Machine Identity Management in Modern Enterprise Security: Concepts, Challenges, and the Role of Privileged Access Management Systems

Erhan Yilmaz
Kron Technologies, Turkiye
erhan.yilmaz@itu.edu.tr (corresponding author)

Received: 11 November 2025 | Revised: 9 December 2025 and 24 December 2025 | Accepted: 3 January 2026

 DOI: https://doi.org/10.48084/etasr.16202

ABSTRACT

The rapid expansion of non-human entities across cloud platforms, microservices, IoT/OT devices, and automated deployment pipelines has positioned machine identities as a main element of enterprise security. These identities, instantiated through cryptographic credentials, such as certificates, SSH keys, and API tokens, enable authentication, authorization, confidentiality, and accountability in Machine-to-Machine (M2M) communication. However, their scale, high turnover, and architectural heterogeneity have outpaced traditional identity governance practices, leading to credential sprawl, inconsistent lifecycle management, and ineffective revocation mechanisms. This study examines the conceptual foundations and lifecycle requirements of machine identities, synthesizing recent research on certificate and key management in distributed environments. It evaluates the evolving role of Privileged Access Management (PAM) systems as policy-driven orchestration and governance layers for machine identities, particularly within zero-trust architectures. The analysis highlights both the strengths and limitations of current PAM implementations, identifying key research directions that include context-aware lifecycle automation, behavioral and attestation-based identity validation, governance for autonomous agents, and post-quantum secure identity infrastructures. Strengthening machine identity governance is, therefore, critical for ensuring the security and operational resilience of contemporary enterprise systems.

Keywords-machine identity; non-human identity; privileged access management; zero-trust architecture; secrets management; identity governance

I. INTRODUCTION

Contemporary enterprise systems are composed of non-human actors, including services, workloads, applications, containers, microservices, IoT/OT devices, CI/CD tools, secrets automation systems, and AI agents [1]. These actors must authenticate to one another to exchange data and perform privileged operations. The cryptographic identities that enable such authentication are commonly termed machine identities. These identities consist of persistent or ephemeral credentials that bind a non-human entity to verifiable cryptographic material and associated attributes, enabling authentication, authorization, confidentiality, and non-repudiation [2]. These types of identities are realized in the form of TLS certificates, SSH keys, API keys, and code-signing certificates [3]. The salience of machine identity management has grown as M2M communications, pervasive encryption, cloud-native architectures, and edge computing have driven an explosion in the number and heterogeneity of non-human principals [4]. This expansion gives rise to a range of governance and operational challenges, most notably including certificate and credential sprawl, inconsistent lifecycle practices, and a lack of effective revocation mechanisms [4, 5].

In this landscape, PAM platforms occupy a strategic integration point. Historically, PAM systems have been optimized for privileged human accounts [6]. However, contemporary PAM systems act as brokers, inventory managers, and rotators of machine secrets (e.g., host and service certificates, SSH keypairs, and API keys), enforcing policy over non-human sessions [4]. By centralizing credential issuance, storage, lifecycle automation, and access governance, PAM solutions can reduce the operational burden associated with managing high volumes of machine identities while also enforcing policy consistency across heterogeneous environments [7, 8]. PAM systems are often characterized as centralized governance components. In practice, however, contemporary deployments are typically implemented as distributed, highly available platforms that integrate with cloud-native services and orchestration frameworks [9, 10]. In this function, PAM operates as a policy coordination and identity orchestration layer rather than a single enforcement

point, a configuration that has been demonstrated to mitigate scalability limitations and single-point-of-failure concerns.

Furthermore, when integrated with Public Key Infrastructure (PKI), Certificate Authorities (CAs), and secrets management platforms, PAM systems can provide continuous verification, least-privilege enforcement, and comprehensive auditability for M2M authentication [5]. This shifts PAM from a traditionally human-centric control plane to a foundational trust anchor in zero-trust architectures where identity is treated as the primary perimeter. This study presents an examination of machine identity governance through four principal objectives. It establishes a clear definitional and architectural foundation for machine identities across cloud-native, hybrid, and industrial computing environments, explicitly distinguishing them from traditional human-centric identity models. Also, it synthesizes research findings and operational guidance from PKIs, device identity frameworks, and zero-trust security models to derive practical lifecycle and control requirements for non-human entities. It evaluates the extent to which contemporary PAM platforms address these requirements, with particular emphasis on certificate and key lifecycle automation, secure secret storage, least-privilege policy enforcement, and monitoring of non-human session activity. Finally, the paper proposes a forward-looking research agenda and evaluative criteria aimed at improving automation maturity, reducing credential sprawl, and enabling scalable, policy-driven identity governance for heterogeneous and dynamic machine populations.

II. CONCEPTUAL FRAMEWORK: MACHINE IDENTITY

Machine identity refers to the set of cryptographically verifiable attributes that uniquely distinguish a non-human entity, such as a device, workload, application, service account, or automated process, from other actors within a computational ecosystem [4]. In contrast to human identities, which are grounded in personal attributes and authenticated through biometric or knowledge-based factors, machine identities rely primarily on cryptographic key materials and associated metadata that bind the identity to a specific operational role or trust domain [2, 3]. The core purpose of machine identity is to enable secure authentication, authorization, confidentiality, integrity, and accountability in M2M interactions [11]. Machine identities possess several defining characteristics:

- Cryptographic Binding: A machine identity is anchored in a keypair or shared secret that permits verifiable authentication [12].
- Non-Interactive Authentication: Authentication typically proceeds without human mediation, often as part of automated or ephemeral workflows [11].
- Scope and Context Dependence: Identity meaning is relative to a trust boundary, such as a Kubernetes namespace, Active Directory domain, or cloud tenancy [2, 13, 14].
- Lifecycle Dynamism: Machine identities are frequently short-lived and tied to scaling events, deployment cycles, or workload scheduling decisions [4].

These characteristics differentiate machine identities from static device identifiers and configuration-layer metadata, positioning them as active components of security enforcement.

A. *Types of Machine Identities*

Machine identities manifest in multiple forms depending on the deployment model and underlying trust architecture. Device identities, for instance, are often anchored in hardware-backed secure storage, enabling attestation of device provenance during manufacturing, onboarding, and runtime operations [15]. Service accounts represent logical identities assigned to applications or backend services, allowing them to access infrastructure resources without human intervention [16]. Another prevalent category of machine identities consists of API keys and session tokens, which authorize automated interactions between software components in microservice-based architectures [17]. In addition, TLS/SSL certificates conforming to the X.509 standard are widely used for server, client, and workload authentication, ensuring encrypted communication and trustworthy endpoint verification [18]. Modern orchestration platforms, such as Kubernetes, further introduce ephemeral workload identities that are dynamically issued at container start-up and revoked upon termination [19]. Together, these identity types illustrate the heterogeneity of contemporary computing environments and underscore the difficulty of establishing unified lifecycle management and governance models. Figure 1 summarizes the primary categories of machine identities.

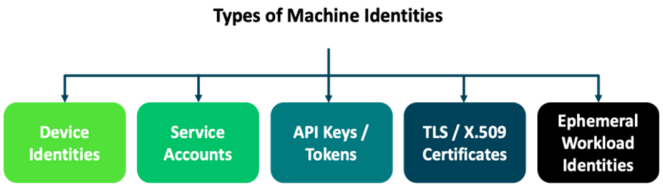

Fig. 1. Types of machine identities.

B. *Cryptographic Foundations*

The foundation of machine identity lies in the application of cryptographic systems designed to provide verifiable trust. PKI remains the dominant trust model, utilizing CAs, registration authorities, and signed certificates to establish hierarchical or mesh-based chains of trust [20, 21]. Within such infrastructures, secure storage of private keys is frequently delegated to Hardware Security Modules (HSMs) or Trusted Platform Modules (TPMs), which mitigate the risk of key extraction or impersonation attacks [22]. In addition to PKI-based models, token-based authentication frameworks, such as OAuth and OpenID Connect, provide standardized methods for distributing identity assertions across distributed systems. These systems are frequently employed in cloud-native architectures, where microservices rely on short-lived signed tokens to assert identity and scope of privilege [23]. Emerging identity systems, such as SPIFFE and its runtime component, SPIRE, extend this approach by assigning workload identities based on attested execution environments rather than static configuration artifacts [24, 25]. Collectively, these

cryptographic mechanisms provide the formal trust guarantees required to securely verify, delegate, and revoke machine identities across heterogeneous and distributed infrastructures.

C. Lifecycle of a Machine Identity

Machine identities progress through a lifecycle consisting of creation, distribution, operational use, rotation, and eventual retirement [26]. During creation, an identity is generated and bound to an entity based on established trust assumptions and security policies. The distribution phase concerns the secure delivery of key materials or certificates to the entity, often supported by enrollment protocols, provisioning workflows, or attestation-based onboarding [27]. Once deployed, the identity is actively used in authentication and authorization processes that support M2M communication, service invocation, or secure session establishment [28]. Over time, cryptographic best practices and risk management requirements necessitate periodic rotation or renewal of identity materials, ensuring that long-lived secrets do not become points of systemic vulnerability [29]. Finally, when a machine identity is no longer needed for workload, device, or privilege purposes, or if it is suspected of being compromised, it must be revoked and removed from trust registries. In distributed environments, effective revocation is operationally challenging and closely related to automation maturity [30]. Therefore, lifecycle management is a key factor in determining security posture in machine identity governance and plays a crucial role in reducing credential sprawl and unauthorized persistence [4]. Figure 2 offers a visual overview of the machine identity lifecycle.

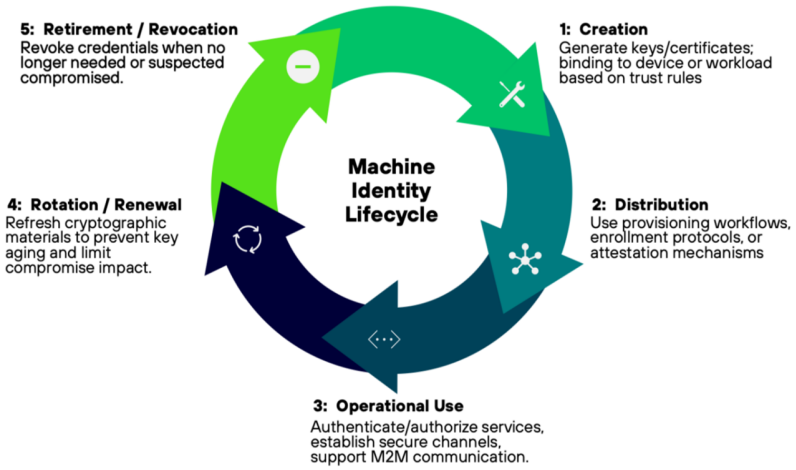

Fig. 2. Machine identity lifecycle.

III. BACKGROUND AND PREVIOUS WORKS

Research in the field of information technology has led to a growing recognition of the important role that machine identities have played in ensuring the security of communication and automated decision-making processes within contemporary computing environments. The studies outlined in Table I demonstrate that the shift from human-centric identity management to frameworks capable of governing non-human entities is both necessary and underway. The proliferation of distributed systems, cloud-native architectures, and autonomous workloads drives this transition. Each of these depends on cryptographic identities to establish trust and authorization boundaries. Several studies emphasize conceptual and architectural treatments of machine identity. Authors in [31] explore machine identity lifecycle management in cloud and AI-driven infrastructures, proposing organizational accountability models for non-human credential ownership. Similarly, authors in [32] argue for reframing identity governance around machine actors in M2M communication environments, noting that conventional access control practices remain biased toward human users. These contributions are important for establishing theoretical and organizational framing, though both studies acknowledge the absence of corresponding automation and enforcement mechanisms at the operational scale.

Other research focuses on technical innovations designed to replace or augment traditional key-based identity models. Authors in [11] propose a behavioral identity approach in which service identity is inferred from observed activity patterns rather than from stored cryptographic secrets. Although conceptually promising, the evaluation is limited to experimental deployments, leaving questions of scalability and feasibility in production cloud environments unresolved. Complementing this line of inquiry, authors in [33] analyze the practical risks associated with secret sprawl in DevOps pipelines and continuous deployment workflows, documenting recurring patterns of credential exposure and operational oversight. Their findings reinforce the need for automated, policy-driven control over both secret and certificate lifecycles. PKI lifecycle governance in distributed and device-centric environments has also been examined. Authors in [34] investigate certificate management practices in industrial IoT systems, identifying structural challenges in enrollment, renewal, and revocation workflows, challenges that closely parallel those observed in cloud-native service mesh environments. Extending this analysis, authors in [35] present empirical measurements showing that IoT vendors frequently

deploy weak, expired, or shared certificates at scale. Finally, authors in [36] propose a microservice-based PKI architecture for constrained environments, demonstrating efficiency improvements in device provisioning while leaving open questions regarding its applicability to highly dynamic and ephemeral workloads. These works illustrate the interplay between conceptual governance models, identity representation mechanisms, and the operational realities of credential lifecycle automation. They also indicate that while awareness of machine identity risks is increasing, systematic and scalable enforcement practices are still emerging. Table I depicts these contributions and their implications.

TABLE I. RESEARCH ON MACHINE IDENTITY AND CERTIFICATE/KEY LIFECYCLE MANAGEMENT

Reference	Study area	Key results / contributions	Limitations
[11]	Dynamic behavioral identity for services	Demonstrates identity derived from operational behavior rather than keys	Prototype-scale only; performance at cloud-scale not evaluated
[31]	Machine identity lifecycle governance in cloud/AI environments	Defines machine identity taxonomies and lifecycle risks; proposes governance framework	Mostly conceptual; lacks implementation validation
[32]	Non-human identity in M2M trust architectures	Argues for shifting IAM security boundary from humans to machine entities	Conceptual; lacks lifecycle automation discussion
[33]	Secure handling of machine secrets in DevOps pipelines	Identifies recurring key leakage patterns in CI/CD workflows	Addresses secret misuse, but not certificate lifecycle automation
[34]	TLS certificate lifecycle automation in industrial IoT networks	Shows lifecycle automation gaps in enrollment, renewal, revocation	Focuses on industrial IoT, not cloud-native ephemeral workloads
[35]	Real-world certificate security in IoT devices	Reveals widespread weak/expired/shared cert usage across vendor products	Observational; does not provide remediation strategies
[36]	Scalable PKI architecture for device identity	Proposes microservice-based PKI with demonstrated efficiency improvements	Evaluated mainly in constrained IoT; less tested in cloud-scale workloads

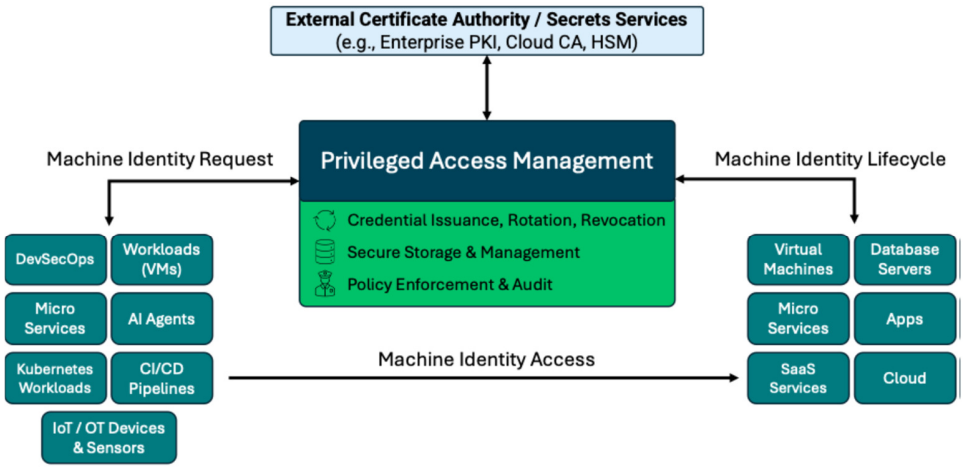

Fig. 3. PAM-centered framework for machine identity governance illustrating credential issuance, lifecycle automation, and policy enforcement across enterprise environments.

IV. ROLE OF PAM IN MACHINE IDENTITY GOVERNANCE

This study does not offer any experimental results, but the analysis employs a structured conceptual evaluation approach. Contemporary PAM systems are evaluated based on a set of analytically derived criteria, synthesized from representative academic studies, industry standards, and architectural practices related to machine identity lifecycle management and zero-trust architectures. These criteria include the degree of credential lifecycle automation across issuance, rotation, and revocation, as well as policy orchestration and least-privilege enforcement for non-human identities. Additional dimensions include scalability for ephemeral workloads, integration with cryptographic infrastructures (e.g., PKI and HSMs), and auditability of M2M interactions. This framework enables a systematic comparison of PAM capabilities in governing machine identities across heterogeneous enterprise environments. The following properties are treated as evaluation criteria for assessing the applicability of PAM platforms to machine identity governance, rather than as an exhaustive list of PAM features. Figure 3 illustrates a high-level conceptual framework in which PAM functions as a centralized policy orchestration and governance layer for machine identities, integrating identity issuance, lifecycle automation, and access enforcement across enterprise infrastructures.

A. Evolution of PAM from Human-Centric to Machine Identity Governance

PAM systems were initially designed to control, monitor, and audit the use of elevated permissions held by human administrators and operators [6]. Early PAM implementations primarily emphasized password vaulting, session recording, and policy enforcement to mitigate the risk of privilege misuse by individuals [37]. As machine identities have become integral to authentication, workload authorization, and service-to-service communication, the role of PAM has expanded

beyond its human-centric origins to encompass the governance of non-human entities [38]. This evolution reflects a broader shift in security architecture identified in recent research, wherein the identity perimeter comprises distributed software components and automated decision-making systems rather than human actors alone [39, 40].

B. Credential Lifecycle Automation as a Core Governance Function

The core contribution of PAM to machine identity governance lies in the consolidation and automation of credential lifecycle functions. The reviewed studies indicate that unmanaged certificate and key sprawl, misconfigured identity provisioning workflows, and inconsistent revocation practices are among the primary sources of security risk in distributed environments. PAM platforms address these challenges by providing centralized storage for machine secrets (e.g., TLS keypairs, SSH private keys, API access tokens), thereby reducing the likelihood of uncontrolled duplication or informal sharing [4]. Integrations with CAs, HSMs, and cloud-native identity services further enable PAM systems to automate key and certificate issuance, rotation, and expiration, aligning operational behaviors with policy requirements [41].

C. Policy Enforcement, Observability, and Zero-Trust Alignment

PAM systems contribute enforcement and observability capabilities that are difficult to achieve through decentralized identity management alone. Machine identities are frequently employed in ephemeral contexts, such as containers, serverless workloads, or automated CI/CD tasks, where traditional session-based monitoring is impractical [4, 11]. By mediating access through short-lived, policy-bound credentials and generating continuous audit records of inter-service communication, PAM platforms help maintain accountability even when identities are transient [5]. This aligns closely with zero-trust principles, in which trust is not inherited but continuously re-established based on identity verification, environment state, and policy evaluation [24].

D. Behavioral Signals and Continuous Identity Validation

The behavioral identity model presented in [11] highlights the potential for identity governance frameworks to incorporate runtime service behavior as a validation signal. PAM systems are well positioned to operationalize such signals, as they already function at the policy enforcement layer between entities executing privileged actions and the systems they access. This positioning places PAM at the convergence of identity authentication, authorization decision-making, and runtime anomaly detection [42]. Consequently, PAM has the potential to evolve into a foundational control layer through which continuous identity verification can be applied consistently across diverse workload and infrastructure environments [5].

E. Limitations of Current PAM Implementations in Dynamic Environments

Limits to current PAM adoption for machine identities have also been indicated. Most existing implementations remain optimized for relatively static or semi-structured environments. While PAM platforms support certificate and key lifecycle automation, most implementations remain oriented toward relatively static or semi-structured machine identity contexts [42, 43]. In contrast, modern service architectures often require identity issuance and revocation at speeds aligned with orchestration events, infrastructure-as-code pipelines, and autoscaling operations [44]. To fully assume the role of machine identity governance systems, PAM solutions must therefore integrate more deeply with deployment workflows and cloud-native identity providers, ensuring that identity lifecycle transitions occur automatically and at scale [45]. PAM systems represent a critical trust anchor for the governance of machine identities [46]. Their ability to centralize credential storage, enforce least privilege, automate lifecycle management, and produce auditable records enables them to mitigate several of the key risks identified in contemporary distributed computing environments [45]. However, the effectiveness of PAM in this role depends on continued alignment with cloud-native operational models and the incorporation of real-time identity validation mechanisms that support changing machine populations [4, 46].

V. FUTURE DIRECTIONS AND RESEARCH OPPORTUNITIES

The increasing centrality of machine identities within distributed computing environments introduces a set of research challenges that extend beyond current operational and architectural practices. While PAM platforms are evolving to accommodate non-human identities, the pace of change in cloud orchestration, autonomous systems, and machine learning pipelines continues to outstrip the maturity of identity governance capabilities. Regarding future research, a first direction concerns the automation and adaptability of lifecycle operations. Existing PAM systems can rotate credentials and renew certificates, but these mechanisms are generally triggered according to static schedules or manual workflows [42]. Future work should investigate event-driven and context-aware lifecycle management, in which identity issuance, renewal, and revocation are dynamically aligned with workload instantiation, scaling actions, policy changes, and observed behavior. This shift implies the need for tighter coupling between orchestration frameworks (e.g., Kubernetes), CA infrastructures, and PAM enforcement layers [5]. Second, there is a substantial opportunity in behavioral and attestation-based identity validation [47]. Integrating behavioral signals, code integrity proofs, runtime attestation, and anomaly detection into identity validation workflows may provide a foundation for continuous authentication. However, such approaches raise open questions regarding computational overhead, false positive rates, and the design of trustworthy baselines in heterogeneous and rapidly changing production environments [48]. Third, the emergence of multi-agent AI systems introduces new classes of machine actors whose behaviors, decision-making models, and execution contexts are not fully deterministic [49]. Traditional certificate-based identity models assume bounded signaling between known endpoints; agentic systems often involve learning processes, inter-agent negotiation, and self-directed task delegation [50].

Research is needed to determine how identity, authorization, and accountability can be enforced when agents act autonomously, including methods for traceable decision provenance and revocation of delegated authority. A fourth research frontier concerns Post-Quantum Cryptography (PQC) and its implications for machine identity scalability [51, 52]. Certificates and tokens will ultimately need to transition to quantum-resistant algorithms; however, PQC primitives introduce larger key sizes, increased computational overhead, and compatibility challenges for embedded and latency-sensitive systems [53]. Further research is therefore required to design hybrid identity architectures that enable secure machine operation during the prolonged transition period in which classical and post-quantum cryptographic infrastructures coexist. Finally, evaluation methodologies represent an underdeveloped dimension of machine identity research. Most existing studies emphasize conceptual models or limited-scale demonstrations, leaving organizations without evidence-based metrics for comparing governance approaches [31, 32]. There is a need for benchmarking frameworks, operational maturity models, and risk quantification techniques that can measure the security, reliability, and economic impact of machine identity governance choices across heterogeneous environments.

In addition, decentralized identity approaches, including blockchain-based and distributed ledger frameworks, have been proposed to address challenges related to trust distribution and credential transparency in machine identity ecosystems [54, 55]. While these approaches may complement centralized governance models, their operational complexity, performance overhead, and integration with existing enterprise identity infrastructures warrant further systematic evaluation [56]. Taken together, these research directions indicate that machine identity management is transitioning from an operational concern to a significant area of security architecture research [4]. Future advances will depend on interdisciplinary work spanning cryptography, distributed systems, cloud automation, behavioral modeling, and AI governance. The role of PAM in this evolution is likely to expand further, as these systems increasingly become the locus of policy enforcement and lifecycle orchestration across diverse categories of machine actors.

VI. CONCLUSIONS

The rapid expansion of non-human actors across cloud-native, automated, and distributed computing environments has shifted the locus of identity governance. Machine identities, in the form of certificates, keys, tokens, and other cryptographic credentials, now function at a scale and velocity that exceed those of traditional human accounts. The results of the present study demonstrate that the resulting security and operational challenges are not merely administrative in nature. Instead, they directly influence the integrity of authentication processes, the enforceability of authorization boundaries, and the reliability of service-to-service communication. The literature reflects an increasing recognition of the security and operational challenges introduced by large-scale machine identity deployments, particularly with respect to lifecycle automation, revocation, and alignment with dynamic workload behavior. In this context, Privileged Access Management (PAM) systems have begun to evolve from tools designed primarily for privileged human access into platforms capable of governing machine credential storage, rotation, and policy enforcement. Their positioning as mediators and auditors of machine identity usage enables them to function as strategic control points within zero-trust security architectures. However, their effectiveness depends on continued integration with orchestration systems and cryptographic infrastructures to ensure that identity governance remains adaptive, policy-driven, and resilient in highly automated environments.

This study suggests that organizations should treat machine identity governance as a primary security function. Security architects and platform teams should prioritize PAM solutions that integrate with Certificate Authorities (CAs), orchestration platforms, and deployment pipelines to enable automated and policy-driven lifecycle management. Assigning explicit accountability for machine identities, enforcing consistent credential controls, and maintaining centralized audit visibility are necessary to limit operational risk and sustain compliance in environments where non-human identities predominate. As machine populations continue to expand, investment in automation maturity and governance alignment becomes a critical managerial consideration. In summary, machine identity management has emerged as a core security capability at the intersection of organizational governance, distributed systems engineering, and cryptographic assurance. PAM platforms contribute to addressing these challenges; however, their long-term effectiveness will depend on sustained architectural adaptation and continued interdisciplinary research. Strengthening machine identity governance is therefore essential not only for mitigating existing operational risks, but also for enabling resilient, scalable, and secure digital infrastructures.

REFERENCES

[1] S. Syed, "Securing Non-Human Identities: Emerging Challenges and Innovative Solutions in Secret Management," *European Modern Studies Journal*, vol. 9, no. 4, Aug. 2025, https://doi.org/10.59573/emsj.9(4).2025.40.

[2] C. Stephanidis and G. Salvendy, *Human-Computer Interaction in Intelligent Environments*. USA: Taylor and Francis Group, 2025.

[3] M. Thelander, "Cheetahs, COVID-19 and the demand for crypto-agility," *Cyber Security: A Peer-Reviewed Journal*, vol. 4, no. 2, pp. 122–134, Dec. 2020, https://doi.org/10.69554/SNXG4181.

[4] Sudheer Kotilingala, "The non-human identity crisis: Managing machine identities in the modern enterprise," *World Journal of Advanced Research and Reviews*, vol. 26, no. 1, pp. 944–954, Apr. 2025, https://doi.org/10.30574/wjarr.2025.26.1.1118.

[5] S. Syed, "Zero Trust Principles and the Evolution of Privilege Access Management Architectures," *Journal of Computer Science and Technology Studies*, vol. 7, no. 7, pp. 859–865, July 2025, https://doi.org/10.32996/jcsts.2025.7.7.94.

[6] J. Garbis and J. W. Chapman, "Privileged Access Management," in *Zero Trust Security: An Enterprise Guide*, J. Garbis and J. W. Chapman, Eds. Berkeley, CA: Apress, 2021, pp. 155–161.

[7] M. J. Haber, "Privileged Attack Vectors," in *Privileged Attack Vectors: Building Effective Cyber-Defense Strategies to Protect Organizations*, M. J. Haber, Ed. Berkeley, CA: Apress, 2020, pp. 1–10.

[8] M. J. Haber, "PAM Architecture," in *Privileged Attack Vectors: Building Effective Cyber-Defense Strategies to Protect Organizations*, M. J. Haber, Ed. Berkeley, CA: Apress, 2020, pp. 173–188.

[9] S. Gadde, "Decentralized Identity Governance in Multi-Cloud Ecosystems: Challenges, Frameworks, and Future Directions," *European Modern Studies Journal*, vol. 9, no. 3, pp. 350–357, July 2025, https://doi.org/10.59573/emsj.9(3).2025.30.

[10] S. V. Anantula, "CHEZ PL: A Scalable Zero-Trust CIAM-PAM Architecture for Large Enterprises," *Journal of Computer Science and Technology Studies*, vol. 7, no. 5, pp. 328–333, June 2025, https://doi.org/10.32996/jcsts.2025.7.5.40.

[11] W. L. Teng and K. Rasmussen, "Actions Speak Louder Than Passwords: Dynamic Identity for Machine-to-Machine Communication," in *Proceedings of the 18th International Conference on Availability, Reliability and Security*, New York, NY, USA, May 2023, pp. 1–11, https://doi.org/10.1145/3600160.3600165.

[12] F. Corella and K. P. Lewison, "Identity-based protocol design patterns for machine-to-machine secure channels," in *2014 IEEE Conference on Communications and Network Security*, July 2014, pp. 91–96, https://doi.org/10.1109/CNS.2014.6997471.

[13] B. Burns, B. Grant, D. Oppenheimer, E. Brewer, and J. Wilkes, "Borg, Omega, and Kubernetes," *Communications of the ACM*, vol. 59, no. 5, pp. 50–57, Apr. 2016, https://doi.org/10.1145/2890784.

[14] C. Hickert *et al.*, "Trust Me, I'm Lying: Enhancing Machine-to-Machine Trust," in *2022 ACM/IEEE 13th International Conference on Cyber-Physical Systems (ICCPS)*, Feb. 2022, pp. 01–02, https://doi.org/10.1109/ICCPS54341.2022.00034.

[15] R. Gallo, H. Kawakami, and R. Dahab, "On Device Identity Establishment and Verification," in *Public Key Infrastructures, Services and Applications*, Berlin, Heidelberg, 2010, pp. 130–145, https://doi.org/10.1007/978-3-642-16441-5_9.

[16] K. I. Iyer, "From Logs to Intelligence: Leveraging Data Science for Service Account Monitoring," *Computer Fraud and Security*, pp. 1202–1209, Apr. 2025, https://doi.org/10.52710/cfs.638.

[17] A. Chatterjee and A. Prinz, "Applying Spring Security Framework with KeyCloak-Based OAuth2 to Protect Microservice Architecture APIs: A Case Study," *Sensors*, vol. 22, no. 5, Jan. 2022, Art. no. 1703, https://doi.org/10.3390/s22051703.

[18] A. Parsovs, "Practical Issues with TLS Client Certificate Authentication." 2013, [Online]. Available: https://eprint.iacr.org/2013/538.

[19] A. Tiwari, N. Singh, and A. Vashishth, "Using container orchestration for rapid ephemeral container use," in *Proceedings of the International Conference on Innovative Computing & Communication (ICICC) 2022*, Jan. 2022, https://doi.org/10.2139/ssrn.4020428.

[20] A. S. Wazan *et al.*, "Trust Management for Public Key Infrastructures: Implementing the X.509 Trust Broker," *Security and Communication Networks*, vol. 2017, no. 1, 2017, Art. no. 6907146, https://doi.org/10.1155/2017/6907146.

[21] J. Huang and D. M. Nicol, "An anatomy of trust in public key infrastructure," *International Journal of Critical Infrastructures*, vol. 13, no. 2–3, pp. 238–258, Jan. 2017, https://doi.org/10.1504/IJCIS.2017.088234.

[22] V. Gopal, S. Fadnavis, and J. Coffman, "Low-Cost Distributed Key Management," in *2018 IEEE World Congress on Services (SERVICES)*, July 2018, pp. 57–58, https://doi.org/10.1109/SERVICES.2018.00042.

[23] N. Naik and P. Jenkins, "An Analysis of Open Standard Identity Protocols in Cloud Computing Security Paradigm," in 2016 IEEE 14th Intl Conf on Dependable, Autonomic and Secure Computing, 14th Intl Conf on Pervasive Intelligence and Computing, 2nd Intl Conf on Big Data Intelligence and Computing and Cyber Science and Technology Congress(DASC/PiCom/DataCom/CyberSciTech), Dec. 2016, pp. 428–431, https://doi.org/10.1109/DASC-PICom-DataCom-CyberSciTec.2016.85.

[24] S. T. Avirneni, "Establishing Workload Identity for Zero Trust CI/CD: From Secrets to SPIFFE-Based Authentication." arXiv, Apr. 20, 2025, https://doi.org/10.48550/arXiv.2504.14760.

[25] H. Cochak, M. Neto, C. Miers, M. Marques, and M. Simplicio Jr., "Enhancing SPIFFE/SPIRE Environment with a Nested Security Token Model:," in *Proceedings of the 14th International Conference on Cloud Computing and Services Science*, Angers, France, 2024, pp. 184–191, https://doi.org/10.5220/0012634400003711.

[26] A. Cameron and O. Grewe, "An Overview of the Digital Identity Lifecycle (v2)," *IDPro Body of Knowledge*, vol. 1, no. 7, Feb. 2022, https://doi.org/10.55621/idpro.31.

[27] F. Martin-Tricot, C. Eichler, and P. Berthomé, "Secure key distribution in heterogeneous interoperable industrial Internet of Things," in *2021 17th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, July 2021, pp. 74–79, https://doi.org/10.1109/WiMob52687.2021.9606265.

[28] A. Esfahani *et al.*, "A Lightweight Authentication Mechanism for M2M Communications in Industrial IoT Environment," *IEEE Internet of Things Journal*, vol. 6, no. 1, pp. 288–296, Oct. 2019, https://doi.org/10.1109/JIOT.2017.2737630.

[29] B. Cimorelli, S. De, A. L. Ferrara, and B. Masucci, "Hierarchical Key Assignment Schemes with Key Rotation," *29th ACM Symposium on Access Control Models and Technologies (SACMAT 2024)*, pp. 171–182, 2024, https://doi.org/10.1145/3649158.3657037.

[30] L. Zhang, G.-J. Ahn, and B.-T. Chu, "A rule-based framework for role-based delegation and revocation," *ACM Trans. Inf. Syst. Secur.*, vol. 6, no. 3, pp. 404–441, May 2003, https://doi.org/10.1145/937527.937530.

[31] A. Wairagade, "Machine Identity Security in Cloud & AI: Ensuring Lifecycle Management, Ownership, and Accountability for Non-Human Identities," *International Journal of Computer Trends and Technology - IJCTT*, vol. 73, no. 2, pp. 80–89, Feb. 2025, https://doi.org/10.14445/22312803/IJCTT-V73I2P110.

[32] O. O. Aramide, "Securing Machine-to-Machine Communications in the Age of Non-Human Identities," *International Journal of Technology, Management and Humanities*, vol. 9, no. 04, pp. 94–117, Dec. 2023, https://doi.org/10.21590/ijtmh.2023090408.

[33] A. Rahman, F. L. Barsha, and P. Morrison, "Shhh!: 12 Practices for Secret Management in Infrastructure as Code," in *2021 IEEE Secure Development Conference (SecDev)*, July 2021, pp. 56–62, https://doi.org/10.1109/SecDev51306.2021.00024.

[34] J. Göppert, A. Walz, and A. Sikora, "A Survey on Life-Cycle-Oriented Certificate Management in Industrial Networking Environments," *Journal of Sensor and Actuator Networks*, vol. 13, no. 2, Apr. 2024, Art. no. 26, https://doi.org/10.3390/jsan13020026.

[35] H. Dong *et al.*, "Behind the Scenes: Uncovering TLS and Server Certificate Practice of IoT Device Vendors in the Wild," in *Proceedings of the 2023 ACM on Internet Measurement Conference*, New York, NY, USA, July 2023, pp. 457–477, https://doi.org/10.1145/3618257.3624815.

[36] S. Sumaidaa, H. Almenhali, R. Ramasamy, O. Voronin, M. Alazzani, and K. Han, "Securing the Device Lifecycle Management: A Scalable and Cost-Efficient Public Key Infrastructure Through Microservices:," in *Proceedings of the 11th International Conference on Information Systems Security and Privacy*, Porto, Portugal, 2025, pp. 342–352, https://doi.org/10.5220/0013171700003899.

[37] M. Nardone, "PAM, Protecting Privileged Accounts and Access Management," in *IAM and PAM Cybersecurity: Securing Identities and Access Management in the Digitalization Era*, M. Nardone, Ed. Berkeley, CA: Apress, 2025, pp. 33–73.

[38] A. Koot, "Introduction to Privileged Access Management (v2)," *IDPro Body of Knowledge*, vol. 1, no. 15, Mar. 2024, https://doi.org/10.55621/idpro.101.

[39] P. Kumar, "Next-generation secure authentication and access control architectures: advanced techniques for securing distributed systems in modern enterprises," *International Journal of Computational and Experimental Science and Engineering*, vol. 11, no. 3, July 2025, https://doi.org/10.22399/ijcesen.3294.

[40] S. L. Narra, "Human-AI Collaboration in Identity Security: When Should AI Decide?," *Journal of Computer Science and Technology Studies*, vol. 7, no. 7, pp. 191–197, July 2025, https://doi.org/10.32996/jcsts.2025.7.7.17.

[41] H. Tuononen, Privileged access management model for a managed service provider. JAMK University of Applied Sciences, 2023.

[42] Vinay Vasanth, "Advancing Enterprise Security: A Framework for AI-Powered Privileged Access Posture Management," *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, vol. 11, no. 1, pp. 88–95, Jan. 2025, https://doi.org/10.32628/CSEIT25111202.

[43] M. Kunz, L. Fuchs, M. Hummer, and G. Pernul, "Introducing Dynamic Identity and Access Management in Organizations," in *Information Systems Security*, Cham, 2015, pp. 139–158, https://doi.org/10.1007/978-3-319-26961-0_9.

[44] A. Vidal, P. H. Gomes, and M. Santos, "Reorchestration: a Reactive Orchestration Architecture," in *2019 IEEE Conference on Network Softwarization (NetSoft)*, June 2019, pp. 498–505, https://doi.org/10.1109/NETSOFT.2019.8806677.

[45] S. A. G. Rao, "Risk Reduction at Scale: Economic Impacts of Automated Identity Governance in Cloud Enterprises Amid Agentic AI Advancements," *Global Business & Economics Journal*, July 2025, https://doi.org/10.70924/f83n6wqz/5ylfneaq.

[46] Sushant Chowdhary, "Protecting the Digital Ecosystem: AI's Dual Role in Machine Identity Security," *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, vol. 11, no. 1, pp. 1986–1996, Feb. 2025, https://doi.org/10.32628/CSEIT251112200.

[47] S. C. Thurupati, "The Evolution of Identity and Access Management: Integrating Biometric and Behavioral Authentication," *IJFMR - International Journal For Multidisciplinary Research*, vol. 6, no. 6, Nov. 2024, https://doi.org/10.36948/ijfmr.2024.v06i06.29986.

[48] H. Ozkan, F. Ozkan, I. Delibalta, and S. S. Kozat, "Online anomaly detection with constant false alarm rate," in *2015 IEEE 25th International Workshop on Machine Learning for Signal Processing (MLSP)*, Sept. 2015, pp. 1–6, https://doi.org/10.1109/MLSP.2015.7324320.

[49] H. Li, Y. Liu, J. Yan, J. Gao, and X. Yang, "Position: Emergent Machina Sapiens Urge Rethinking Multi-Agent Paradigms." arXiv, July 01, 2025, https://doi.org/10.48550/arXiv.2502.04388.

[50] K. Huang *et al.*, "A Novel Zero-Trust Identity Framework for Agentic AI: Decentralized Authentication and Fine-Grained Access Control." arXiv, May 28, 2025, https://doi.org/10.48550/arXiv.2505.19301.

[51] A. Tsili, K. Kordolaimis, K. Krilakis, and D. Syvridis, "A Scalable Framework for Post-Quantum Authentication in Public Key Infrastructures." arXiv, Apr. 16, 2025, https://doi.org/10.48550/arXiv.2504.12062.

[52] L. H. Mahdi and A. A. Abdullah, "Fortifying Future IoT Security: A Comprehensive Review on Lightweight Post-Quantum Cryptography," *Engineering, Technology & Applied Science Research*, vol. 15, no. 2, pp. 21812–21821, Apr. 2025, https://doi.org/10.48084/etasr.10141.

[53] M. Raavi, P. Chandramouli, S. Wuthier, X. Zhou, and S.-Y. Chang, "Performance Characterization of Post-Quantum Digital Certificates," in *2021 International Conference on Computer Communications and Networks (ICCCN)*, July 2021, pp. 1–9, https://doi.org/10.1109/ICCCN52240.2021.9522179.

[54] M. Alizadeh, K. Andersson, and O. Schelén, "Comparative Analysis of Decentralized Identity Approaches," *IEEE Access*, vol. 10, pp. 92273–92283, 2022, https://doi.org/10.1109/ACCESS.2022.3202553.

[55] Bhaskara Garnimitta, "Blockchain-Enabled Decentralized Identity Management: A Novel Framework for Microservices Architecture," *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, vol. 11, no. 1, pp. 1929–1936, 2025, https://doi.org/10.32628/cseit251112209.

[56] B. Alangot *et al.*, "Decentralized Identity Authentication with Auditability and Privacy," *Algorithms*, vol. 16, no. 1, Jan. 2023, Art. no. 4, https://doi.org/10.3390/a16010004.