

Autonomous Identity Lifecycle Governance: An Architecture For Post-Termination Identity Sanitation

Varuni H. Sukhtankar

Independent Security Researcher, Mumbai, Maharashtra, India

Abstract—Autonomous AI agents and Non-Human Identities (NHIs) introduce lifecycle security challenges beyond conventional identity provisioning, authorization, and deprovisioning. During execution, an autonomous identity may create credentials, delegated authority, sessions, dependencies, and trust relationships that persist beyond its intended operational lifetime. This paper proposes the Autonomous Identity Lifecycle (AIL), a lifecycle-oriented architecture for governing autonomous identities from intent and identity synthesis through execution, termination, Identity Sanitation, and verified closure. AIL treat's identity termination and verified closure as distinct security outcomes. Identity Sanitation addresses the identification and appropriate disposition of residual security state, while Closure Evidence and Assurance provides evidence for validating that closure conditions have been satisfied. The architecture also introduces Identity-BOM as a lifecycle provenance structure for tracing identities, authorizations, delegations, dependencies, trust relationships, credentials, sessions, and related security artifacts throughout the identity lifecycle. In addition, the established principle of Least Persistence is applied through Identity Time-to-Live (I-TTL) to constrain the temporal lifetime of identity-related authority. The proposed architecture provides a foundation for governing post-termination identity security state; empirical validation and prototype implementation remain future work.

Index Terms—Autonomous Identity Lifecycle, Identity-BOM, Identity Governance, Identity Sanitation, Non-Human Identity.

I. INTRODUCTION

A. Background and Motivation

Identity governance has traditionally focused on the creation, authentication, authorization, and management of digital identities. These controls determine who or what receives access, under what conditions, and with what authority. The emergence of

autonomous systems, artificial intelligence (AI) agents, and Non-Human Identities (NHIs) introduces a broader identity governance challenge. Unlike traditional identities, autonomous identities may be created, delegated, exercised, and terminated at machine speed while continuously interacting with services, credentials, tokens, secrets, and dependent systems. [1]

Existing identity governance models primarily address how trust is established, managed, and exercised throughout an identity's operational lifetime. Although existing lifecycle approaches include deactivation and deprovisioning, less attention is given to how the security and trust state created during autonomous operation should be resolved when the identity terminates. An autonomous identity may terminate while credentials, permissions, delegated authority, active sessions, or trust relationships remain. [2], [3]

Termination therefore does not necessarily establish that the security state associated with the identity has also been resolved. As autonomous identities become increasingly embedded within enterprise environments, this persistence creates a need for controls that extend beyond provisioning and deprovisioning to address post-termination security state and its verified disposition. [2], [3]

Effective governance therefore requires not only managing identity and authority during operation, but also establishing how the security state associated with an autonomous identity is sanitized and verified when its operational lifetime ends.

B. Research Gap

The rapid growth of Non-Human Identities (NHIs) has expanded identity security beyond traditional governance concerns, with existing efforts addressing machine identity proliferation, credential management, orphaned access, excessive privileges, and machine-to-machine trust [3], [4].

Existing identity lifecycle and non-human identity approaches provide mechanisms for identity provisioning, authorization, credential management, and deprovisioning [1], [2], [4]. However, autonomous identities introduce an additional lifecycle challenge.

During execution, an autonomous identity may establish trust relationships, exercise delegated authority, and create dependencies that persist beyond its operational lifetime. Consequently, termination of the identity does not necessarily establish the disposition of the security state created during execution [2], [3].

The resulting challenge is therefore not limited to controlling identity creation or access, but also to determining what security state remains after an autonomous identity terminates and whether that state has been appropriately resolved.

Existing approaches provide controls for identity lifecycle operations; however, the systematic treatment of post-termination security state, its disposition, reconciliation, and verification of closure requires further architectural consideration.

This gap motivates an architecture that treats post-termination sanitation and verified closure as distinct objectives rather than automatic consequences of identity termination. This paper proposes the Autonomous Identity Lifecycle (AIL), a lifecycle architecture for governing autonomous identities from creation through verified closure, with mechanisms for lifecycle provenance, post-termination state disposition, closure validation, and evidence generation.

C. Contributions

This paper addresses the challenge of governing and verifying the post-termination security state of autonomous identities and makes the following contributions:

- **Identity Sanitation:**

A post-termination security control that extends identity governance beyond deprovisioning by identifying and addressing residual credentials, delegated authority, sessions, dependencies, trust relationships, and other security artifacts, and supporting subsequent closure validation.

- **Autonomous Identity Lifecycle (AIL):**

A five-phase security architecture that provides the lifecycle context required to govern autonomous identities from intent and identity synthesis through execution, termination, Identity Sanitation, and verified closure.

- **Identity-BOM (Identity Bill of Materials):**

A lifecycle provenance structure for tracking relationships among identities, authorizations, delegations, dependencies, trust relationships, credentials, sessions, and other security artifacts throughout autonomous execution and termination. Unlike conventional inventory views, the Identity-BOM emphasizes traceability of security state through termination and closure. **Least Persistence through Identity Time-to-Live (I-TTL):** A lifecycle principle that extends the Principle of Least Privilege into the temporal dimension by governing not only the authority granted to an identity, but also the duration for which that authority may exist.

The paper further defines three lifecycle security constructs:

- **Identity Residue:** The footprint of an identity after termination, consisting of residual authority, trust relationships, credentials, sessions, and security artifacts that persist beyond its lifecycle.
- **Shadow Realm:** An emergent security layer formed by the accumulation of Identity Residue that persists beyond the lifecycle of the identities that created it.
- **Identity Trojan:** A legitimate autonomous identity that completes its intended purpose while leaving behind exploitable capabilities.

Together, these contributions establish an identity-centric approach to post-termination security, in which termination stops the autonomous identity, Identity Sanitation addresses the associated residual security state, Identity-BOM provides lifecycle provenance for that state, and verified closure establishes that the identity's security lifecycle has been appropriately concluded.

D. Paper Organization

The remainder of this paper is organized as follows. Section II reviews related work and background in Identity Governance and Administration (IGA), machine and Non-Human Identity (NHI) governance,

and the lifecycle governance gap addressed by this paper. Section III introduces the proposed security constructs of Identity Residue, Shadow Realm, and Identity Trojan, and presents real-world evidence illustrating the underlying post-termination security-state condition. Section IV presents the Autonomous Identity Lifecycle (AIL) architecture, covering architecture overview, intent and identity synthesis, control and escalation, execution, termination and Identity Sanitation, and closure evidence and assurance.

Section V introduces the Identity-BOM, defining its architecture and role in lifecycle provenance and closure validation. Section VI presents Least Persistence as the architectural principle governing the temporal lifetime of autonomous identities and their authority.

Section VII examines the enterprise implications of post-termination identity security, including governance, compliance, financial risk, and industry applications and operational contexts. Section VIII discusses the benefits, challenges and trade-offs, and limitations of the proposed approach. Section IX outlines future research directions, and Section X concludes the paper. The Acknowledgment section recognizes contributions and support relevant to this work, and the References section provides the sources cited throughout the paper.

II. BACKGROUND AND RELATED WORK

A. Identity Governance and Administration

Identity Governance and Administration (IGA) manages digital authority across the identity lifecycle through provisioning, authorization, access review, policy enforcement, and deprovisioning. Its traditional model centers on the identity and the access state associated with it: what the identity can access, under what conditions, and for how long. [2], [3]

As enterprise environments have become more automated, however, identities increasingly do more than hold access. They establish sessions, receive delegated authority, interact with dependent services, and participate in machine-to-machine authorization flows. The identity remains the primary object of governance, while the relationships, dependencies, and security state created through its activity may extend beyond that object. [1]

B. Machine and Non-Human Identity Governance

This shift has expanded identity governance into machine and Non-Human Identity (NHI) populations. Applications authenticate to services, workloads access resources, automation platforms execute tasks, and autonomous systems interact directly with enterprise infrastructure. Machine identity management and NHI governance therefore increasingly address identity proliferation, credential management, excessive privilege, orphaned access, visibility, and lifecycle oversight. [3], [4]

These controls address an important part of the problem: ensuring that machine identities are known, authorized, and governed while they operate. Autonomous identities introduce a further lifecycle consideration because their operation can create trust relationships, delegated authority, and dependencies across connected systems that may persist beyond the identity's operational lifetime. [1]

C. Lifecycle Governance Gap

The resulting question is not simply whether the autonomous identity can be terminated, but whether the security state created during its operation terminates with it. An identity may be revoked while a delegated authorization remains active, a dependency continues to function, or a previously established trust relationship continues to influence another system. [1], [2], [3]

This creates a lifecycle boundary that conventional identity termination does not necessarily resolve. Identity termination establishes the end of the originating identity; it does not, by itself, establish the closure of the authority and trust relationships created during execution. Determining what remains, how those states should be disposed of, and how closure can be demonstrated therefore requires explicit post-termination sanitation and closure validation. [1], [2]

III. LIFECYCLE CLOSURE CHALLENGES IN AUTONOMOUS IDENTITY GOVERNANCE

A. Identity Residue

In this work, Identity Residue is defined as the footprint of an identity after termination. It encompasses residual authority, permissions, dependencies, and trust artifacts that continue to exist or influence systems after the originating identity has ceased to operate.

Identity Residue is not inherently malicious. It represents a post-lifecycle condition in which security state created through identity execution persists beyond the intended lifetime of the originating identity. Such residue may include residual permissions, delegated authority, dependent interactions, authorization chains, sessions, credentials, or trust relationships.

When these conditions remain isolated, they may be individually identifiable and manageable. As autonomous identities increase in number and complexity, however, residual effects may extend across systems and accumulate beyond the boundary of any single identity [4]. This creates a security condition in which the originating identity is no longer active, while security state associated with its previous operation remains present elsewhere in the environment.

B. Shadow Realm

The Shadow Realm describes this accumulation of Identity Residue at the system level. It is an emergent security layer formed by the collective persistence of authority, trust relationships, dependencies, and security artifacts after the identities that created them have terminated.

Unlike Identity Residue, which describes the footprint of an individual identity, the Shadow Realm describes the condition that emerges when those residual states persist across systems, applications, automation platforms, and machine-to-machine interactions. As these relationships become distributed across an environment, they may become increasingly difficult to identify, validate, and govern through conventional identity management.

The resulting condition is a divergence between the trust an organization believes exists and the trust that continues to operate in practice. Residual permissions, delegated authority, dependent interactions, and inherited trust may remain active without a clear owner, purpose, or lifecycle boundary.

The security significance increases when residual capability is not merely present but remains available for unintended use, creating an exploitable post-termination security state.

C. Identity Trojan

An Identity Trojan is a legitimate autonomous identity that completes its intended purpose while leaving

behind exploitable capabilities. Unlike a traditional malicious implant or persistence mechanism, an Identity Trojan originates from authorized and expected identity activity.

The defining characteristic is not malicious intent during execution, but the persistence of residual authority, trust, dependencies, or security artifacts beyond the intended lifecycle of the identity. The originating identity may have been properly deprovisioned or revoked while capabilities established during its operation remain available within the environment.

Identity Residue describes what remains; the Shadow Realm describes how residual conditions accumulate; an Identity Trojan describes the point at which that residual capability becomes exploitable.

This progression shifts the problem from identity termination to post-termination security exposure. A capability that survives the identity that created it may provide an unauthorized access path, enable privilege misuse, or influence downstream systems after the originating identity has ceased to exist.

This progression is illustrated in Fig. 1.




Fig. 1. Conceptual progression from Identity Residue to the Shadow Realm and Identity Trojan

D. Real-World Evidence

The lifecycle condition addressed by this paper is evident in real-world security incidents, even without autonomous AI identities. Such incidents show that access, credentials, keys, and trust relationships can persist beyond an identity's intended lifecycle.

Klue (2026): Klue reported that a threat actor leveraged a previously compromised GitHub Personal Access Token (PAT) to introduce unauthorized code into its integration service and collect third-party integration credentials, including Salesforce OAuth access and refresh tokens. [5]

Coupang (2025 incident; 2026 investigation): An investigation by the Korean Ministry of Science and Information and Communications Technology (MSIT) reported that Coupang did not invalidate a signing key used by a former developer for its user authentication system after the developer left the company. The former employee subsequently used the signing key to forge electronic access badges and bypass the normal authentication process. [6]

CISA (2024): Cybersecurity and Infrastructure Security Agency (CISA) and the Multi-State Information Sharing and Analysis Center (MS-ISAC) reported that a threat actor compromised the account of a former employee whose account had not been immediately disabled after departure. The actor used the account to access an internal Virtual Private Network (VPN), navigate the organization's on-premises environment, and conduct Lightweight Directory Access Protocol queries against a domain controller. [7]

Although these incidents were not AI-driven, they demonstrate the lifecycle condition addressed by this paper: identity termination and authority termination do not always occur simultaneously. Autonomous identities increase the relevance of this condition because they can create trust relationships, delegated authority, and machine-to-machine dependencies during execution.

These cases support the underlying lifecycle problem, not the AIL architecture itself. This distinction motivates treating post-termination security state and its verified disposition as explicit security objectives rather than assumed outcomes of identity termination.

IV. AUTONOMOUS IDENTITY LIFECYCLE (AIL)

A. Architecture Overview

The Autonomous Identity Lifecycle (AIL) is a lifecycle-oriented security architecture for governing autonomous identities from intent through verified closure. Rather than treating an autonomous identity as a static object, AIL governs it through defined stages in which authority is established, exercised, terminated, and resolved. [1]

AIL maintains lifecycle provenance through an Identity-BOM, recording the identity, authorization, delegation, dependencies, and trust relationships created or exercised throughout its lifecycle.

The resulting architecture is shown in Fig. 2.




Fig. 2. Autonomous Identity Lifecycle (AIL) architecture and cross-cutting lifecycle provenance model.

Each phase establishes a distinct control boundary within the identity lifecycle. Intent and Identity Synthesis define the purpose, authority requirements, and lifecycle parameters of the identity. Control and Escalation govern how authority is assigned and adjusted. Execution constrains how the identity operates within its authorized scope while establishing the security state that must be addressed at termination.

Termination and Identity Sanitation address the security state created during execution and establish the conditions for closure. Closure Evidence and Assurance governs the generation, preservation, retention, and eventual disposal of evidence required to demonstrate how the identity was governed and how closure was achieved, with retention determined by applicable regional regulatory, legal, and organizational requirements.

B. Intent and Identity Synthesis

Every autonomous identity should begin with a defined purpose. The Intent and Identity Synthesis phase establishes the objective the autonomous actor is intended to achieve, the authority required, the resources within scope, and the conditions governing its lifecycle.

Traditional identity management often begins with identity creation followed by access assignment. [2], [3] AIL introduces a preceding governance step in which the intended function, scope, authority, risk, and lifecycle are defined before the identity is created. Identity creation therefore becomes an outcome of a defined objective rather than an independent administrative action.

Intent serves as the governing specification for identity synthesis. It provides the parameters for determining identity attributes, permissions, lifecycle constraints, and governance controls. In this phase, three components are responsible for identity synthesis. The Objective defines the mission, clearly establishing the purpose for which the identity will exist. The supporting capabilities, permissions, systems, services, and Identity Time-to-Live (I-TTL) required for the mission are determined by the Planner, where I-TTL defines the maximum period for which the identity is authorized to exist, operationalizing the Least Persistence principle introduced later in the architecture. The Identity Compiler synthesizes the Planner's requirements into a purpose-built

autonomous identity with defined scope, authority, and lifetime.

The resulting identity enters execution with defined security and lifecycle boundaries that establish the expected conditions against which its subsequent security state can be evaluated at termination.

C. Control and Escalation

The Control and Escalation phase governs how the identity's authority changes during execution while remaining bounded by policy and risk.

The Policy Engine defines the authorization boundary and adaptive rules that govern how the identity's authority may change. The Risk Engine evaluates context, anomalies, and privilege escalation, continuously reassessing whether the identity's authority remains appropriate. Human-on-the-Loop (HOTL) provides human oversight without requiring human involvement in every action, intervening when risk thresholds are exceeded or policy ambiguity requires escalation.

Unlike Human-in-the-Loop (HITL), HOTL preserves autonomous execution and calls for human intervention only when defined conditions require it. Authorization therefore adapts to changing conditions at the tempo of autonomous execution while remaining bounded by policy, with human involvement conditional rather than continuous.

D. Execution

Execution begins with the Synthetic Identity produced by the Identity Compiler during the Intent and Identity Synthesis phase. Once authority has been established and brought under control, the autonomous identity operates within its defined scope. The Execution phase governs how that authority is exercised across the resources, services, and dependencies required to achieve the intended objective.

Execution may involve authentication, service interactions, delegated actions, or access to dependent systems. These activities remain subject to the authorization boundaries established in the preceding phases, including the identity's defined purpose, scope, permissions, and Identity Time-to-Live (I-TTL).

As execution progresses, the identity may establish sessions, delegated relationships, dependencies, and trust relationships with other systems. These interactions extend the security state created during

execution beyond the individual actions performed by the identity, creating residual security state that must be identified and addressed when execution ends.

E. Termination and Identity Sanitation

Termination ends the identity's ability to execute once the objective is complete or its authorized lifetime expires. This marks the end of operational activity, but not necessarily the end of the security state created during execution [2], [3]. This distinction is formalized in (1).

Identity Termination $\neq$ Verified Closure (1)

Termination establishes that the identity has stopped operating, whereas verified closure establishes that the associated post-termination security state has been appropriately resolved.

Residual artifacts that remain after termination can accumulate into the Shadow Realm, extending security exposure beyond the identity's intended lifetime. A legitimate identity that completes its mission while leaving behind an exploitable capability represents an Identity Trojan.

Identity Sanitation addresses the appropriate disposition of the security state associated with the terminated identity, including credentials, delegated privileges, tokens, sessions, secrets, trust relationships, dependencies, and execution contexts created or used during execution. The sanitation process determines whether these artifacts should be revoked, removed, expired, retained, or otherwise accounted for in accordance with applicable lifecycle requirements.

F. Closure Evidence and Assurance

The Closure Evidence and Assurance phase governs the evidence generated throughout the identity lifecycle and preserves what is required to demonstrate how the identity was governed and how its post-termination security state was resolved.

The Sanitation Report records the outcome of Identity Sanitation and the evidence supporting its disposition, using evidence gathered across the identity lifecycle. The Identity-BOM provides the lifecycle record and supporting provenance that trace the identity's credentials, authorization, delegation, dependencies, and trust relationships during sanitation and evidence generation.

Retention preserves the resulting evidence according to applicable organizational, regulatory, legal, and

audit requirements. Disposal securely removes the evidence when its defined retention period ends. Closure evidence therefore remains subject to protection, retention, and disposal requirements beyond the operational lifetime of the identity it documents. [2], [3]

V. IDENTITY-BOM

A. Definition and Architecture

An autonomous identity does not operate in isolation. Throughout its lifecycle, it can interact with services, agents, credentials, and delegated authorities, creating a dynamic web of machine-to-machine identity and trust relationships.

An Identity Bill of Materials (Identity-BOM) provides a lifecycle-aware record of this security context. It links the originating identity to the security relationships and artifacts associated with its operation, creating a structured view of the security state formed through its lifecycle.

The identity is therefore part of a broader security picture. Effective post-termination sanitation requires visibility into the relationships and artifacts associated with the identity, including those that may persist after the identity terminates. While a Software Bill of Materials (SBOM) maps the software components that comprise a software artifact, an Identity-BOM maps the security context associated with an autonomous identity and its lifecycle.

B. Lifecycle Provenance

The proposed Identity-BOM is a lifecycle provenance structure for representing the security relationships associated with an autonomous identity. It records identity metadata together with authorization, delegation, dependency, trust, credential, session, and related security artifacts that may be created or modified during execution.

Existing work has explored machine-identity inventories that enumerate non-human identities, their permissions, lifecycle state, and ownership. The Identity-BOM proposed in this work extends this inventory-oriented perspective toward autonomous lifecycle termination by emphasizing the provenance and disposition of security relationships across the complete lifecycle.

Within the AIL architecture, the Identity-BOM is therefore not intended merely as a static inventory. It

provides a traceability mechanism that can be consulted during Identity Sanitation to identify security artifacts associated with the terminated identity and during closure validation to determine whether expected lifecycle relationships have been appropriately terminated, revoked, expired, or otherwise dispositioned.

The proposed contribution is consequently the integration of lifecycle provenance with autonomous identity termination, sanitation, and closure assurance, rather than the claim that an inventory of machine identities is itself a new concept.

C. Role in Closure Validation

Termination ends the identity's operational lifetime, but closure requires verification that the security state associated with that identity has been appropriately resolved. The Identity-BOM provides the context needed to perform this validation by linking the identity to the credentials, authorization, delegation, dependencies, and trust relationships associated with its lifecycle.

During Identity Sanitation, the Identity-BOM enables these relationships and artifacts to be traced to their originating identity and lifecycle context. This allows the sanitation process to identify residual security state, determine its appropriate disposition, and record the actions taken to resolve or retain it.

The resulting sanitation state can then be compared with the expected closure conditions established during the identity lifecycle. This allows AIL to distinguish identity termination from verified closure: termination confirms that the identity has stopped operating, while closure validation establishes that the associated security state has been appropriately resolved.

VI. ARCHITECTURAL PRINCIPLES

A. Least Persistence

Least Persistence is an established security principle in which temporary security relationships, authorizations, or resources should not remain active beyond the period for which they are required. ETSI describes least persistence as designing temporary relationships so that they do not persist after use and applying expiration to security relationships where appropriate. [8]

In this work, the principle of least persistence is applied specifically to autonomous identity lifecycle governance. Autonomous identities may establish authorization, delegation, session, trust, or credential relationships during execution. If these relationships remain valid after their operational purpose has ended, they may constitute residual identity-related security state.

To operationalize this principle within the proposed, AIL architecture, this work introduces Identity Time-to-Live (I-TTL). I-TTL represents the intended temporal validity of an identity-related security relationship and provides a lifecycle control that can be evaluated during termination and Identity Sanitation.

The contribution is therefore not the introduction of least persistence itself, but its application to autonomous identity lifecycle governance and its integration with termination, sanitation, Identity-BOM provenance, and closure validation.

VII. ENTERPRISE IMPLICATIONS

A. Governance, Compliance, and Financial Risk

Autonomous business processes create a separation between business completion and identity closure. A financial transaction may reach its intended outcome while the identity that executed it retains authority, credentials, sessions, or delegated relationships. The transaction is complete; the security state is not.

This creates a governance condition: enterprise accountability must extend beyond the business event to the termination state of the identity that produced it. AIL addresses this boundary by associating business activity with lifecycle provenance and requiring the post-termination security state to reach verified closure. The resulting evidence can establish not only that an action was authorized, but that the authority through which it was performed was subsequently resolved.

This distinction matters because residual authority can outlive the purpose that justified it and may be misused by threat actors, causing unauthorized transactions, financial loss, operational disruption, and reputational damage. AIL treats business completion and identity closure as separate governance states, requiring closure validation rather than inferring it from completion of the transaction.

B. Industry Applications and Operational Contexts

Autonomous identities can operate beyond conventional Information Technology environments, including customer services, healthcare workflows, financial systems, and operational technology. Although the systems and regulatory conditions differ, each environment introduces the same lifecycle requirement: autonomous authority must remain bound to its intended purpose and operational lifetime, with the security state created during execution addressed and closure verified after execution.

The following scenarios illustrate how this lifecycle requirement manifests across different enterprise environments:

Technology, Media & Telecommunications (TMT) – AI Customer Support Agent: A short-lived agent accesses customer relationship management (CRM) platforms and customer-data application programming interfaces (APIs) to resolve customer issues. AIL bounds authority to the required scope and duration and addresses residual credentials, sessions, and data-access relationships when the interaction ends.

Consumer Products & Healthcare (CHS) – AI Medical Claims Processing Agent: An agent accesses electronic health record (EHR) and insurance systems to process claims. AIL governs the identity throughout the workflow and provides lifecycle evidence for termination and closure, supporting regulatory accountability and audit requirements.

Advanced Manufacturing, Mobility & Infrastructure (AMI) – Smart Factory Maintenance Agent: An autonomous agent interacts with Digital Twins, Internet of Things (IoT), and operational technology/industrial control systems (OT/ICS) for equipment monitoring and maintenance. AIL bounds machine authority to the operational window and addresses orphaned access and trust relationships when the maintenance activity concludes.

Financial Services (FS) – AI Loan Underwriting Agent: An autonomous agent participates in loan underwriting by interacting with banking systems, financial-data services, and risk engines under delegated authority. Its execution may create sessions, credentials, and trust relationships across these systems. AIL constrains the identity's authority to the underwriting objective and authorized lifetime, while post-execution Identity Sanitation addresses the security state that persists after the underwriting

activity ends and Closure Evidence provides the basis for verifying its resolution.

Across these environments, the business or operational task may end before the identity state does. AIL makes this boundary explicit by treating completion of the activity and verified closure of the autonomous identity as separate lifecycle conditions.

VIII. DISCUSSION

A. Benefits

AIL extends identity governance toward post-termination security-state management by treating the security state created during execution as an explicit governance concern. Separating termination from Identity Sanitation and closure validation provides a structured approach for addressing credentials, delegated authority, sessions, dependencies, and trust relationships that may persist beyond an identity's operational lifetime.

AIL also introduces a temporal dimension through Least Persistence and Identity Time-to-Live (I-TTL), constraining not only the authority granted to autonomous identities but also the duration for which that authority remains valid.

Identity-BOM and Closure Evidence and Assurance provide lifecycle provenance and evidence for validating closure. Together, they support traceability from identity creation and authorization through execution and termination, enabling residual security state to be examined within the lifecycle rather than reconstructed after an incident.

B. Challenges and Trade-offs

AIL introduces additional lifecycle controls that increase governance and implementation complexity. Identity-BOM maintenance, sanitation, and closure evidence require integration across identity providers, authorization systems, credential stores, sessions, and dependent services.

In heterogeneous enterprise environments, complete visibility into identity-related security relationships and artifacts may not be available, limiting the completeness of post-termination sanitation and closure validation.

Least Persistence introduces a trade-off between security and operational continuity. Shorter identity lifetimes may reduce residual exposure but increase provisioning, renewal, and workflow-interruption

risks. Appropriate Identity Time-to-Live therefore depends on mission duration, system behavior, risk, and policy requirements.

Human-on-the-Loop introduces a further trade-off between autonomous scale and human oversight. Escalating uncertain or exceptional conditions can preserve autonomous execution, but effectiveness depends on reliable policy evaluation, risk assessment, and escalation criteria. Excessive escalation may recreate human bottlenecks, whereas insufficient escalation may permit inappropriate authority changes.

C. Limitations

AIL is an architectural model rather than a validated implementation or quantitative security-control evaluation. This paper does not establish empirically that Least Persistence reduces security incidents, nor does it provide measured performance for sanitation, closure validation, or HOTL escalation.

The effectiveness of Identity Sanitation also depends on the ability to discover and control the security artifacts and relationships created during execution. Systems that lack sufficient lifecycle visibility may limit the completeness of sanitation and closure validation.

Finally, AIL addresses lifecycle-oriented identity risk rather than the broader threat landscape. It is not intended to replace controls for phishing, malware, authentication attacks, insider threats, or network-based attacks. Its scope is narrower: ensuring that autonomous identity authority and its associated post-termination security state are governed through termination, sanitation, and verified closure.

IX. FUTURE RESEARCH DIRECTIONS

The proposed architecture establishes a lifecycle model for governing and verifying post-termination security state, but its effectiveness requires empirical validation.

Future work should first define measurable assurance criteria for Identity Sanitation and closure validation, including sanitation success rate, closure coverage, and time to verified closure. Experimental evaluation should then assess Identity-BOM across heterogeneous identity and authorization domains, particularly for tracing and reconciling credentials,

delegated authority, sessions, dependencies, and trust relationships.

Least Persistence also requires quantitative evaluation to determine whether shorter identity and authority lifetimes reduce residual security state without disrupting operational continuity. Candidate measures include identity persistence duration and residual artifact count. HOTL should similarly be evaluated under autonomous operating conditions to determine appropriate escalation thresholds for risk, uncertainty, privilege escalation, and policy ambiguity.

Finally, a controlled prototype should evaluate the complete AIL lifecycle from identity synthesis through execution, termination, sanitation, and evidence generation. Such experiments would provide empirical evidence for the architectural assumptions presented in this paper.

X. CONCLUSION

Autonomous identity governance requires a lifecycle boundary beyond execution. An identity may complete its objective while the authority and security state created during execution remain unresolved. When such residual states accumulate, they can contribute to a Shadow Realm of persistent authority, trust relationships, dependencies, and security artifacts.

The Autonomous Identity Lifecycle (AIL) addresses this condition by connecting identity synthesis, controlled execution, termination, Identity Sanitation, and closure evidence within a single lifecycle. Identity-BOM provides provenance for tracing associated security state, while Least Persistence constrains the duration of identity authority.

The central contribution is a lifecycle model in which termination stops the identity, sanitation addresses residual security state, and verified closure confirms its resolution. By treating closure as an explicit security outcome, AIL provides a foundation for governing identity residue and its accumulation as autonomous identity ecosystems scale.

REFERENCES

- [1] Identity Lifecycle Management Working Group, Federal Chief Information Security Officer Council ICAM Subcommittee, "Identity Lifecycle Management Playbook," ver. 1.4, Mar. 31, 2026.

- [Online]. Available:
<https://www.idmanagement.gov/playbooks/ilm/>
- [2] Federal Chief Information Security Officer Council, Identity, Credential, and Access Management Subcommittee, and General Services Administration, Office of Government-wide Policy, “Digital Worker Identity Playbook,” ver. 1.3, Apr. 2025. [Online]. Available:
<https://www.idmanagement.gov/playbooks/dw/>
- [3] Cloud Security Alliance and Oasis Security, “The State of Non-Human Identity and AI Security,” Jan. 26, 2026. [Online]. Available:
<https://cloudsecurityalliance.org/artifacts/state-of-nhi-and-ai-security-survey-report>
- [4] OWASP Foundation, “OWASP Non-Human Identities Top 10,” 2025. [Online]. Available:
<https://owasp.org/projects/non-human-identities-top-10>
- [5] J. Smith, “CrowdStrike investigation summary and security improvements,” Klue, Jul. 1, 2026. [Online]. Available:
<https://klue.com/blog/crowdstrike-investigation-summary-and-security-improvements>
- [6] Ministry of Science and ICT, Republic of Korea, “Investigation results on the data breach by a former Coupang employee,” Feb. 10, 2026. [Online]. Available:
<https://english.msit.go.kr/eng/bbs/view.do?bbsSeqNo=42&mId=4&mPid=2&nttSeqNo=1221&sCode=eng>
- [7] Cybersecurity and Infrastructure Security Agency and Multi-State Information Sharing and Analysis Center, “Threat actor leverages compromised account of former employee to access state government organization,” AA24-046A, Feb. 15, 2024. [Online]. Available:
<https://www.cisa.gov/sites/default/files/2024-02/aa24-046a-threat-actor-leverages-compromised%20account-of-former-employee.pdf>
- [8] ETSI, “Cyber Security (CYBER); Standards mapping and gap analysis against regulatory expectations,” ETSI TR 103 990 V1.1.1, Mar. 2024.