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

Vol. 12-No.1 (2026) pp. 1170-1179
http://www.ijcesen.com

ISSN: 2149-9144

Research Article

AI Identity for Agentic Systems: Integrating PAM and IAM for Secure Autonomy

Sushant Chowdhary*

Dr. A.P.J. Abdul Kalam Technical University, India
* **Corresponding Author Email:** mrsushant.chowdhary@gmail.com - **ORCID:** 0000-0002-5247-1440

Article Info:

DOI: 10.22399/ijcesen.4985
Received : 29 December 2025
Revised : 20 February 2026
Accepted : 22 February 2026

Keywords

Agentic AI Identity,
Privileged Access Management,
Identity and Access Management,
Autonomous Systems Governance,
Non-Human Identity Management

Abstract:

Agentic AI systems, as autonomous agents that make decisions and take actions in enterprise environments, raise identity governance questions regarding agents that maintain state, collaborate, and request escalation of privileges to perform complex activities, different from the questions that human actors, service accounts, and API tokens raise in identity governance contexts. The article outlines the existing challenges of establishing customary IAM for agentic AI and proposes a thorough framework building on IAM and PAM elements for holistic oversight of agentic AI. This framework entails considering agentic AI as first-class identity subjects, being responsible for their lifecycle, and dynamically provisioning their credentials. It also provides context-aware policy enforcement and continuous behavior monitoring. Augmenting IAM's identity controls with PAM's privileged operations governance satisfies operational autonomy and security requirements, increases accountability and compliance, and ensures the trustworthiness of autonomous AI operations while countering privileged credential theft, privilege escalation, and audit trail deficiencies stemming from customary identity-centric approaches to PCC management of autonomous AI operations.

1. Introduction

Agentic AI are autonomous agents that execute actions and make decisions on enterprise applications, cloud services, and application programming interfaces. Agentic AIs are different from human users who have a consistent identity and predictable access patterns. Agentic AI also introduces new non-human identities that can dynamically change in ways that violate security assumptions. As agentic systems become ubiquitous, from automated workflow orchestrators to decision-making systems, governance models will need to balance security, accountability, and compliance while enabling effectiveness. Agentic AI systems can autonomously sense and reason via global optimization and perform tasks across diverse domains such as healthcare, finance, manufacturing, and smart cities with minimal human intervention [1]. Currently designed with human actors with fixed roles and supervised actions in mind, identity models do not cater to autonomous agents that reside in distributed systems, need fast, real-time responses, and require elevated privilege for complex, value-creating work. Authentication, authorization, and accounting mechanisms struggle to be effective with agentic systems. 80% of security incidents result from a breach of privileged account credentials [2] and require implementing credential lifecycle management, session recording, and just-in-time provisioning of privileged privileges. Privileged identity management is essential for organizations. This article presents a holistic PAM and IAM framework to address the challenges presented by agentic AI identities from the perspective of identity management. It is argued to treat autonomous agents as first-class identity subjects, rather than as implicit system components, establishing oversight mechanisms that balance operational autonomy and security and privacy requirements.

2. The Challenge of Non-Human Identity in Agentic Systems

2.1 Defining Agentic AI Identity

Unlike customary software programs, which are characterized by logical control flows, agentic AI systems are autonomous, goal-directed, and use decision-making and workflow automation techniques to execute tasks [3]. LLMs like GPT-4, PaLM, LLaMA, etc. analyze the context to break down high-level goals and generate tasks to reach them. They adapt their performance based on their output and results, all without human intervention [3]. This level of autonomy raises questions of identity, not covered by service accounts or API tokens.

Apart from authentication credentials, agentic identity includes behavioral characteristics, rights to make decisions, risk profiles, and accountability chains [5]. Such agents have state across sessions, have long-term memory, can create sub-agents at runtime, and can coordinate with existing agents. In consequence, an identity framework should support decentralized orchestration and the ability to change tasks dynamically over time [3]. For example, to provide privileged access for short-lived operations, just-in-time provisioning of credentials can be enabled with verifiable machine identities and explicit trust anchors, rather than with static permissions [4]. Persistent and ephemeral machine identities bind non-human entities to verifiable cryptographic material; they are used in TLS certificates, SSH certificates, API keys, and code-signing certificates [4].

2.2 Risks of Inadequate Identity Governance

Because agentic systems do not impose clear identity governance, they are vulnerable to system-level attacks. It has been observed that over three of the harmful instructions are executed without proper sandboxing. However, with container-based sandboxing, almost all malicious commands are blocked, raising privacy and availability issues surrounding agent runtime environments [5]. However, the long lifetime of API keys was problematic, as agent processes were operating autonomously in distributed environments, creating a long-lived attack surface with potentially compromised agents [4].

Because agentic AIs are autonomous, credential exposure is harder to contain. Once they use the provisioned credentials, they will not have the same ability as people to detect suspicious prompts. They will be an attractive target [3]. First, systemic risks relate to agentic/fused systems' capabilities to operate/coordinate across agents. This can involve collusion amongst corrupted agents, cascading errors between systems, memory poisoning in cognition, or distributed execution patterns that are less visible to an observer [5]. Second, weak identity controls result in less strong accountability chains and a lack of visibility. This property can become a vulnerability if there will be increased use of forensic tools to determine attribution to identifiable agents [3].

However, the rapid growth of machine-to-machine communication, cloud-native architectures, and edge computing has driven the emergence of certificate and credential sprawl, inconsistent lifecycle practices, and weak revocation methodologies. These trends have led to an increase in the number and diversity of non-human principals, leading to governance challenges for customary identity and access management (IAM) [4].

2.3 Limitations of Traditional Access Models

Customary IAM technology (e.g., Microsoft Entra ID, Okta, ForgeRock) has evolved from username/password-based authentication to include more complex capabilities such as multi-factor authentication (MFA), single sign-on (SSO), and federated identity management. However, customary IAM technology is designed for humans, and although it can also support non-human entities such as service accounts, managed identities, and workload identities for machine-to-machine transactions, it is not suitable for supporting dynamic workloads and autonomous agentic decision-making.

Most RBAC models are based on static permissions. However, in agentic systems, permissions must be contextualized: they should be tailored to task requirements and environmental contexts, as well as risks [3]. The situation is aggravated when agents are able to create sub-agents, change their own tasks, or delegate their computing resources to external environments [3]. Static IAM policies are not sufficient. However, service accounts and API keys cannot provide the granularity needed for agentic operations to take place across multiple security domains within a single workflow, and having long-lived static credentials violates the principle of minimal credential exposure [4].

Existing models do not support periodic credential rotation, continuous verification, and revocation of credentials based on real-time risk signals, which would allow agents to operate independently and asynchronously without explicit human oversight [3, 4]. If they do not perform lifecycle management for non-human identities, unused credentials and permissions will amass, leading to credential sprawl and orphaned access. Most systems lack automated identity governance and policy enforcement and behavioral analytics conducive to agentic

environments [3]. PKI, CAs, and secrets management platforms integrated with PAM systems give continuous verification, least-privilege enforcement, and full auditability, achieving 75% better protection of static credentials, 60% better audit traceability, and 40% better anomaly detection accuracy based on a set of key performance indicators (KPIs) [3].

3. Identity and Access Management for Agentic AI

3.1 Establishing Agentic Identity as Governed Entities

If IAM for agentic AI systems is to treat them as governed identities, the identity lifecycle must include provisioning, modification, suspension, and deprovisioning. This is analogous to the identity lifecycle of human users, which IAM is concerned with [6]. Specifically, the three pillars of IAM (authentication, authorization, and identity lifecycle management) must be extended to include the needs of non-human actors [6]. Identity provisioning for agents should take into consideration authentication mechanisms like certificate-based authentication, biometric equivalent cryptographic authentication, machine-appropriate multifactor authentication methods, and operational metadata like purpose, scope, business function it belongs to, and risk classification of the specific agent [6].Each agent identity must be bound with attestations for business justification, entitlements, and accountability chains from direct agency to full accountability for a business unit. Identity registries must support polymorphic agentic identities to support a single logical agent with multiple runtime identities across different systems and varying levels of governance while still ensuring an identity correlation. Organizations that are deficient in lifecycle management can suffer from orphaned accounts, stale access, or failure to meet compliance requirements [6]. An effective provisioning workflow ideally requires the approval of both technical and business authorities, achieving separation of duties for agentic identities. This ensures that agentic capabilities are consistent with the organization's risk tolerance, as industry research shows that more than 80% of data breaches are tied to credential compromise and privilege abuse [6].

3.2 Role-Based Access Control for Autonomous Agents

Role-based authorization in an agentic system needs to be specified with sets of permissions to correspond to agent capabilities. AI-powered IAM automation also provides advantages for administrators, such as role-mining functions that recommend minimally privileged access by analyzing previous authorization decisions. This considerably reduces the risks associated with over-permissioned principals [6]. Create agent roles based on functional roles (data access, API invocation, workflow orchestration) and scope their permissions to the least privileges necessary for their duties.

Attribute-based extensions are designed to evaluate permissions in context. Context can be data attributes, stages of an operation, risk indicators about the environment, user behavior, checks on hardware, location, and threat information [6]. Furthermore, role hierarchies accommodate agents with different degrees of freedom, from a highly limited assistive agent to a fully autonomous agent, but constrained by guardrails. Permission boundaries must account for transitive access issues where agents interact with multiple systems and cannot allow for permission escalation across systems. In addition, the AI system may learn to automate the creation of access control policies. It can also detect patterns where agents repeatedly request the same resources, recommending policy changes to streamline this process without compromising security [7]. Periodic access reviews and recertification processes ensure that agent roles are suitable to the changing business.

3.3 Policy Enforcement and Conditional Access

Policy-driven access frameworks provide mechanisms to implement Risk-Adaptive Access Control (RAdAC), where decisions on access are made based on threat assessment and business conditions that change dynamically in real time [6]. Risk intelligence evaluates requests and login attempts for validity and risk and factors in things like the physical location from where requests are made, device fingerprints, and behavioral information [6]. Multiple risk assessments via application programming interfaces inform access decisions based on prior access tendency information, threat patterns, user session observation, and behavioral anomaly detection [6].

This means that when a trusted agent accesses the hub from managed devices in controlled geographies, exceptions to MFA requirements can be made with respect to the corporate access policy [6]. In contrast, when accessing the hub from external geographies or unregistered devices, corporate access policy can be enforced. All attempts to access sensitive resources during an active phishing attempt are denied regardless of

successful authentication with standard means [6]. Reinforcement learning algorithms can employ context or similar forms of decision-making based on the process of incrementally modifying thresholds based on feedback [6].

Behavioral adaptation enables continual authentication, tracking patterns of behavior in terms of access, such as when the system is used or checking the entity's credentials during an active session [6]. Machine learning models can identify anomalies during access, such as agents being seen to operate from unknown IPs or at times when this is unexpected, causing MFA or other pauses in access [6]. Threat intelligence feeds can automatically update policies when newly identified vulnerabilities affecting the agentic systems are found. Logging of policy evaluation decisions creates an audit trail of access denials and conforms to data-protection regulations such as GDPR, HIPAA, and SOX [6] (once they apply to agentic systems in production).

3.4 Lifecycle Management and Identity Recertification

During the agentic identity lifecycle, agent access is monitored, and recertification workflows periodically ensure that agentic identity continues to serve a valid purpose and that the associated access rights remain appropriate from both business and technical perspectives. AI predictive analytics can also be used in the domain of continuous improvement. When an AI notices a pattern in which certain groups of agents are frequently trying to access certain resources that are beyond what they are assigned, it can inform the administrator of a possible privilege escalation and allow administrators to reassign access before attackers exploit it [7].

Automated discovery mechanisms identify orphaned or dormant agent identities and alert the administrator so that they can be deprovisioned to minimize attack surface. Over 80% of breaches involve stolen or compromised credentials [6]. Automated version control of agent identity allows for autonomous updates of the system, the permissions granted to reflect agent capabilities, and compliance audits of historical records. AI-driven distributed domain management enables the movement of the risk metadata between the security domains of neighboring organizations, allowing global threats to be identified while respecting local business constraints [6]. In deprovisioning, credentials are revoked, attached resources are removed, active user sessions are killed, and agent identities are deleted from downstream systems to avoid orphan access rights.

4. Privileged Access Management for Sensitive Operations

4.1 Dynamic Secret Management and Credential Rotation

Privileged Access Management is the mechanism to secure the highly sensitive accounts that agentic systems may need in order to have elevated privileges to operate on enterprise infrastructure. Privileged accounts (system administrators, third-party vendors, or automated services) intrinsically have elevated privileges like system-wide change, sensitive data access, and control of mission-critical processes, which make them an attractive target of compromise [8]. These accounts can be compromised by insider attacks or by credentials being stolen through hacking, leading to mass data exposure, system outages, and regulatory consequences; it is estimated that insiders are responsible for 22% of all security incidents [8].

Dynamic secret management eliminates the need for static credentials hard-coded in agent configurations. Secrets are provided on demand to agents that require them to perform privileged actions. The credential exposure can be decreased further by implementing credential rotation policies, which automatically generate new secrets after predefined time intervals or on specific events such as suspected compromises or completion of privileged sessions. Agents request credentials from broker services, which authenticate the agent and verify authorization policies, returning short-lived credentials for specified target systems and actions [8]. This avoids credential persistence in agent memory or configuration stores, thereby limiting the number of injected credentials that could be collected from successful agent compromise.

Secrets management systems maintain audit log data of credential issuance events, usage events, or credential rotation events, allowing for the detection of suspicious privileged access patterns that may reflect agent compromise or malicious activity. AI-driven anomaly detection is another promising technology for secrets management that leverages machine learning and deep learning techniques on large volumes of access and usage data [9]. Detection provides a measure of how accurately a system is able to recognize whether an activity is malicious or not. Accuracy rates exceed 92% [8]. To limit the number of false positives, their techniques include using an adaptive learning algorithm that considers agent role, time of day, and the context of the action [8].

4.2 Just-in-Time Privilege Elevation

JIT privilege elevation provides agentic systems with temporary elevated privileges (in specific modes, like administration), which are released after the requested task is completed. Agents provide JIT privilege elevation requests along with reason codes (such as pre-approved workflows and business processes), and PAMs use contextualized access policies to determine whether to grant JIT privilege elevation. [8] However, without context-aware access policies, security concerns are multiplied, and it is difficult to distinguish between legitimate and suspicious actions in real-time [8].

When escalation is requested, rules can be defined to trigger workflows for approving high-risk operations or requiring approval when actions would have a direct impact on the business or increase the risk profile considerably. Elevation sessions are configured to have a maximum lifetime, after which elevated privileges are automatically revoked, even if agents forget to drop them. Healthcare has seen a 40% reduction in Electronic Medical Records (EMR) access violations within the first 90 days of deploying context-aware enabled policies [8].

Session recording enables a replay of any actions taken during agent privilege elevation for auditing purposes, compliance, and security review. Workflow orchestration can enable privilege elevation rules to give additional context about the various stages an agent is going through, enabling privilege only to be granted at the time it is needed and supporting least privilege. Automated compliance checks provide over 95% coverage for core access control requirements, such as authentication, authorization, session recording, policy enforcement, and audit logging [8].

4.3 Session Monitoring and Anomaly Detection

Session visibility is achieved through the continuous monitoring of agentic systems' privileged sessions so that any suspicious activity and policy violation can be detected in near real time. Session telemetry comprises a detailed record of privileged activity of accessed resources, invoked commands, data exfiltration, and effects on systems. By analyzing user activity and access logs and examining system behavior, AI models can find anomalies that indicate a potential security breach or other malicious behavior [9]. Behavioral analytics establishes the baseline pattern for each agent identity and looks for deviations that could indicate compromised agents or systems or emerging new attack patterns.

They may check for anomalous access patterns, volumes, target systems, and behaviors by comparing the present behavior of agents to their expected behavior (what they would be doing) on the current system and what they have done in the past. They can detect unauthorized access attempts and lateral movements not otherwise apparent [9]. The high detection rate (92%), precision rate (90%) and recall rate (93%) allows effective and timely insider threat or privileged account abuse detection, without overwhelming the information security team with too many false positives [8].

Alerts generated can notify security operations teams of high-risk anomalies, which can in turn terminate that session, revoke that user's credentials, or suspend the agent identity pending further investigation. Compliance monitoring engines can also notify security operations teams in real-time if there is an unauthorized access or policy violation [8]. Integration with security information and event management (SIEM) solutions connects privileged access events to other enterprise security data, exposing enterprise systems where several agentic systems operate to threats. In advanced DevSecOps environments, SAST, DAST, and SCA integrated with one another are known to catch 98% of vulnerabilities in the development pipeline [10].

4.4 Audit Trails and Compliance Reporting

Agentic systems must maintain a complete audit log of privileged agent actions for compliance and audit trail purposes. Audit logs must include each action description, business justification, approval chain, policy checks, and information about the environment at invocation [8]. Immutable, cryptographic audit logs (e.g., using blockchain) can also be employed for forensic analysis and compliance auditing purposes [8]. The role of AI in IAM helps organizations in regulatory compliance by automatically checking compliance, recording activities, and monitoring behaviors of agents/accessors within organizations [9]. This allows organizations to comply with industry regulations and maintain documentation, which can be useful for regulatory review later on, and helps identify non-compliance quickly [9].

Reporting capabilities consolidate information regarding privileged access across agents and systems. This includes reporting on privilege escalation or potentially noncompliant or inefficient workflows. Such reporting capabilities enable audit-ready reporting in real time, reinforcing the tool's value in achieving continuous and auditable compliance with industry standards. Automated frameworks achieve over 95% coverage across core access control baseline requirements [8]. Compliance frameworks increasingly expect autonomous systems to be treated as agents with

privileged access and call for audit trails showing appropriate controls.

As an example for advanced security architectures, the average cost of an AI-based scan as part of continuous assessment (12%-18% CPU overhead) and average latency of enforcement of Zero Trust policies for resource access below 40 milliseconds indicate that there are no bottlenecks in the flow of authorization [10]. Alerts, integrated remediation workflows, and AI-driven remediation tools reduced MTTR of critical vulnerabilities from weeks or months to 1-3 days [10]. Retention policies balance compliance and business requirements with storage costs. Older audit data is archived to save storage and costs, while recent activity remains available for rapid responses typical in security monitoring and incident response.

5. Integrating IAM and PAM for Comprehensive Governance

5.1 Complementary Control Layers

Identity and Access Management (IAM) combined with Privileged Access Management forms a layered defense of the governance of an agentic system. IAM is a framework of business processes, policies, and technologies for managing electronic or digital identities. IAM is about ensuring that the right entities have the appropriate access to the right resources at the right times for the right reasons [11]. Identity management systems are based upon four components: Administration, the creation and management of user accounts with up-to-date permission levels; Authentication, the verification of identity by means of passwords, biometrics, certificates, and other such means; Authorization, the granting of access to resources after authentication has taken place; Auditing, monitoring the integrity of the whole IAM system [11]. They make it possible for access to a system to be secure, auditable, and policy-compliant [11].

PAM manages the use of privileged accounts by controlling the automation of credential management and session management. These software offerings use secure access control, auditing, communication, and logging of all privileged credentials. PAM can be used to manage and control local and domain administrative accounts and client and service accounts for network devices, operating systems, application-to-DB (A2DB) accounts, and databases. This separation supports appropriate security controls for the distinct governance layers and coherence in management of autonomous processes. For example, IAM policies can grant an agent permission to perform database maintenance operations, and PAM controls can ensure that an agent only retrieves credentials when performing an authorized operation via secure channels and in a controlled session with real-time anomaly detection.

IAM and PAM alone are insufficient. IAM does not assure accountability over who issues the credential or how the credential is used once obtained. PAM does not enforce policies around identity during privileged access. The IAM-PAM architecture model combines the IAM and PAM approaches. IAM controls the identity for every link in the chain from the user to the cloud. PAM controls elevated access privileges issued and monitored during a session [11].

5.2 Policy Alignment and Orchestration

IAM and PAM policies should align whenever possible, minimizing friction, operational overhead, and the potential for security misconfigurations. The security policy represents the security requirements of the autonomous systems and establishes the controls to protect the system resources. In addition, it defines the rights and permissions of the different principals or groups and is a security model that specifies the constraints and access rights that the system must enforce [11]. Policy orchestration functions combine authorization decisions from the two layers to ensure that the actions of agents comply with IAM role-based permissions and PAM-based privilege management before executing them.

Unified policy languages help administrators address governance needs in both identity management and privilege management together, thus reducing complexity and preventing misconfiguration. Governance and organizational structure improve access control, especially since day by day, new threats target unauthorized access and steal valuable secret data with devastating business and economic consequences [12]. Policy synchronization mechanisms proactively propagate changes applied to IAM and PAM platforms to ensure that agent roles, business requirements, and security threats are addressed.

Integration architectures provide bidirectional communications between IAM systems and PAM systems. IAM systems send agent identity, role, and policy context to PAM systems. PAM systems send privileged access telemetry to IAM systems, informing risk and access review decisions. IAM formalizes the use and management of identities in all application areas, simultaneously securing distributed environments [12]. Even when federated identity management relies on identity providers

from a limited set of cloud infrastructure providers, and privacy, integrity, and non-repudiation are guaranteed by certified public keys from a public key infrastructure (PKI) [12], configuration of orchestration platforms may establish precedence rules to allow security-critical IAM to take precedence over PAM with operational requirements.

5.3 Unified Audit and Accountability

Integrating IAM and PAM audit streams thus allows for a full audit trail throughout an agentic system's operational lifecycle. The ACCESS control segment is concerned with authentication of users and applications as well as authorization based on roles ([11]). Unified audit repositories correlate events such as provisioning, role assignments, and policy updates with privileged access events such as credential usage and sensitive operations.

This gives the ability to trace agent authorization to specific privileged operations and investigations and can trace autonomous behavior to governing policies and business justifications. PAM implementations typically provide detailed enterprise-wide audit logging and notification of all privileged accounts. This gives any organization the ability to reduce privileged risks and meet compliance framing requirements by increasing authority and ownership over authorized credentials. [12] Unified dashboards summarize agentic system governance status, identity status, live sessions, privilege usage, and policy compliance, all in a single overview.

Under regulatory requirements, accountability frameworks typically use a single collection of audit records that store the agent's identity, the organizational unit that owns this agent, and the business process that authorizes the agent to perform the action. Compliance management is one of the core function zones of IAM [12]. Forensic capabilities can recreate the entire history of agent operations by aggregating identity lifecycle and privileged access events to address questions such as how an agent obtained a level of privilege, what they did with it, and whether it was policy-compliant.

5.4 Scalability and Operational Efficiency

Integrated IAM-PAM systems need to be able to scale to enterprise use cases that may have hundreds or thousands of agentic systems running on distributed infrastructure in parallel. Automation, to help reduce the administrative burden on PAM administrators, can be achieved by self-service provisioning workflows, policy-based access decisions, and automatic credential rotation. IAM can be used across a range of disjoint projects and apply to entire organizations, spanning private infrastructure and multiple cloud providers, both public and private [12]. IAM is often more enterprise-wide in its approach, as it identifies entities, validates cloud objects, and governs access to them using policy rules [12].

PAM implementations may also be expected to have advanced reliability, disaster recovery, and time-to-recover capabilities when applications fault or the underlying equipment loses the connection that can cause breakdowns [12]. PAM may be delivered as Software as a Service (SaaS) fully hosted in the Cloud, with specialist super nodes driving and merging policies and events and systems hosted by PAM vendors in multi-tenancy installations [12]. This also reflects the increased confidence organizations have in storing PAM cloud credentials, administration tools, and policies in the cloud [12].

Federation capabilities enable agentic identities and privilege controls to operate and be enforced across multiple security domains and cloud-hosted and partner-based systems, all while being managed from a centralized location. Performance improvements ensure that authorization and credential retrieval latencies are low enough to allow real-time autonomous decision-making. The architecture supports heterogeneous agentic models, including cloud-native agents, containerized workloads, serverless functions, and edge workloads, and scales to provide audit and monitoring for high-velocity telemetry streams from simultaneous agent sessions without sacrificing detection accuracy or introducing operational blind spots [11, 12].

Table 1: *AI-Enhanced IAM Capabilities for Agentic Systems [5, 6]*

IAM Function	Traditional Approach	AI-Enhanced Approach	Key Benefit
Authentication	Static credentials, periodic verification	Continuous behavioral monitoring, adaptive MFA	Real-time anomaly detection
Authorization	Manual role assignment, static policies	Automated role-mining, dynamic permission adjustment	Minimally privileged access
Access Control	Rule-based, uniform enforcement	Risk-Adaptive Access Control (RAdAC)	Context-aware security decisions

Threat Detection	Signature-based, reactive	Predictive analytics, behavioral baselines	Proactive risk mitigation
Policy Management	Manual updates, periodic reviews	Self-optimizing, feedback-driven adaptation	Continuous policy alignment
Compliance	Manual audits, periodic recertification	Automated tracking, real-time reporting	Breach prevention

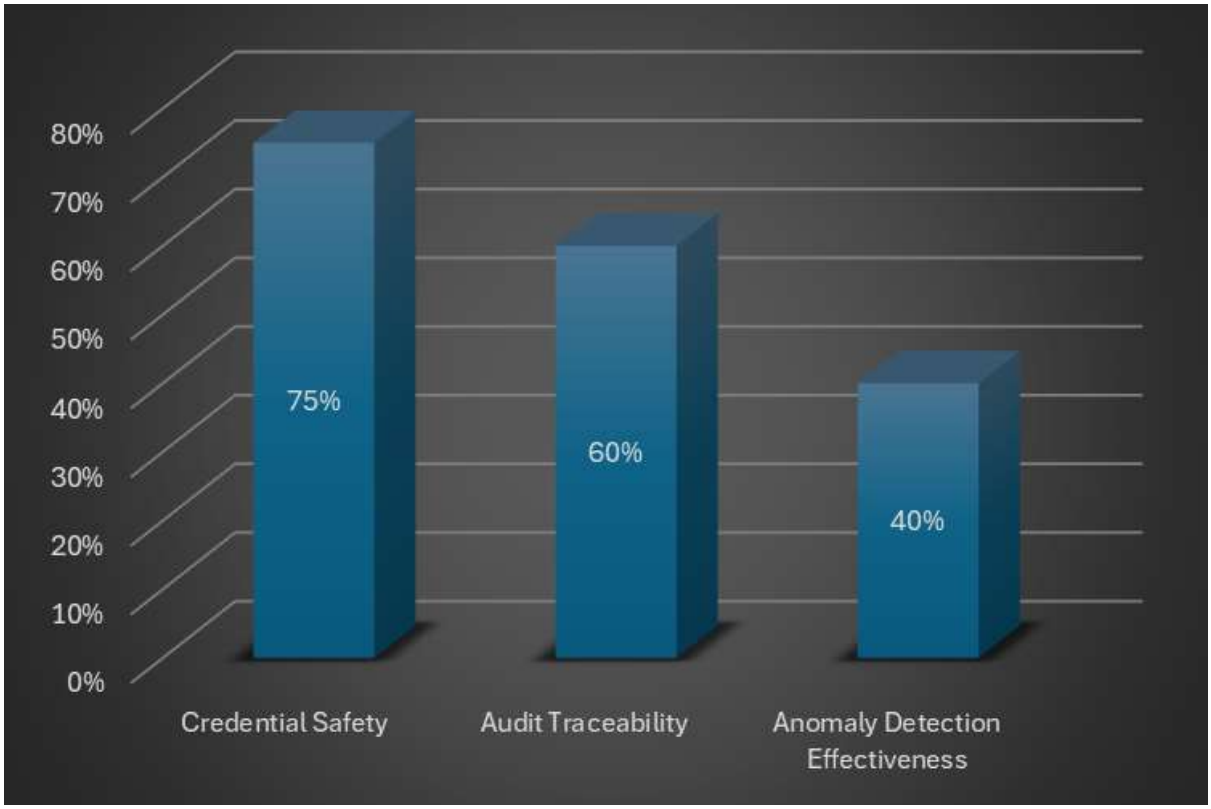

Figure 1: *Security Improvements Due to PAM Systems [3]*

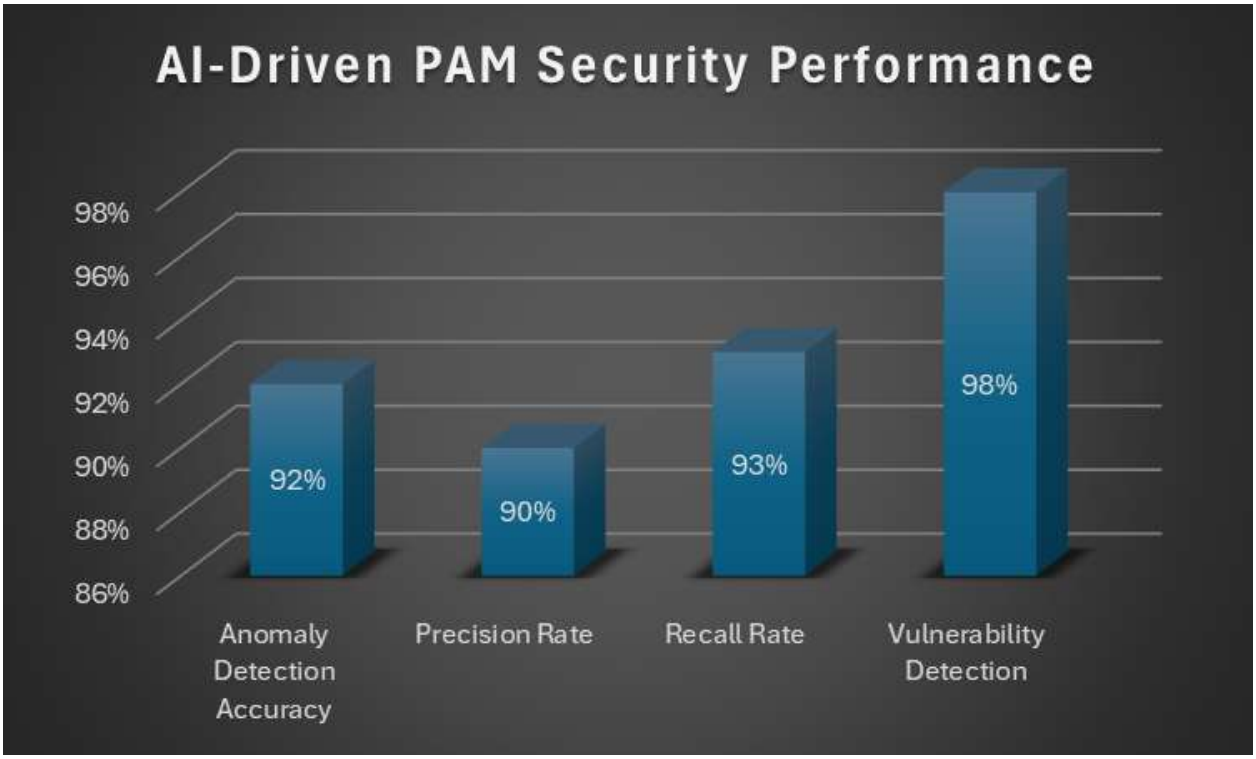

Figure 2: *AI-Driven PAM Security Performance: Detection Accuracy [8, 10]*

***Table 2**: IAM-PAM Integration Framework Components [8, 11, 12]*

Component	IAM Function	PAM Function	Integration Benefit
Administration	User account creation, role management, permission updates	Privileged account lifecycle, elevated access provisioning	Unified identity governance across privilege levels
Authentication	Identity verification via passwords, biometrics, certificates	Privileged session authentication, credential validation	Layered authentication for sensitive operations
Authorization	Resource access rights post-authentication	Privilege elevation, time-bound access grants	Context-aware access decisions spanning both layers
Auditing	Identity lifecycle monitoring, policy compliance	Privileged session recording, credential usage tracking	Comprehensive audit trails with 95%+ compliance coverage
Policy Enforcement	Role-based access control, attribute-based policies	Privilege controls, session monitoring	Coordinated policy orchestration
Compliance Management	Regulatory adherence, access reviews	Privileged risk reduction, credential ownership	Automated compliance checks across frameworks

6. Conclusions

IAM and PAM must be merged with agentic AI to manage agents that autonomously authorize and perform privileged operations across distributed IT infrastructure. IAM should be leveraged to create policies for agent identity management (including identity attributes like identity roles), policy-based access control decisions, and identity lifecycle management (including periodic recertification of the agent's identity). For example, user account management, dynamic secret provisioning, just-in-time elevation of privilege, monitoring, and logging and auditing privileged session activity. PAM implements a defense-in-depth security system because both frameworks fail to deliver complete governance. The combination of both frameworks establishes necessary supervision to maintain safe and compliant operations that can undergo auditing. Organizations can achieve operational advantages through agentic AI because it functions as autonomous identities while maintaining trust and accountability together with responsible AI control measures.

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] Soodeh Hosseini and Seilani Hossein, "The role of agentic AI in shaping a smart future: A systematic review," Array, 2025. Available: https://www.sciencedirect.com/science/article/pii/S2590005625000268

[2] Andras Cser, "The Forrester Wave™: Privileged Identity Management, Q4 2018," Forrester Research, 2018. Available: https://softprom.com/sites/default/files/The%20Forrester%20Wave%E2%84%A2_%20Privileged%20Identity%20Management%2C%20Q4%202018.pdf

[3] Badal Bhushan, "An Explainable Zero Trust Identity Framework for LLMs, AI Agents, and Agentic AI Systems," EuroLexis Research Index Library For Open Access Journals, 2025. Available: https://www.ijcaonline.org/archives/volume187/number46/bhushan-2025-ijca-925777.pdf

[4] Erhan Yilmaz, "Machine Identity Management in Modern Enterprise Security: Concepts, Challenges, and the Role of Privileged Access Management Systems," Engineering, Technology & Applied Science Research, 2026. Available: https://etasr.com/index.php/ETASR/article/download/16202/6268

[5] Sahaya Jestus Lazer et al., "A Survey of Agentic AI and Cybersecurity: Challenges, Opportunities, and Use-Case Prototypes," arXiv preprint arXiv:2601.05293, 2026. Available: https://arxiv.org/pdf/2601.05293

[6] Ramanan Hariharan, "AI-Driven Identity and Access Management in Enterprise Systems," International Journal of IoT, 2025. Available: https://inlibrary.uz/index.php/Ijiot/article/download/114074/115754

[7] Surendra Vitla, "The Future of Identity and Access Management: Leveraging AI for Enhanced Security and Efficiency," Journal of Computer Science and Technology Studies, 2024. Available: https://al-kindipublishers.org/index.php/jcsts/article/download/8619/7322

[8] Oluchukwu Modesta Oluoha et al., "Designing advanced digital solutions for privileged access management and continuous compliance monitoring," World Scientific News, 2025. Available: https://worldscientificnews.com/wp-content/uploads/2025/05/WSN-203-2025-256-301.pdf

[9] Prakash Somasundaram, "Unified Secret Management Across Cloud Platforms: A Strategy for Secure Credential Storage and Access," International Journal of Computer Engineering and Technology, 2024. Available: https://www.researchgate.net/profile/Prakash-Somasundaram/publication/379435761_Unified_Secret_Management_Across_Cloud_Platforms_a_Strategy_for_Secure_Credential_Storage_and_Access/links/6608f46b390c214cfd2b056d/Unified-Secret-Management-Across-Cloud-Platforms-a-Strategy-for-Secure-Credential-Storage-and-Access.pdf

[10] Ian Coston et al., "Enhancing secure software development with AZTRM-D: An AI-integrated approach combining DevSecOps, risk management, and zero trust," Applied Sciences, 2025. Available: https://www.mdpi.com/2076-3417/15/15/8163

[11] Samson Oruma et al., "Architectural views for social robots in public spaces: business, system, and security strategies," International Journal of Information Security, 2024. Available: https://link.springer.com/content/pdf/10.1007/s10207-024-00924-x.pdf

[12] Shadma Parveen et al., "Integration of identity governance and management framework within universities for privileged users," International Journal of Advanced Computer Science and Applications, 2021. Available: https://www.researchgate.net/profile/Sultan-Ahmad/publication/353079051