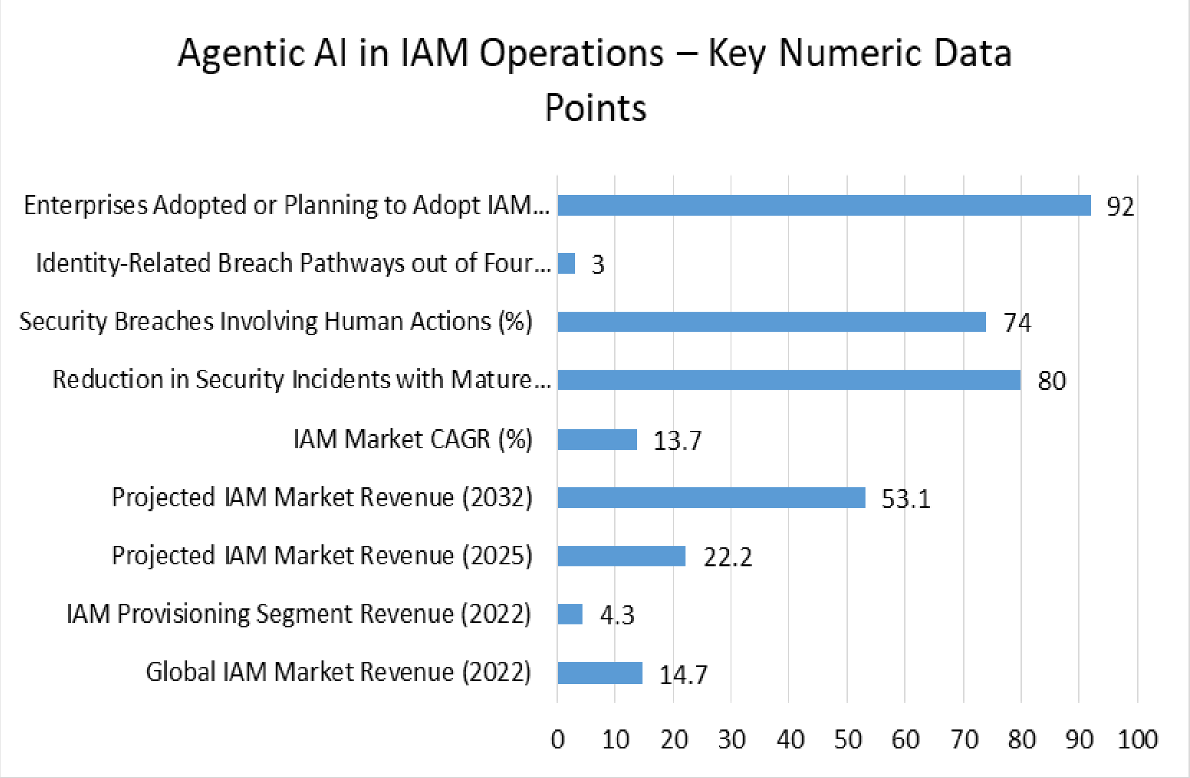

***Figure 1:** Agentic AI in IAM Operations—Key Numeric Data Points [3, 4].*

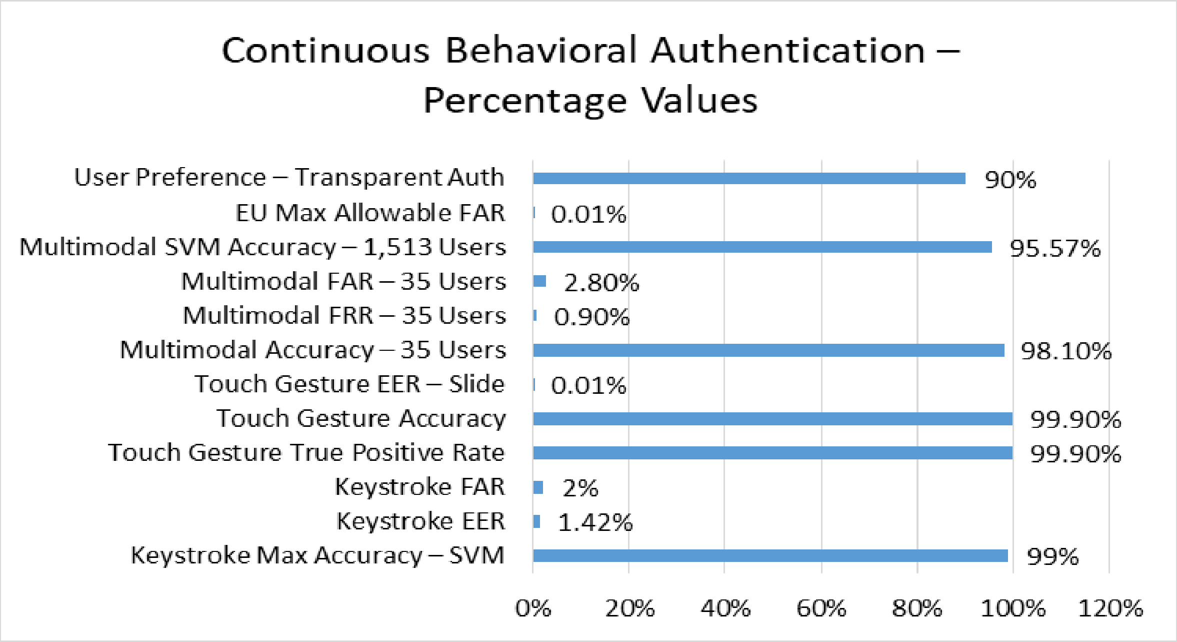

Figure 2: *Continuous Behavioral Authentication – Percentage Values [5, 6].*

Table 1: *Just-in-Time Trust and Decentralized Identity—Key Concepts [7, 8].*

Component	Mechanism	Financial Services Application
Just-in-Time PAM	Privileges granted only for authorized task duration and revoked upon completion	Trading system credential access scoped to transaction execution window
Zero Standing Privileges	Elimination of always-on privileged account states across enterprise	Database administrative accounts activated only during authorized operations
Continuous Adaptive Trust	Contextual evaluation of session state, identity assurance, and task authorization	Payment processing interface access constrained to verified transaction context
Decentralized Identity	Cryptographically verifiable credentials anchored in distributed ledger infrastructure	KYC compliance attestation without transmitting full identity documentation
Selective Disclosure	Proof of qualifying credential attribute without revealing underlying personal data	Regulatory eligibility verification across cross-border transaction chains

Table 2: *Machine Identity Governance and Platform Convergence—Key Concepts [9, 10].*

Threat/Component	Root Cause	Governance Response
Overprivileged Access	Broad permissions granted on assumption of predictable machine behavior	Just-in-Time and Just-Enough Privilege enforcement with auto-revocation
Lack of Observability	Absence of real-time monitoring and behavioral baseline tracking	Continuous audit logging with anomaly detection across all machine identities
Shadow Machine Identities	Untracked credentials operating outside IAM visibility with no expiration controls	Automated discovery, identity attribution, and inventory management across environments
Lifecycle Management Gaps	Static credentials without rotation, ownership, or consistent offboarding processes	Automated credential rotation with short validity periods and policy-enforced expiry
Governance Fragmentation	Siloed access control across cloud platforms, databases, and CI/CD tooling	Centralized unified policy engine applying consistent least-privilege controls across human and non-human identities

7. Conclusions

In financial services, agentic AI, continuous behavioral authentication, just-in-time privilege, blockchain-anchored decentralized credentials, and unified governance platforms have transformed Identity and Access Management from a compliance-driven authentication function into a clever adaptive security layer. This is one of the most consequential architectural shifts in enterprise cybersecurity. In this approach, therefore, access control is done at machine speed and in real-time and applies least-privilege security principles to each class of entity, whether human or non-human. It shows measurable improvement in fraud prevention, privilege reduction, and compliance effectiveness in the domains of innovation above when mature IAM architectures are deployed in each area. Machine identity governance is a structurally underaddressed risk domain where non-human identities already outnumber human identities by orders of magnitude, i.e., without governance of scale or behavior that would normally be applied to human identities. Gaps in explainable artificial intelligence for auditable access decisions, quantum-safe cryptographic primitives, and cross-institutional privacy-preserving intelligence are also the most pressing areas for further research. IAM's evolution from a necessary operational cost to a calculated enabler of a secure, efficient, and trustworthy digital financial ecosystem will be determined by how financial institutions operating across these gaps respond.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.
- **Use of AI Tools:** The author(s) declare that no generative AI or AI-assisted technologies were used in the writing process of this manuscript.

References

[1] Ahmad Ismail et al., "Mapping One Decade of Identity Studies: A Comprehensive Bibliometric Analysis of Global Trends and Scholarly Impact," MDPI, 2025. [Online]. Available: https://www.mdpi.com/2076-0760/14/2/92

[2] SecurityScorecard, "IAM in 2025: Identity and Access Management Best Practices," 2025. [Online]. Available: https://securityscorecard.com/blog/iam-in-2025-identity-and-access-management-best-practices/

[3] Tajammul Pangarkar, "Identity and Access Management Statistics By Security," Market.us Scoop, 2006. [Online]. Available: https://scoop.market.us/identity-and-access-management-statistics/

[4] Mary Marshall, "What is Identity and Access Management (IAM)? Complete 2025-2026 Guide for Enterprise Security," Avatier, 2025. [Online]. Available: https://www.avatier.com/blog/iam-complete-guide-for-enterprise-security/

[5] Mohammed Abuhamad et al., "Sensor-based Continuous Authentication of Smartphones' Users Using Behavioral Biometrics: A Contemporary Survey," arXiv, 2020. [Online]. Available: https://arxiv.org/pdf/2001.08578

[6] Shrooq Alshomrani and Shancang Li, "PUFDCA: A Zero-Trust-Based IoT Device Continuous Authentication Protocol," Wireless Communications and Mobile Computing Volume 2022, DOI: https://doi.org/10.1155/2022/6367579 [Online] Available: https://onlinelibrary.wiley.com/doi/pdf/10.1155/2022/6367579

[7] Matt Miller, "Just-In-Time Privileged Access Management (JIT PAM): The Missing Piece to Achieving 'True' Least Privilege & Maximum Risk Reduction," BeyondTrust, 2019. [Online]. Available: https://www.beyondtrust.com/blog/entry/just-in-time-privileged-access-management-jit-pam-the-missing-piece-to-achieving-true-least-privilege-maximum-risk-reduction

[8] Seyed Mohammad Hosseini et al., "Blockchain-Based Decentralized Identification in IoT: An Overview of Existing Frameworks and Their Limitations," MDPI, 2023. [Online]. Available: https://www.mdpi.com/2079-9292/12/6/1283

[9] Apono, "Machine Identity Management: How to Discover, Manage, and Secure," 2025. [Online]. Available: https://www.apono.io/blog/machine-identity-management/

[10] Identity Management Institute, "Quantum Threats to Identity and Access Management." [Online]. Available: https://identitymanagementinstitute.org/quantum-threats-to-identity-and-access-management/