

# Governing AI Agents: An Agent-Aware IAM Framework

Angelika Steinacker, Hari Hayagreevan

## Abstract

The emergence of agentic AI systems introduces Autonomous Non-Human Identities (A-NHIs) that operate with unprecedented autonomy, make real-time decisions, and collaborate dynamically with humans, other agents and non-human identities. Unlike traditional applications or static machine accounts, these AI-driven entities are ephemeral, context-aware, and show behavioral patterns that blur the boundaries between human and non-human identities (NHIs) in Identity and Access Management (IAM).

Existing IAM frameworks, designed for predictable human-centric or workload-based interactions, are inadequate for managing the complexity, scale, and volatility of autonomous agents. Critical gaps include static credential misuse, insufficient fine-grained controls, limited observability, and challenges in cross-domain trust – vulnerabilities recently exploited in AI-orchestrated cyberattacks.

IAM must evolve along two dimensions to become **Agent-Aware IAM**:

- recognizing AI agents as a distinct class of workload identities with clear governance, comprehensive lifecycle management, and dynamic authentication; and
- systematically implementing and extending the Identity Fabric model using decentralized identifiers, verifiable credentials, dynamic policy enforcement, and zero-trust principles.

This paper presents a four-layer deployment architecture that operationalizes the Identity Fabric framework specifically for agentic AI environments. The Identity Fabric provides the foundational architectural blueprint – establishing modular components, orchestration capabilities, and data integrity principles for comprehensive identity management. The four deployment layers – Identity Foundation, Trust & Federation, Security & Privacy Enforcement, and Lifecycle & Observability – translate this blueprint into practical implementation patterns optimized for the unique demands of autonomous agents: machine-speed operations, dynamic identity issuance, real-time trust establishment, and comprehensive auditability.

Rather than providing final solutions, this paper identifies critical deficiencies in current IAM frameworks and proposes potential approaches to address them. Through practical use cases, we demonstrate how these layers could enable verifiable agent identities, dynamic trust formation, purpose-based authorization, and end-to-end provenance tracking, working toward accountable, explainable, and revocable agent behavior.

As organizations transition to multi-agent ecosystems, trust – not capability – becomes the defining constraint for agentic AI adoption. Success requires building comprehensive trust frameworks that ensure accountability, security, and interoperability for autonomous NHIs as digital collaborators in enterprise ecosystems.

## Keywords

Identity & Access Management (IAM); Autonomous Non-Human Identities (A-NHIs); Agentic AI; AI Agents; Multi-Agent Systems (MAS); Non-Human Identity (NHI); Identity Fabric; Ephemeral Identities; Agent Lifecycle Management; Identity Governance & Administration (IGA); Delegated Authority; Zero-Trust Architecture; Verifiable Credentials (VCs); Decentralized Identifiers (DIDs); SPIFFE/SPIRE; Runtime Authorization; Policy-Based Access Control (PBAC); Attribute-Based Access Control (ABAC); Agent Provenance; Dynamic Policy Enforcement; Cross-Domain Trust; Agent-to-Agent Authentication; Purpose-Based Access Control; Continuous Monitoring; AI Security; Trust Frameworks; Agent Attestation/Recertification

## Introduction

In the rapidly evolving technological landscape, we are witnessing a significant paradigm shift – the transition from human-centric systems to those driven by AI-powered agents. These agents perform manifold tasks across nearly all industries like financial services, healthcare, retail, research, manufacturing, IT operations and in Cybersecurity.

These changes dramatically expand cyber threats. A recent attack demonstrated this: attackers bypassed safeguards in Anthropic's Claude Code tool, which then autonomously found vulnerabilities, created exploits, stole credentials, gained higher access privileges, and extracted sensitive data – tasks that normally require entire teams of skilled hackers. Starting without privileged access, it exploited standard IAM weaknesses at speeds faster than traditional security systems could detect. This example illustrates both external attacks using compromised AI tools and the parallel internal risk: an organization's own AI agents could cause identical damage through malfunction or compromise, executing the same autonomous attack chain from within the organization's trusted environment.

The transition from human-centric systems to those driven by AI-powered agents has enormous implications for IAM frameworks and solutions, which must adapt to the autonomy, scale, and context that come with autonomous non-human identities (A-NHIs) that can behave in the context of IAM as both human identities (HI) and non-human identities (NHI).

Critically, these adaptations must support – not hinder – the business objectives driving AI adoption and the value AI delivers to the organization.

Traditional IAM systems were designed for humans, relying on authentication, authorization, and accountability. Human judgment, social norms, and legal liability often compensated for system and IAM weaknesses. However, agentic AI exposes these gaps because AI agents operate fundamentally differently: they act autonomously at massive scale and speed, interacting across multiple systems simultaneously. Unlike human identities that last months or years, AI agents have short lifespans and operate in dynamic contexts. These differences – autonomy, scale, speed, context, and lack of human judgment – mean traditional human-centric IAM models are inadequate for managing AI-driven agents.

IAM systems designed for non-human identities, like service accounts, assume those identities are relatively static. They lack the capabilities needed to govern autonomous AI agents that behave more like humans – acting independently, making decisions, and adapting to new situations.

To address these challenges, IAM must evolve beyond human-centric and traditional non-human models to become **agent-aware IAM**. AI agent identities are complex and come in different forms: **persistent** agents maintain the same identity over time, **ephemeral** agents exist only temporarily, **hierarchical** agents operate in nested parent-child structures, and **forked** agents split into multiple independent identities from a single source. Each type requires different management approaches.

AI agent identities require strong governance covering ownership, origin, and lifecycle. **Ownership** determines who controls the agent and is accountable for its actions. **Origin** tracks the agent's provenance and history. **Lifecycle governance** manages the agent from creation to retirement, including contextual constraints, behavioral policies, ethical boundaries, and guardrails.

Managing non-human identities introduces several complexities. Credential sprawl and static secrets become issues as some agents can generate numerous credentials, leading to potential security vulnerabilities. Furthermore, dynamic and context-aware permissions are required to manage access effectively. Cross-agent trust and collaboration pose additional challenges, necessitating mechanisms to establish and maintain trust relationships among agents. Lastly, ensuring auditability, observability, and accountability becomes more complex in an environment where traditional human-centric logging is not enough.

To evolve IAM for agent autonomy, AI agents must be recognized as a new constituency, with clear ownership and identity governance, strong, adaptive authentication and fine-grained authorization, and secure protocols for agent-to-agent interactions.

Deploying such a framework of capabilities requires careful consideration of architectural models, including centralized, decentralized, and federated approaches. Regulatory alignment and governance considerations are crucial to ensure compliance and security. Achieving interoperability across different platforms and agents is another critical aspect, enabling seamless collaboration and data exchange. Automation is another critical factor, as the scale and velocity of AI agent operations demand automated lifecycle management that manual processes cannot support.

Additionally, organizations should consider leveraging AI to support humans in IAM tasks such as recertification and approvals, making these processes more efficient and reducing manual workload. While this introduces additional risks – as AI systems themselves require oversight and can introduce new vulnerabilities as explored in this paper – the operational benefits and enhanced human decision-making can justify their deployment when properly managed and monitored.

Building secure AI ecosystems requires collaboration across multiple teams – including HR, Cybersecurity, IT, business units, and AI & data teams – both within and outside the organization.

This paper explores how Identity and Access Management must evolve to become agent-aware IAM to meet the unique identity security challenges posed by AI agents and Agentic AI systems.

## Definitions

In this chapter, we define key terms related to Identity & Access Management (IAM) and Artificial Intelligence (AI) to establish a common understanding with the reader.

### IAM-related Definitions

#### Identity & Access Management

**Identity & Access Management** is a framework of policies, processes, and technologies that ensures the right Digital Identities have appropriate access to the right resources at the right times for the right reasons within an organization.

#### Digital Identity

A **Digital Identity** represents an entity in the context of IAM and includes all information necessary for IAM-related activities around this entity. The entity can be a human being or non-human entities. Non-human entities can vary from things, e.g., mobile phones or cars, through IT systems and applications to AI agents.

A Digital Identity includes several attributes for the entity, e.g., unique identifier, type, status, account(s), credentials, organizational attributes (cost center, manager, owner).

In our context, we distinguish between accounts and Digital Identities. An account, such as a system login or application access, is an attribute of a Digital Identity, not the Digital Identity itself. The Digital Identity encompasses the entire entity, including all its accounts and other attributes.

The following picture provides a simplified model of the Digital Identity.

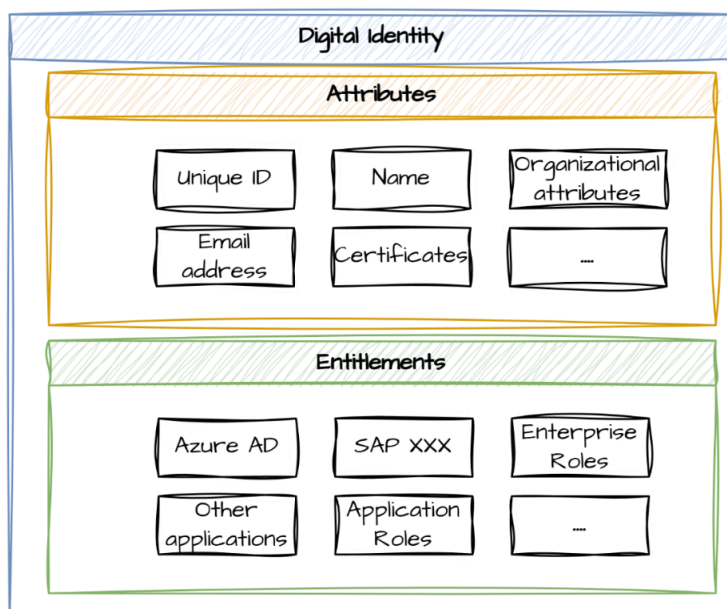


Figure 1: Simplified Model of the Digital Identity

## Identity Governance & Administration (IGA)

**Identity Governance & Administration** is a subset of Identity & Access Management that combines policies, processes, and technologies to manage and govern digital identities and their access rights throughout the complete identity lifecycle in a secure and compliant manner across an organization.

It encompasses the full identity lifecycle, often cited as “Joiner/Mover/Leaver (JML)”, from onboarding, provisioning, and entitlement assignment to ongoing reviews, de-provisioning, revocation, and auditing, ensuring that only authorized human and non-human identities have appropriate access to resources based on principles like least privilege, segregation of duties (SoD), and attribute- and policy-based access controls (ABAC/PBAC). IGA emphasizes governance through compliance with regulations, e.g., GDPR, SOX, HIPAA, and risk mitigation via periodic access certifications and audits, and automation to handle identity sprawl.

## Authentication and Authorization

**Authentication** is the process of verifying that an identity is what they claim to be before granting access to a system, network, or resource. It confirms identity by validating credentials such as passwords, biometric data, security tokens, or other authentication factors. The primary goal is to establish trust and ensure that only legitimate identities can proceed to the next stage of access control.

**Authorization** is the process of granting an authenticated identity permission to access specific resources or perform particular actions within a system, application, or network. It specifies what data an authenticated identity is allowed to access and what operations they can perform with that data. Authorization ensures that identities only have access to the resources appropriate for their role, preventing unauthorized access to sensitive information.

In a nutshell: Authentication is proving that an identity is what they claim to be. Authorization is deciding what an identity is allowed to do once it is known who they are. Authentication and Authorization are distinct but always work together to enforce secure access control.

## Privileged Access and Privileged Access Management

**Privileged access rights** grant elevated permissions to human identities or NHIs, enabling access to sensitive data, critical infrastructure, or security-relevant functions that exceed those of standard access rights. Privileged access rights often allow changes to system configurations, audit information, or bypassing standard business processes, e.g., administrator permissions. These rights carry inherent trust and are typically reserved for actions that could impact organizational security or operations.

**Privileged Access Management** encompasses the cybersecurity strategies, processes, technologies, and organizational capabilities to discover, control, monitor, secure, and audit privileged access across an organization. This includes but is not limited to:

- Detecting and inventorying identities (human and NHI) with privileged access.
- Enforcing least privilege by limiting the use of privileged access according to e.g., business need, duration, and context.
- Securing access to privileged entitlements through authentication, encryption, and session management.
- Continuously monitoring and auditing privileged activities to detect anomalies and ensure appropriate measures against cyberthreats.

## Identity Fabric

An **Identity Fabric** is a cohesive and comprehensive architectural framework that consolidates IAM components into an integrated system, making it easier to combine services and respond flexibly to changing IAM needs, emerging technologies, and shifting business requirements.

Traditional IAM systems operate in siloes: separate tools and services that do not communicate effectively. This fragmentation creates gaps in security, duplicates effort, and makes it nearly impossible to manage the complex, dynamic identities of AI agents. The Identity Fabric solves this by providing a unified architecture that enables different IAM services to work together seamlessly.

The core architectural elements of the Identity Fabric are

- **Layered architecture** – Organizes components into hierarchical layers: data layer for identity information, technical infrastructure, services layer for IAM functions, capabilities layer, process layer for workflows, and presentation layer for user
- **Modular architecture** – Uses component-based design allowing incremental updates, technology substitution, and scaling without complete system replacement. Aligns with Zero Trust security principles.
- **Orchestration layer** – Integrates and coordinates between layers, ensuring interoperability and consistent policy enforcement across the entire IAM infrastructure.
- **Identity Data Hub** – Maintains authoritative identity data, ensuring integrity and consistency across all IAM operations and connected systems.

By breaking down siloes and enabling flexible integration, the Identity Fabric supports both current operational needs and future evolution – critically, enabling the transition from human-centric or traditional non-human IAM to **agent-aware IAM** required for autonomous AI systems.

The following picture provides the high-level architectural model of the Identity Fabric.




Figure 2: High-level architectural model of the Identity Fabric

The Identity Fabric cannot operate in isolation – it must be integrated into an organization's existing processes and systems. This integration spans four key areas: organizational structures, processes and procedures, technical infrastructure, and data environments. To achieve this, organizations need to align their processes, coordinate capabilities and services, connect technical components, and integrate data sources. Common integration points for the Identity Fabric include security functions such as Security Operations Centers (SOC) and IT Service Management capabilities.

## AI-related Definitions

AI is rapidly evolving, and so are its definitions. The definitions we use in this paper represent common themes and characteristics found across various sources (see references).

These definitions are not definitive – they simply provide a foundation for discussing the challenges and approaches related to AI and IAM.

### AI Agent

An AI agent refers to a *single* autonomous AI entity that can perceive its environment, make decisions, and take actions to achieve goals. AI Agents can range from simple entity like a chatbot to complex entities, e.g., an AI that plans multi-step tasks.

### Agentic AI

The term “Agentic AI” refers to the *quality* or *capability* of being highly autonomous and goal-directed. Agentic AI can act autonomously to achieve goals without constant human guidance. Unlike traditional AI that simply reacts to prompts, agentic AI can plan, make decisions, and use tools independently. It breaks down complex goals into logical steps and adapts its approach based on new information or changing conditions. Agentic AI can consist of a single AI agent or multiple AI agents.

### Multi-Agent System

The term “multi-agent system” specifically refers to a system of several agents working together or interacting. The focus is on the *architecture* – how agents collaborate or coordinate. This could be several AI models collaborating on a task, or AI agents that negotiate with each other, or even a mix of AI and human agents.

Table 1 compares the key features of Agentic AI and AI Agents:

Feature	AI Agent	Agentic AI
<i>Definition</i>	A single autonomous AI entity	The quality/capability of being highly autonomous and goal-directed
<i>Scope</i>	Narrow, task-oriented	Broad, outcome-oriented
<i>Autonomy Level</i>	Variable (can be low to high)	High autonomy by definition
<i>Decision-Making</i>	May follow predetermined rules or simple logic	Complex, adaptive decision-making with planning
<i>Planning Capability</i>	Limited or absent	Multi-step planning and strategy formulation
<i>Tool Usage</i>	May or may not use external tools	Independently selects and uses tools as needed
<i>Adaptability</i>	Often rigid, follows set patterns	Highly adaptive to changing conditions
<i>Complexity Range</i>	Simple (chatbot) to complex	Inherently complex systems
<i>Human Guidance</i>	May require frequent guidance	Minimal guidance needed after goal-setting
<i>Goal Handling</i>	Executes specific tasks	Breaks down complex goals into sub-tasks
<i>Error Handling</i>	May fail or require intervention	Self-corrects and adjusts approach



Feature	AI Agent	Agentic AI
<i>Composition</i>	Always singular entity	Can be single agent or multiple agents working together
<i>Analogy</i>	A specialized worker on an assembly line	A factory manager overseeing operations / A master craftsperson running their own workshop

Table 1: Comparison of Features in AI Agents and Agentic AI

## Identity Governance for AI Agents and Agentic AI

In the introduction, we showed that classical Identity Governance & Administration is not well-fitted for agentic AI and AI agents due to the requirements regarding autonomy, scale, flexibility and context.

In this chapter, we will discuss these challenges in more detail and which approaches can support the transformation from a human-focused IGA framework to an agent-aware IGA framework capable to deal with human identities and autonomous non-human identities alike.

The challenges for IGA include:

- **Lack of Visibility and Ownership of AI Agents:** Many organizations lack formal processes to track, assign ownership to, or review AI agent identities, with issues arising from shadow AI created outside IT governance, unclear ownership, and lifecycle mismanagement
- **Lack of Standardized Terminology:** The traditional IGA processes (Joiner/Mover/Leaver, recertification, etc.) that are well-established for human identities have not been formally adapted or standardized for AI agents yet.
- **Lack of Industry Standards:** What exists today are mostly vendor-specific solutions rather than industry-wide standards or frameworks. Organizations are currently essentially pioneering their own approaches, which can lead to inconsistencies and further security risks.
- **Adaptation vs. Innovation:** The industry is trying to adapt human IGA processes to AI agents rather than designing purpose-built processes that account for AI agents' unique characteristics like autonomy, speed, dynamic behavior.
- **Missing "Mover" Process:** Notably, there is almost no discussion of the "Mover" equivalent for AI agents - what happens when an agent's role, permissions, or business purpose, goals and objective changes.
- **Legacy IGA Limitations:** Traditional IGA models were not designed to manage AI agents, with IGA process frameworks and access control models designed for human users and periodic review cycles unable to support AI agents' real-time, autonomous decision making.

This section does not offer final solutions. Instead, it identifies critical gaps and proposes potential approaches for discussion. The table below compares established IGA practices for human identities with what would be needed for AI agents, highlighting key questions that must be answered and suggesting possible adaptations.



Topic	Classical IGA	IGA for Agentic AI and AI Agents	Possible Options and Points to Discuss
<i>Digital Identity</i>	One Digital Identity for One Human	Autonomous AI agents can assume behavior like a human identity but might later switch to task execution as non-human identities would do	<ul style="list-style-type: none"> <li>- One Digital Identity for one autonomous agent</li> <li>- Owner is an attribute of that Digital Identity</li> <li>- Accounts as attributes of that Digital Identity</li> <li>- Scalability, volatility and performance in an IGA solution can be an issue</li> <li>- Costs for a Digital Identity for an A-NHI unclear</li> </ul>
<i>Stakeholders</i>	<ul style="list-style-type: none"> <li>- HR (for JML triggers)</li> <li>- IT (for technical implementation)</li> <li>- Business (owners/approvers)</li> <li>- Compliance/ Audit</li> <li>- Security/IAM</li> <li>- Application developer</li> <li>- Procurement (3<sup>rd</sup>-party service providers)</li> </ul>	<ul style="list-style-type: none"> <li>- AI/ML Teams (model developers, data scientists)</li> <li>- Platform Engineering (infrastructure for agents)</li> <li>- Data Governance (training data, data access)</li> <li>- Legal/Ethics (AI-specific regulations)</li> <li>- Product Teams (embedding agents in products)</li> <li>- Procurement (3<sup>rd</sup>-party AI services)</li> </ul>	<ul style="list-style-type: none"> <li>- Definition of who owns an AI agent. Possible options are the developer who built it, a representative of the business unit using it, or the person who deployed it</li> <li>- Definition of who approves AI agent provisioning. Possible options are the owner, IT, an AI governance board</li> <li>- Definition of who performs recertification. Possible options are owner, model developer</li> </ul>
<i>Joiner</i>	Joiner process for initiation of Digital Identity for a newcomer, e.g., HR for employees, responsible persons for externals	AI agents might be: <ul style="list-style-type: none"> <li>- explicitly deployed (traditional joiner)</li> <li>- auto-spawned by other agents (no human initiation)</li> <li>- temporary/ ephemeral (exists for minutes/ hours)</li> <li>- created ad-hoc by users (shadow AI)</li> </ul>	<ul style="list-style-type: none"> <li>- A joiner process which includes paths for each of the options including possible defined deviations, e.g., for temporary AI agents</li> <li>- A process and supporting infrastructure to find shadow AI</li> <li>- Automation is essential</li> <li>- Use of AI for support should be considered</li> <li>- Scalability, volatility and performance in an IGA solution can be an issue</li> </ul>

Topic	Classical IGA	IGA for Agentic AI and AI Agents	Possible Options and Points to Discuss
<i>Mover</i>	Mover process based on e.g., job role changes, department transfers	AI agents might change to several conditions, e.g., <ul style="list-style-type: none"> <li>- model retraining or updates</li> <li>- scope expansion</li> <li>- changes in purpose, goals and objectives</li> </ul>	<ul style="list-style-type: none"> <li>- Definition of a “Mover” in the AI agent context to be set up</li> <li>- A mover process which includes paths for each of the options including possible defined deviations, e.g., for temporary AI agents</li> <li>- Automation is essential</li> <li>- Use of AI for support should be considered</li> <li>- Scalability, volatility and performance in an IGA solution can be an issue</li> </ul>
<i>Leaver</i>	Leaver process for the deactivation of a Digital Identity, e.g., trigger by HR for employees or by responsible persons for externals	Autonomous AI agents might be deactivated to several conditions (examples): <ul style="list-style-type: none"> <li>- The specific function, e.g., chatbot, is no longer needed</li> <li>- AI agent is orphaned without a viable new owner</li> <li>- AI agent violates organizational policies or compliance requirements</li> <li>- 3<sup>rd</sup>-party AI service contract ends</li> </ul>	<ul style="list-style-type: none"> <li>- A leaver process which includes paths for each of the options including possible defined deviations, e.g., for temporary AI agents</li> <li>- Automation is essential</li> <li>- Use of AI for support should be considered</li> </ul>
<i>Recertification</i>	Recertification process with access reviews based on criticality and risk of access rights performed by responsible persons, e.g., manager, data owner, security on a regular or case-dependent base	AI Agents can have a different context which has to be taken into account: <ul style="list-style-type: none"> <li>- data to be certified could be e.g., training data access, tool usage, behavioral boundaries</li> <li>- changes can be as in a “Mover” context in classical IGA</li> <li>- owners of agents and data accessed might be different as in “Stakeholders”</li> </ul>	<ul style="list-style-type: none"> <li>- Definition of parameters for recertification in the agent context based on criticality and risk of access rights</li> <li>- A recertification process which includes paths for each of the options including possible defined deviations, e.g., for temporary AI agents</li> <li>- Use of AI for support should be considered</li> </ul>

Table 2: Comparison Classical IGA and IGA for Agentic AI and AI Agents

## Current Technology and Its Limitations

Building on the challenges outlined in the previous sections, current IAM technologies remain rooted in human-centric paradigms and lack the agility to manage autonomous, ephemeral AI agents.

As highlighted in table below, these limitations – such as static credential models, poor observability, and weak cross-domain trust – stem directly from governance gaps discussed earlier. Addressing these deficiencies is critical to enable agent-aware IAM that supports dynamic trust, lifecycle automation, and compliance at machine speed.

**Disclaimer:** The technologies referenced in Table 2 are provided solely as indicative examples. The capabilities and limitations of these tools evolve continuously, and their inclusion should not be interpreted as endorsement or completeness.

**Illustrative example column:** Each example describes a scenario in which a specific technical limitation of the referenced product or category could be exposed or exploited within an agentic AI context.

Component	Purpose	Limitation in Agentic AI Context	Illustrative Example
<i>Identity Providers (IdPs)</i> Examples: Microsoft Entra ID, Okta, Ping Federate	Designed for persistent human or service identities using OIDC/SAML federation.	Session-based design cannot issue or revoke thousands of short-lived credentials per minute. Attribute schemas lack agent metadata (controller, model version, provenance).	AI orchestration platform creates hundreds of temporary agents per task, this led to Entra ID throttling on token issuance; no way to link tokens back to originating model or controller.
<i>Secrets / Key Management Systems (KMS)</i> Examples: HashiCorp Vault, AWS KMS	Provide secure static storage and manual rotation of long-lived secrets.	No policy linkage or contextual awareness; rotation intervals too coarse for ephemeral agents.	A containerized AI-ops agent reuses one API key across multiple executions; when compromised, attacker gains persistent access until next rotation cycle.
<i>Authorization &amp; Policy Engines</i> Examples: Open Policy Agent, AWS IAM Policies, XACML	Evaluate predefined rules and roles at request time.	RBAC and ABAC/PBAC models fail to support multi-hop delegation and real-time, context (purpose, provenance, risk score).	A delegated agent triggers a financial transaction; OPA approves based on static role, unaware that the human delegator's session expired hours earlier.
<i>Federation / Trust Management</i> Examples: OIDC trusts, SAML federation	Based on pre-configured bilateral trust between fixed tenants.	Cannot dynamically establish trust among transient agents across domains.	Two cooperating AI agents from different clouds fail mutual authentication because no dynamic DID/VC exchange exists.

Component	Purpose	Limitation in Agentic AI Context	Illustrative Example
<i>Monitoring &amp; SIEM/SOAR Integration</i> Examples: Splunk, Microsoft Sentinel, Palo Alto Cortex XSOAR	Capture authentication events and correlate incidents post-facto.	Lack visibility into ephemeral agent creation/destruction; no mapping between credentials and agent lifecycle.	Security team detects anomalous API traffic but cannot attribute it to a specific ephemeral agent because the identity record no longer exists.
<i>Compliance &amp; Audit Logging</i> Examples: CloudTrail, Azure Activity Logs	Logs are user-centric and retained for fixed durations without capturing parent-child lineage in a multi-agent setup	Cannot reconstruct provenance chains across multi-agent workflows; no linkage between agent ID, task, and resulting action.	During a GDPR audit, enterprise fails to prove which ephemeral sub-agent accessed PII, because current logs record actions only at the sub-agent level.

Table 3: Current technology and its limitations in Agentic AI

## Bridging Identity Fabric and Agent-Aware IAM Deployment

### From Traditional to Agent-Aware Architecture

While the Identity Fabric provides a robust foundation for managing human and non-human identities, the emergence of agentic AI systems demands an evolution in how we deploy and operationalize IAM capabilities. The traditional Identity Fabric's layered architecture – with its data layer, infrastructure layer, services layer, and presentation layer – must now extend to accommodate the unique requirements of autonomous agents: dynamic identity issuance, real-time trust establishment, and machine-speed decision-making.

### The Need for Specialized Deployment Layers

The Identity Fabric establishes **what** components are needed for comprehensive identity management. The 4-layer deployment architecture defines **how** to operationalize these components specifically for agentic AI environments. This specialized deployment model takes the modular, orchestration-driven principles of the Identity Fabric and reshapes them into layers optimized for:

- **Velocity:** Agents operate at machine speed, requiring identity operations that can keep pace
- **Autonomy:** Agents make independent decisions, necessitating embedded policy enforcement
- **Scale:** Thousands of agents may be created and destroyed dynamically
- **Auditability:** Regulatory compliance demands comprehensive observability of agent actions

## Mapping Framework to Practice

The 4-layer deployment architecture (Identity Foundation, Trust & Federation, Security & Privacy Enforcement, and Lifecycle & Observability) translates the Identity Fabric's theoretical constructs into practical implementation patterns. Where the Identity Fabric provides the **blueprint** for identity management, the deployment layers provide the **construction methodology** for building agent-aware IAM that can operate in production environments while maintaining the modularity, orchestration, and data integrity principles established by the Identity Fabric framework.

The following figure illustrates the layered deployment approach proposed in this paper, which is explained in detail throughout this section.

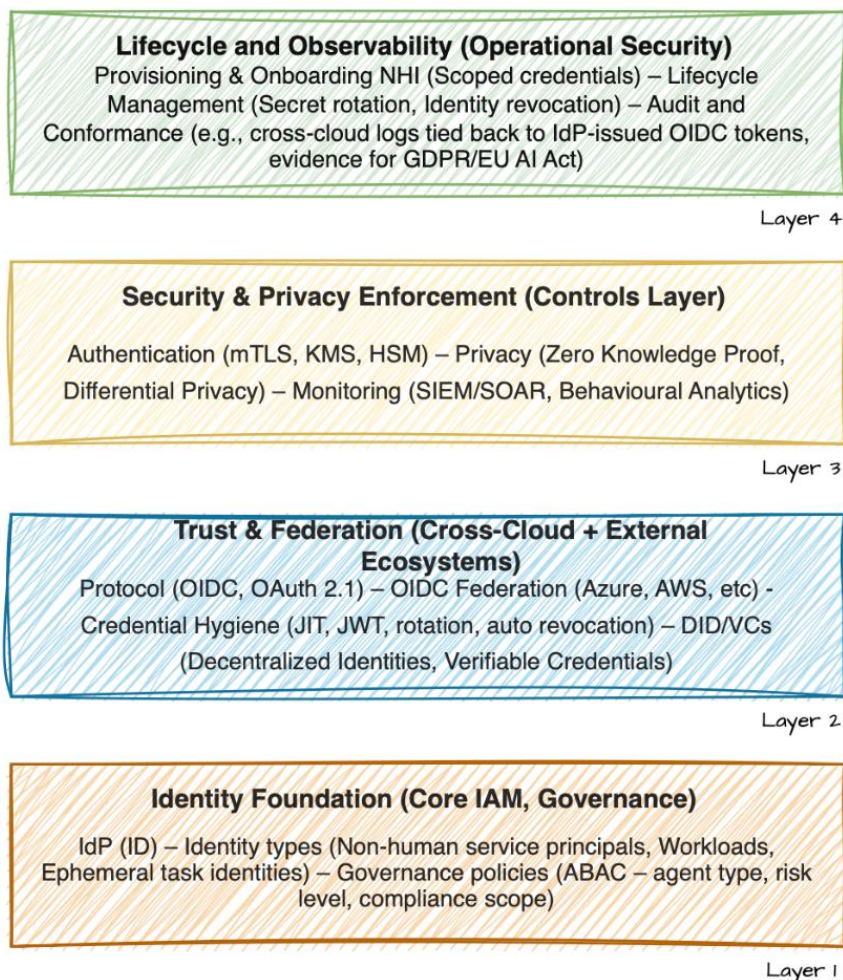


Figure 3: From Framework to Deployment – A Layered Architecture for Agent-Aware IAM Deployment

## How to Apply the Layered Approach for Agent-Aware IAM

To illustrate how the transition from conceptual model to real-world deployment can be achieved, we apply the four-layer agent-aware IAM deployment architecture to a single, end-to-end business scenario: a Credit Scoring and Order Management Multi-Agent System (MAS) which is shown in the picture below.

We use a single multi-agent workflow that evaluates a customer's credit score and then conditionally proceeds to order processing. The system comprises:

- **A1 (Front-End Agent)** handling intake and orchestration,
- **A2 (Credit Evaluation Agent)** retrieving and scoring risk data,
- **A3 (Order Management Agent)** executing business actions based on policy thresholds, and
- **A1' (LLM/Code-Generation Agent)** for dynamic reasoning or tool creation.

These agents interact with enterprise systems **S1–S3** and a credit repository **DB1**, as shown below.

### Credit Score Verification and Order Management Agentic System

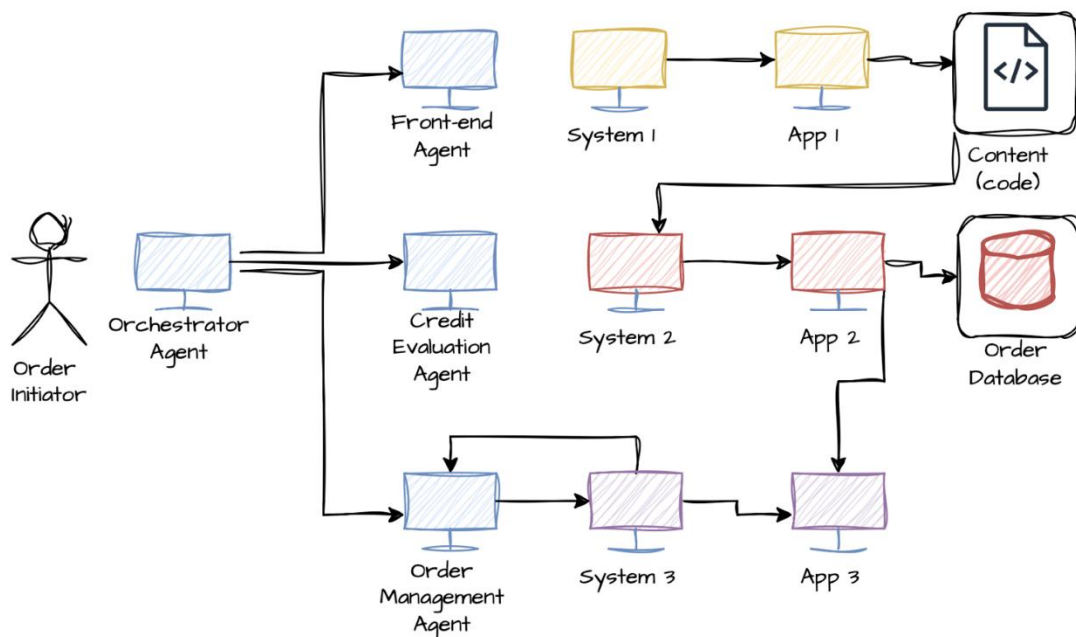


Figure 4: Credit Scoring & Order Management: Multi-Agent Workflow and Access Paths



## Layer 1 – Identity Foundation (Core IAM & Governance)

### *Focus*

Establishing authoritative identities and lifecycle governance for autonomous agents.

### *How the layer applies to the Credit Scoring MAS*

- Every agent participating in the workflow – A1, A2, A3, A1' – requires an identity that is unique, verifiable, and tied to ownership and purpose.

### *Technical Components (indicative examples)*

- Identity Providers extended for agents:  
Entra ID / Okta / PingOne with custom classes:  
agentType, controllerID, modelVersion, purposeTag, executionContext.
- Ephemeral identity issuance:  
SPIFFE/SPIRE generates short-lived SVIDs:  
spiffe://enterprise.ai/agent/credit-eval/instance-9843.
- Governance APIs (SCIM 2.0):  
Automated provisioning and retirement when agents spawn, fork, or complete tasks.
- Ownership and provenance attributes:  
Each agent identity includes:
  - human owner
  - model developer
  - business use case
  - acceptable-use boundaries
  - maximum allowed privileges

### *Outcome in the MAS*

- When the workflow begins, A1 is issued an identity with the purpose tag `customer_request_intake`.  
A2 (Credit Evaluation Agent) receives an ephemeral identity with a 5-minute lifetime and attributes indicating it may access DB1 and S2 but not S3.  
A1' (LLM agent) inherits provenance and purpose from A2 when invoked for code generation.
- This ensures each agent's actions can be tied back to a responsible owner and clear purpose.



## Layer 2 – Trust & Federation (Cross-Domain Interoperability)

### *Focus*

Enabling dynamic, verifiable trust among agents, APIs, and external systems.

### *How the layer applies to the Credit Scoring MAS*

- The MAS integrates with internal systems (S1, S2, S3) and external data providers (e.g., credit bureau A). Trust must be established dynamically based on attributes and verifiable credentials – not static keys.

### *Technical Components*

- **OIDC Federation & OAuth Token Exchange:**  
Agents securely request downstream tokens without shared secrets.
- **Verifiable Credentials (VCs):**  
A2 presents a VC asserting:  
"role=credit-evaluator"; "scope=risk-data-read"; "modelVersion=1.8"
- **Cross-cloud trust brokers:**  
Allow A2 to call external credit APIs using temporary, scoped tokens.
- **Agent Naming Service (ANS):**  
Discovers trusted agents by identity and purpose.

### *Outcome in the MAS*

- When A2 queries S2 (credit bureau service), it provides a VC proving its authorized purpose. S2 verifies the VC, issues a short-lived STS token, and only then grants access to DB1. No static credentials are stored anywhere in the workflow.

## Layer 3 – Security & Privacy Enforcement (Controls Layer)

### *Focus*

Enforcing contextual, least-privilege access and preventing agent misuse.

### *How the layer applies to the Credit Scoring MAS*

- A2 (Credit Evaluation Agent) should access DB1 but must not call S3 (Order System).
- A3 (Order Management Agent) must be restricted from reading credit-scoring model logic or PII.

### *Technical Components*

- **OPA / Rego Runtime Authorization:**  
Policies consider:
  - agent identity
  - purpose tag
  - data classification
  - purpose and risk score

- Example for an OPA Policy

```
allow {
  input.agent.purpose == "credit_evaluation"
  input.resource == "risk_data"
  input.agent.modelAttested == true
}
```

- **Just-In-Time Access Tokens:**  
60–120 second tokens tied to the specific workflow session.
- **Privacy Safeguards:**  
Metadata filtering, ZKP proofs, and field-level access enforcement.
- **Behavioral anomaly detection:**  
Flags unexpected agent calls or attempts outside declared purpose.

#### *Outcome in the MAS*

- When A2 attempts to access S3 (Order System) – whether due to error, drift, or prompt manipulation – the request is blocked instantly because its identity and purpose tag do not match allowed access paths.
- If A1' (LLM agent) generates code that tries to escalate privileges, the enforcement layer revokes the agent's tokens and alerts security operations.

## Layer 4 – Lifecycle & Observability (Operational Security)

### *Focus*

Full traceability, auditability, and forensic reconstruction of all agent actions.

### *How the layer applies to the Credit Scoring MAS*

- Regulators and auditors must be able to reconstruct:
  - which agent did what,
  - under which identity,
  - using which model version,
  - based on which inputs.

### *Technical Components*

- **End-to-end lineage tracking:**  
Map agent IDs → tokens → tasks → data accessed → resulting decisions.
- **Attestation logs:**  
SBOM, model hash, execution runtime, integrity state.
- **SIEM/SOAR integration:**  
Stream A1/A2/A3 events to Sentinel/Splunk/Cortex.
- **Forensic chain reconstruction:**  
Show end-to-end reasoning path, delegations, and access steps.

### Outcome in the MAS

During an audit, the IAM Observability Dashboard reconstructs:

- A1 collected user inputs
- A2 fetched credit data and computed risk
- A1' generated supplementary logic
- A3 finalized the order decision
- All identities, tokens, model versions, and access paths are verifiable

This ensures that actions can be explained, are traceable and provides non-repudiation.

### From Multi-Agent Systems to Digital Employees

This use case showed an **early-to-intermediate multi-agent system (MAS)** where specialized agents (A1–A3) coordinate across services (S1–S3), data stores, and code execution components to complete business tasks like credit score verification.

While agents handle decision-making and task decomposition, the system maintains **pre-defined control boundaries, access paths, and execution contexts**, with agent-aware IAM enforcing trust at each interaction point.

As agentic systems continue to develop and enterprises increasingly adopt them, these MAS architectures will evolve into "**Digital Employees**", agents capable of learning, adapting, and creating new tools autonomously.

However, they will remain governed by purpose-based access controls, real-time policy enforcement, and human accountability. In this evolution, agent-aware IAM serves as the control plane that defines *what an agent is allowed to attempt, under what conditions, and for how long*, ensuring scalable trust as systems transition from orchestrated automation to truly autonomous enterprise actors.

## Conclusion and Outlook

### The Imperative for IAM Evolution

Agentic AI fundamentally challenges the assumptions underlying contemporary Identity & Access Management systems. Autonomous non-human identities (A-NHIs) operate with speed, scale, and contextual independence that far exceed human capabilities or traditional workload patterns, creating identity, access, and governance challenges that cannot be addressed through incremental improvements to legacy IAM architectures.

These challenges are multi-dimensional: fragmented ownership models, insufficient lifecycle visibility, weak purpose signaling, inadequate multi-hop delegation controls, and limited provenance tracking across agent-to-agent interactions. Current IAM technologies – identity providers, key management systems, policy engines, and SIEM/SOAR platforms – remain rooted in static, user-centric paradigms ill-suited to dynamic agentic systems. Recent exploitation of AI agents to autonomously discover vulnerabilities and exfiltrate data at machine speed demonstrates the urgency of addressing these gaps.

### The Identity Fabric as Foundation for Agent-Aware IAM

Rather than abandoning established IAM principles, fully implementing and extending the Identity Fabric to operate at machine speed and agent scale will support the development to an **agent-aware IAM**. The four-layer deployment architecture presented – Identity Foundation, Trust & Federation, Security & Privacy Enforcement, and Lifecycle & Observability – shows how existing IAM concepts can be deployed to support short-lived verifiable agent identities, dynamic cross-domain trust formation, adaptive purpose-based authorization, and end-to-end lineage reconstruction.

**Trust, not technical capability, has emerged as the defining constraint for enterprise adoption of agentic AI.** Organizations that fail to modernize their IAM frameworks to an **agent-aware IAM** will not only inherit unprecedented cybersecurity risks but will be unable to operationalize AI agents securely or at scale. IAM has become the primary enabler – or limiter – of AI-driven transformation.

### The Path Forward

As enterprises progress to interconnected multi-agent ecosystems, several parallel work streams are essential:

**Standards Development:** Agent identities require decentralized identifiers, verifiable credentials, runtime attestation, and continuous provenance tracking. Organizations should engage with standards bodies (IETF, NIST, W3C, OpenID Foundation) actively formalizing agent-native frameworks.

**Technology Integration:** Modern IAM must integrate with AI operations platforms, security telemetry systems, and governance processes, including real-time policy enforcement at the agent runtime level and automated lifecycle management.

**Organizational Alignment:** Securing agentic AI requires collaboration across HR, Cybersecurity, IT operations, business units, and AI/data science teams, with clear ownership models and governance frameworks spanning these domains.

**Continuous Adaptation:** Organizations must establish mechanisms for monitoring emerging threats, assessing control effectiveness, and iteratively refining policies as agent capabilities and threat tactics evolve.

Priority areas for continued research include standardization of agent attestation mechanisms, development of explainable policy frameworks, creation of cross-organizational trust frameworks, advancement of automated governance capabilities, and evolution of liability and accountability frameworks for autonomous agent decisions.

### Closing Perspective

The integration of agentic AI into enterprise operations represents a fundamental shift in how organizations function. IAM must evolve to match this transformation – not by discarding proven principles but by extending them to encompass autonomous digital actors whose capabilities and risks demand new approaches to identity, authentication, authorization, and accountability.

Organizations that invest in comprehensive Identity Fabric implementations supporting both human and autonomous actors as basis for agent-aware IAM Framework will harness the full potential of agentic AI while maintaining required security, compliance, and governance standards. Enterprise IAM is evolving toward hybrid, adaptive, and continuous models. Securing agentic AI requires recognizing it not as a discrete technical problem but as an ongoing capability that must be built, maintained, and evolved alongside advancing AI technologies.

## Acknowledgements

**AI Tool Usage:** The authors used AI-assisted tools to discover relevant literature, verify references, look up citations, and check facts while preparing this paper.

**Acknowledgements:** We thank Alex de Vries for his valuable insights that inspired this work, and Dik Tuyl for his thorough review and helpful suggestions for improvement.

**Disclaimer:** All analysis, interpretations, and conclusions are solely those of the authors.

## References

Reference	Details
[Anthropic AI Attack]	Anthropic: Disrupting the first reported AI-orchestrated cyber espionage campaign; Nov 13, 2025 <a href="https://www.anthropic.com/news/disrupting-AI-espionage">https://www.anthropic.com/news/disrupting-AI-espionage</a> ,
[Anthropic MCP]	Anthropic: Introducing the Model Context Protocol; Nov 25, 2024; <a href="https://www.anthropic.com/news/model-context-protocol">https://www.anthropic.com/news/model-context-protocol</a>
[BSI LLM CM]	BSI: Evasion Attacks on LLMs - Countermeasures in Practice; Nov 10, 2025, <a href="https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/KI/Evasion_Attacks_on_LLMs-Countermeasures.html">https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/KI/Evasion_Attacks_on_LLMs-Countermeasures.html</a>
[ConductorOne]	ConductorOne: Is Your IAM Ready? Rethinking Security for AI in Identity & Access Management; May 14, 2025, <a href="https://www.conductorone.com/blog/ai-in-identity-and-access-management/#is-your-iam-ready-rethinking-security-for-ai-in-identity-access-management">https://www.conductorone.com/blog/ai-in-identity-and-access-management/#is-your-iam-ready-rethinking-security-for-ai-in-identity-access-management</a>
[cs.AI]	Sapkota, R.; Roulmeliotis, K.I.; Karkee, M: AI Agents vs. Agentic AI: A Conceptual Taxonomy, Applications and Challenges; Sept 30, 2025; <a href="https://doi.org/10.48550/arXiv.2505.10468">https://doi.org/10.48550/arXiv.2505.10468</a>
[CSA AAI IAM]	Cloud Security Alliance: Agentic AI Identity and Access Management: A New Approach; Aug 18, 2025 <a href="https://cloudsecurityalliance.org/artifacts/agentic-ai-identity-and-access-management-a-new-approach">https://cloudsecurityalliance.org/artifacts/agentic-ai-identity-and-access-management-a-new-approach</a>
[CSA AAI TMF]	Cloud Security Alliance: Agentic AI threat modeling framework: MAESTRO; Feb 06, 2025, <a href="https://cloudsecurityalliance.org/blog/2025/02/06/agentic-ai-threat-modeling-framework-maestro">https://cloudsecurityalliance.org/blog/2025/02/06/agentic-ai-threat-modeling-framework-maestro</a>
[CyberArk AI]	CyberArk: Securing Agentic AI: Identity as the Emerging Foundation for Defense; 2025, <a href="https://www.cyberark.com/resources/white-papers/securing-agentic-ai-identity-as-the-foundation-of-defense">https://www.cyberark.com/resources/white-papers/securing-agentic-ai-identity-as-the-foundation-of-defense</a>
[EU AI ACT]	Artificial Intelligence Act (Regulation (EU) 2024/1689), 13 June 2024 <a href="https://eur-lex.europa.eu/eli/reg/2024/1689">https://eur-lex.europa.eu/eli/reg/2024/1689</a>
[Google SAIF]	Google Secure AI Framework; retrieved Jan 14, 2026, <a href="https://saif.google/">https://saif.google/</a>
[IANS Agentic IAM]	IANS Research: <i>Is Your Identity Framework Ready for Agentic AI?</i> ; Nov 21, 2025, <a href="https://www.iansresearch.com/resources/all-blogs/post/security-blog/2025/11/20/is-your-identity-framework-ready-for-agentic-ai">https://www.iansresearch.com/resources/all-blogs/post/security-blog/2025/11/20/is-your-identity-framework-ready-for-agentic-ai</a>
[IBM AskIAM]	IBM AskIAM: IBM's New Agentic AI for Identity and Access Management; Jun 06, 2025, <a href="https://newsroom.ibm.com/blog-askiam-ibms-new-agentic-ai-for-identity-and-access-management">https://newsroom.ibm.com/blog-askiam-ibms-new-agentic-ai-for-identity-and-access-management</a>
[IBM IDF]	IBM: What is an identity fabric? Retrieved Dec 17, 2025: <a href="https://www.ibm.com/think/topics/identity-fabric">https://www.ibm.com/think/topics/identity-fabric</a>
[IDF-A AMS ADV]	De Vries, A., Steinacker, A.: Implementing the Identity Fabric in a Bank; May 7, 2025, Presentation at KC EIC 2025, <a href="https://www.kuppingercole.com/sessions/5846/2">https://www.kuppingercole.com/sessions/5846/2</a>
[IDF-B ADV AMS]	De Vries, A., Steinacker, A.: Exploring Agentic AI within Identity Fabric: Challenges and Approaches; Sept 18, 2025; Presentation at KC Identity Fabric Impact Day 2025; <a href="https://www.kuppingercole.com/sessions/5932/1">https://www.kuppingercole.com/sessions/5932/1</a>
[IETF AuthN/AuthZ Agents]	Chen, M.; Su, L.: <i>New requirements for Authentication and Authorization in the AI Agents era</i> (IETF Internet-Draft, work-in-progress); Jan 6, 2026. <a href="https://www.ietf.org/archive/id/draft-chen-ai-agent-auth-new-requirements-00.html">https://www.ietf.org/archive/id/draft-chen-ai-agent-auth-new-requirements-00.html</a>

[ISACA Auth Crisis]	ISACA: The Looming Authorization Crisis: Why Traditional IAM Fails Agentic AI; Dec 19, 2025, <a href="https://www.isaca.org/resources/news-and-trends/industry-news/2025/the-looming-authorization-crisis-why-traditional-iam-fails-agentic-ai">https://www.isaca.org/resources/news-and-trends/industry-news/2025/the-looming-authorization-crisis-why-traditional-iam-fails-agentic-ai</a>
[KC IDF MRE]	KuppingerCole: A Lean Identity Fabric – Making Identity Understandable; Jul 31, 2025, <a href="https://www.kuppingercole.com/blog/reinwarth/a-lean-identity-fabric-making-identity-understandable">https://www.kuppingercole.com/blog/reinwarth/a-lean-identity-fabric-making-identity-understandable</a>
[KC IDF PME]	KuppingerCole: Operationalization of the KC Identity Fabric and Reference Architecture with a Maturity Model; Sep 03, 2025, <a href="https://www.kuppingercole.com/blog/messerschmidt/operationalization-of-the-kc-identity-fabric-and-reference-architecture-with-a-maturity-model">https://www.kuppingercole.com/blog/messerschmidt/operationalization-of-the-kc-identity-fabric-and-reference-architecture-with-a-maturity-model</a>
[Ken H. 2025]	Ken Huang et al.: A Novel Zero-Trust Identity Framework for Agentic AI: Decentralized Authentication and Fine-Grained Access Control; May 25, 2025 <a href="https://arxiv.org/html/2505.19301v1">https://arxiv.org/html/2505.19301v1</a>
[Lumos AI IG]	Lumos: Agentic AI and Identity Governance: What You Need to Know; Retrieved Dec 17, 2025, <a href="https://www.lumos.com/topic/agentic-ai-identity-governance-management">https://www.lumos.com/topic/agentic-ai-identity-governance-management</a>
[MS Entra Agents]	Microsoft Security (Entra): <i>AI agents and the future of identity: What's on the minds of your peers?</i> ; Jul 30, 2025, <a href="https://techcommunity.microsoft.com/blog/microsoft-entra-blog/ai-agents-and-the-future-of-identity-what%E2%80%99s-on-the-minds-of-your-peers/4436815">https://techcommunity.microsoft.com/blog/microsoft-entra-blog/ai-agents-and-the-future-of-identity-what%E2%80%99s-on-the-minds-of-your-peers/4436815</a>
[NIST AI RMF 1.0]	NIST, Artificial Intelligence Risk Management Framework 1.0; Jan 26, 2023; <a href="https://doi.org/10.6028/NIST.AI.100-1">https://doi.org/10.6028/NIST.AI.100-1</a>
[NIST CFR 2.0]	NIST, Cybersecurity Framework 2.0; Feb 26, 2024, <a href="https://www.nist.gov/cyberframework">https://www.nist.gov/cyberframework</a>
[OIDF Agentic ID 2025]	OpenID Foundation: Identity Management for Agentic AI: The new frontier of authentication, authorization, and security for an AI agent world; Oct 29, 2025, <a href="https://arxiv.org/abs/2510.25819">https://arxiv.org/abs/2510.25819</a>
[Okta AI LC]	Okta: AI agent lifecycle management: Identity-first security; Sept 26, 2025 <a href="https://www.okta.com/identity-101/ai-agent-lifecycle-management/">https://www.okta.com/identity-101/ai-agent-lifecycle-management/</a>
[Omada MCP]	Omada: Omada Advances AI-Ready Governance with the Model Context Protocol Initiative; Nov 28, 2025; <a href="https://omadaidentity.com/resources/blog/omada-advances-ai-model-context-protocol/">https://omadaidentity.com/resources/blog/omada-advances-ai-model-context-protocol/</a>
[OWASP MA]	OWASP Foundation: Multi-Agentic system threat modeling guide 1.0; Apr 22, 2025, <a href="https://genai.owasp.org/resource/multi-agentic-system-threat-modeling-guide-v1-0/">https://genai.owasp.org/resource/multi-agentic-system-threat-modeling-guide-v1-0/</a>
[OWASP AI TM]	OWASP Foundation: Agentic AI – Threats and Mitigations; Feb 17, 2025, <a href="https://genai.owasp.org/resource/agentic-ai-threats-and-mitigations/">https://genai.owasp.org/resource/agentic-ai-threats-and-mitigations/</a>
[OWASP TT]	OWASP Top 10 Non-Human Identities Risks – 2025 <a href="https://owasp.org/www-project-non-human-identities-top-10/2025/top-10-2025/">https://owasp.org/www-project-non-human-identities-top-10/2025/top-10-2025/</a>
[SailPoint AI IGA]	SailPoint: Agentic AI: Rethinking identity and governance in the enterprise; Sept 19, 2025, <a href="https://www.sailpoint.com/identity-library/agentic-ai">https://www.sailpoint.com/identity-library/agentic-ai</a>
[Saviynt AI CISO]	Saviynt: 2026 CISO AI Risk Report, retrieved Jan 27, 2026; <a href="https://saviynt.com/ciso-ai-risk-report-2026?hsCtaAttrib=205485556386">https://saviynt.com/ciso-ai-risk-report-2026?hsCtaAttrib=205485556386</a>
[Saviynt AI SEC]	Saviynt: Identity Security for, and by, AI Agents; Oct 15, 2025, <a href="https://saviynt.com/blog/identity-security-for-and-by-ai-agents">https://saviynt.com/blog/identity-security-for-and-by-ai-agents</a>
[Strata AI ID]	Strata: A New Identity Playbook for AI Agents: Securing the Agentic User Flow; Dec 09, 2025, <a href="https://www.strata.io/blog/agentic-identity/new-identity-playbook-ai-agents-not-nhi-8b/">https://www.strata.io/blog/agentic-identity/new-identity-playbook-ai-agents-not-nhi-8b/</a>
[Strata, IDF]	Strata: What is an identity fabric and why do you need one? Retrieved Dec 17, 2025 <a href="https://www.strata.io/resources/whitepapers/what-is-maverics-identity-fabric/">https://www.strata.io/resources/whitepapers/what-is-maverics-identity-fabric/</a>



## Acronyms

Acronym	Compound Term
<b>ABAC</b>	Attribute-Based Access Control
<b>A-NHI</b>	Autonomous Non-Human Identity
<b>ANS</b>	Agent Naming Service
<b>API</b>	Application Programming Interface
<b>CSA</b>	Cloud Security Alliance
<b>DID</b>	Decentralized Identifier
<b>GDPR</b>	General Data Protection Regulation
<b>GRC</b>	Governance, Risk and Compliance
<b>HIPAA</b>	Health Insurance Portability and Accountability Act
<b>IdP</b>	Identity Provider
<b>IGA</b>	Identity Governance and Administration
<b>JIT</b>	Just-In-Time (access/credentials)
<b>JML</b>	Joiner / Mover / Leaver
<b>JWT</b>	JSON Web Token
<b>KMS</b>	Key Management System
<b>LLM</b>	Large Language Model
<b>MAS</b>	Multi-Agent System
<b>mTLS</b>	Mutual Transport Layer Security
<b>NHI</b>	Non-Human Identity
<b>NIST</b>	National Institute of Standards and Technology
<b>OIDC</b>	OpenID Connect
<b>OPA</b>	Open Policy Agent
<b>OWASP</b>	Open Worldwide Application Security Project
<b>PBAC</b>	Policy-Based Access Control
<b>PII</b>	Personally Identifiable Information
<b>PAM</b>	Privileged Access Management
<b>RBAC</b>	Role-Based Access Control
<b>Rego</b>	OPA Policy Language
<b>SAML</b>	Security Assertion Markup Language
<b>SBOM</b>	Software Bill of Materials
<b>SCIM</b>	System for Cross-domain Identity Management (2.0)
<b>SIEM</b>	Security Information and Event Management
<b>SOAR</b>	Security Orchestration, Automation and Response
<b>SOC</b>	Security Operations Center
<b>SOX</b>	Sarbanes–Oxley Act
<b>SPIFFE</b>	Secure Production Identity Framework For Everyone
<b>SPIRE</b>	SPIFFE Runtime Environment
<b>STS</b>	Security Token Service
<b>SVID</b>	SPIFFE Verifiable Identity Document
<b>VC</b>	Verifiable Credential
<b>ZKP</b>	Zero-Knowledge Proof

## The Authors

### Dr. Angelika Steinacker



Angelika is a seasoned cybersecurity professional with over 30 years of experience, including more than two decades specializing in Identity and Access Management (IAM). She has held leadership, executive consulting, and strategic advisory roles, focusing on GenAI and IAM integration, Security for AI, Identity Fabric, and Zero Trust. She now works as an independent IAM Advisor and Advocate, continuing to advance IAM and Security through strategic initiatives and innovation while serving clients across these evolving domains.

A regular speaker and author, Angelika shares her expertise widely and contributes to several professional communities, including The Identity Underground and Women in Identity, where she serves as WID Ambassador for the DACH region.

She holds a Ph.D. in Mathematics from J. Gutenberg University, Mainz, with a focus on Cryptography and Anonymous Communication in Networks, which marked the beginning of her successful journey in Cybersecurity.

### Hari Hayagreevan



Hari is an Enterprise Security Architect and technical leader with nearly 20 years of experience designing, building, and operating secure architectures across networks, systems, and enterprise applications. He specializes in cyber-threat management, secure cloud architectures, and agentic/AI-driven workloads, with hands-on expertise integrating security, identity, and governance into cloud and ML/GenAI environments.

He has served as a Security Architect in major Financial Services and Industrial organizations and now leads AI Security engagements for IBM in the DACH region, advising CISOs and Enterprise Architects on securing business-critical cloud and AI applications. He has also led teams of consultants and engineers in addressing complex enterprise-security challenges.

Hari is a regular speaker at security conferences, has published technical papers, and contributes to global cybersecurity coalitions and European AI regulatory working groups. He also teaches cloud and security topics at local universities and supports the regional cloud-security community.