

April 2026



ICIT

Identity Security In the Critical Path for Agent Deployment

Jim Routh

Fellow, Institute for Critical Infrastructure Technology

www.icitech.org

Table of Contents

Executive Summary	03
I. The Impending Collision: AI Ambition vs. Legacy Identity	04
II. The New Architectural Paradigm: From Identity Governance to Identity Security	06
III. The Strategic Shift: Dynamic Provisioning & Continuous Validation	07
IV. The Economic Engine: Self-Funding Transformation	09
V. Conclusion: A Call to Action for the CISO	10
About	11



About ICIT

The Institute for Critical Infrastructure Technology (ICIT) is a nonprofit, nonpartisan, 501(c)3 think tank with the mission of modernizing, securing, and making resilient critical infrastructure that provides for people's foundational needs. ICIT takes no institutional positions on policy matters. Rather than advocate, ICIT is dedicated to being a resource for the organizations and communities that share our mission. The views and opinions expressed in this essay are solely those of the author(s) and do not necessarily reflect the official policy or position of ICIT. Any assumptions made within the analysis are not reflective of the position of any entity other than the author(s). To learn more, please visit www.icitech.org



Thank you to our Strategic Partner **CyberRisk Alliance** | cyberriskalliance.com

Objective:

Create a paper that guides the financial services sector and other sectors to evolve practices toward identity security (and away from identity governance platforms) as an essential component for AI Agent deployment at enterprise scale. Identify alignment with the practical approach used by Wells Fargo.

Executive Summary

Enterprises, large and small, are under significant pressure to leverage AI capabilities to fundamentally improve business opportunities by both lowering operating costs and driving business growth. Critical path for agent deployment at scale includes a fundamental redesign of identity security capabilities.

Legacy identity governance platforms and processes were designed to manage human identity access by humans making decisions: provision, certify, and deprovision. The consequences of this record-keeping architecture include increased costs and wait time as the business grows. Enterprises have a backlog of application integration projects. Existing processes are unable to handle non-human identities, which outnumber human identities by as much as 80 to 1 today, and this is projected to be 400 to 1 in a few years due to agent deployment.

The architecture of identity security to meet today's requirements, including the implementation of AI agents, must be a data lake of entitlement usage attributes that enable every identity (human and non-human) to be registered, risk-scored, and with policy applied to block specific transactions while enabling others.

The result of the redesign of identity security is an increase in the volume of transactions at a lower cost, with higher satisfaction for stakeholders, supporting the concept of "least privilege" to improve cyber resilience. The business case for this transformation is enabled through dynamic provisioning (lowering costs while increasing capacity) and ultimately realized with a layer of continuous validation applied to privilege access management that operates in real time. Controls will be enforced using AI agents to govern the capabilities of AI agents operating within the enterprise, in addition to agents from third parties.

I. The Impending Collision: AI Ambition vs. Legacy Identity

There is exceptional pressure from business stakeholders to adopt and deploy AI agents and the necessary infrastructure to achieve significant improvements in enterprise productivity. Major firms like NVIDIA are aiming for a ratio where 50,000 employees work alongside 100 million AI assistants. Other companies focus on decreasing cycle times for tasks like strategic planning by 70-85%. Others are pursuing intelligent document processing to classify unstructured data using AI.

The rapid rate of AI adoption is unlike any other technology breakthrough in human history, impacting how people work, live, learn, and entertain. The achievement of the transformational objectives at enterprise scale requires a fundamental change to identity security supporting non-human digital identities. This is necessary both to discover and register the existence of AI agents, along with applying policy to govern what agents cannot do (guardrails), and what they can and should do (policy).

Identity governs the platforms and processes designed for humans to manage access to systems for other human identities. These governance platforms were essentially record-keeping systems to track transactions, grant access to specific systems, certify that access periodically, and revoke it when it was no longer needed. Record-keeping systems are relatively straightforward, providing compliance reporting for audit transactions for both internal and external auditors and regulators. Most enterprises have implemented specific systems to automate the management of entitlements to privileged users. Systems administrators need additional access to sensitive systems and entitlements to provide effective management of the IT assets. Power users who administer access to distributed data sources and database administrators all require privileged access to perform their roles. Privileged Access Management (PAM) applications were either purchased or built to satisfy these requirements that continue to grow as the business grows. Today, close to 90% of PAM systems within enterprises only support human identities and not non-human identities.

The migration to cloud computing and SaaS applications enables the use of APIs for machine-to-machine connectivity, which has increased the use of service accounts. The implementation of AI agents puts more pressure on enterprise management of non-human identities that can act on behalf of humans to make decisions, process data, and publish output. The growth of non-human identities, including AI agents, makes legacy identity governance platforms and workflows obsolete. Systems and workflows that provide data to humans to interpret, add context to, and then take specific actions are expensive and time-consuming. These systems are restricting business growth. Many of these legacy systems were built decades ago, when technology infrastructure was less complex, and most transactions were initiated with human judgment.

Experts expect that enterprises will have a ratio in the near term of several hundred non-human identities to a single individual identity. Enabling this change requires a new architecture to handle the growth in non-human identities. The requirements for this architecture include a data lake to capture entitlement attributes and how entitlements are used by all identities (human and non-human). The activity of entitlement usage is vital to establishing patterns. The patterns can be useful to determine a risk profile for each identity. The risk

profile can be applied to entitlement requests to grant entitlement access, lowering costs while satisfying stakeholders with immediate access.

Business pressure for positive outcomes from investments in AI is increasing across industries and enterprises. The deployment of AI agents is accelerating the use of non-human identities. Estimates today, following the deployment of cloud computing, are 80 non-human identities to every human identity. Estimates within two years following the adoption of AI agents will be several hundred non-human identities to human identities. Cloudflare CEO Matthew Prince predicted that AI bot traffic will exceed human traffic online by 2027. A recent study by Human Security confirms that this milestone has already been met.

Why do enterprises continue to rely on processes and platforms designed to manage human identities? AI deployment in a majority of enterprises is taking place in unsanctioned and ungoverned ways today. No registration processes exist for AI agents to discover, record, and ultimately manage AI agent identities. The maturation of standards, like the Model Context Protocol (MCP) is beneficial for enterprises deploying agents, but it cannot enable agent identity registration and policy management. Agent deployment without identity registration increases the volume of unknown assets deployed, thereby extending the attack surface for threat actors.

Conventional IAM governance platforms are obsolete in design and insufficient to support the need for identity security for AI agents. Newer architectures, additive to AI governance platforms, available today, remain limited in the use of activity data to reduce the need for human-initiated transactions in identity management. Enterprises have to clearly define a different set of business requirements for a transformation of the identity management platforms, workflows, and business processes to keep up with AI. The call for a distinctly different set of requirements for identity security should start with the CISO to influence the redesign of identity security for non-human identities.

The history of IAM governance dates back to the 1960s when Fernando Corbató created the first password at MIT used for a file-sharing system for research scientists. He was promoted in 1964, likely leading to the creation of the identity administrator role to manage access requests for the file share system. In subsequent years, this model was used for all enterprises with systems in place. More administrators were added as the enterprises grew. All of the enterprise functions, processes, and platforms required human judgment to drive transactions. Human-driven transactions are expensive, time-consuming, and prone to delays and mistakes. This is true for any function and certainly true for identity management. Technical requirements for IAM evolved into a set of record-keeping functions for auditors and regulators. Unfortunately, this history in IAM has led to an unsustainable model, unable to keep up with the demands of AI and business growth.

II. The New Architectural Paradigm: From Identity Governance to Identity Security

Enterprises need a data lake architecture that captures entitlement usage attributes across many different domains, enabling the establishment of attribute patterns. The attribute activity pattern represents a baseline for normal behavior for an individual user, an API transaction, an AI agent, and an MCP server. Measuring pattern deviation is relatively straightforward and results in a number that can be used as a risk score. The more significant the pattern deviation, the higher the risk and the greater the likelihood that the real-time behavior is coming from a threat actor using compromised credentials and not the actual person.

Using streaming data to compare real-time data to established patterns for specific attribute activity is not difficult and can be accomplished with deterministic models. Large Language Models and/or AI agents are not required for this architecture, but they can help in the implementation.

The key to make this data lake architecture effective is to capture activity data in real time streaming components that can then be compared to established patterns. Starting with a few attributes and consistently growing the attributes will consistently improve the results of determining pattern deviation more accurately. A weighting can be applied to each attribute for higher accuracy (GPS data = 10%, online entitlement access = 75%). Then, to use the streaming data in comparison with the pattern to trigger automated workflows based on a pre-determined threshold. For example, if the pattern deviation is 75% or higher, then the entitlement in use gets revoked through the automated workflow. If the pattern deviation score is higher than 60 but less than 75%, then an incident record is opened and fed to the SEIM. Identity security administrators evolved into analysts who make adjustments to the thresholds based on trend analysis.

The most significant benefit of reducing dependence on human-driven transactions for Identity Security is the speed of response and the associated cost reduction for response. Viewing the automated transactions in this new architecture increases the opportunities to capture the reduction in operating costs by lowering dependence on humans to initiate transactions, and applying the savings toward identity security redesign implementation costs. CISOs can justify the costs for the implementation of the new architecture and retooling of identity management staff by harvesting the savings in operational cost. Implementation costs are lower as new processes and infrastructure are implemented, and the improvement in speed and cost contributes to annual reductions in cost.

A second important benefit is that the use of automated workflows triggered from pattern deviation (in milliseconds) improves the speed of response to increase protection from threat actor attacks using AI. The velocity of threat actor attacks will continue to improve enabled by AI.

CISOs have an opportunity to document the current cost of core identity transactions (provisioning, certifications, and de-provisioning) that are driven by humans using judgment and calculate the transaction cost. This provides a baseline for improvement when automated transactions operate for pennies while human-initiated transactions cost hundreds of dollars for each transaction. The improvement in unit cost alone for these core transactions supports a compelling business case for the new architecture.

III. The Strategic Shift: Dynamic Provisioning & Continuous Validation

One of the initial use cases to reduce the unit cost of identity transactions is to implement dynamic provisioning at scale. The legacy architecture for IAM and associated business processes provides a workflow of requests for additional entitlements from network users and obtains the specific approvals from data and application owners, in addition to the requester's leader. There are weeks of wait time (until the approvals are made) as requests sit in email inboxes until they are acted upon. Dynamic provisioning starts with the creation of a risk score for every registered user based on their use of existing entitlements housed in the data lake of entitlements. All requests for additional entitlements coming from users with a pre-determined risk score below an established threshold get provisioned in minutes with no human review. This reduces the backlog of requests, delivers capabilities in minutes vs. weeks, applies an effective risk score for consistently better risk management at a lower cost, resulting in higher user satisfaction. Some enterprises consider dynamic provisioning to provide a trifecta of benefits:

- ① **Lower unit cost of transaction**
- ② **Immediate access to entitlements, shrinking the backlog, while offering higher user satisfaction**
- ③ **More consistent application of risk management and resilience**

Some enterprises start with dynamic provisioning to provide opportunities to harvest benefits early in the implementation of the new architecture. Others focus on the benefits of speed for entitlement enablement and the associated productivity gains. Many of the newer data lake-based identity platforms offer an additional benefit of deploying least privilege, shrinking the attack surface, and eliminating the need for certification workflow in email. The platform records the enablement of entitlements from dynamic provisioning and then identifies all entitlements that have not been used within 90 days and automatically initiates a workflow to revoke the specific entitlement due to non-use. A key benefit and byproduct of this use of least privilege is the elimination of annual certifications for the entitlements not in use. Typically, senior executives carry the burden to review entitlements for their respective teams and often the information provided is insufficient for them to understand, resulting in the retention of entitlements that are neither used nor needed.

One of the most significant benefits of a data lake architecture for identity security comes from the addition of continuous validation to the privilege access management capabilities that exist today. The majority of PAM solutions available in the market or built by enterprises use a password vault to assign specific and sensitive entitlements to privileged users for specific periods of time (4 hours, 8 hours, 2 days, 5 days, etc.). This enables users to perform their specialized role within specific windows of time, reducing the attack surface by avoiding permanent use of sensitive entitlements.

Enterprises that add a layer of continuous validation of the identity using the privileged entitlements for the duration of their entitlement period. Their online usage attributes are compared to the established patterns and any significant pattern deviation detected triggers the revocation of the entitlement. One of

the key benefits of this capability is to frustrate threat actors using RaaS to escalate privilege to exfiltrate data and disrupt data replication and recovery infrastructure. Continuous validation, instead, disrupts the RaaS threat actors that use AI in their attacks, designed for success in minutes. The additive friction of continuous validation will force threat actors to pursue easier targets.

The results of continuous validation capabilities additive to conventional PAM functionality, include:

- ① **Improved speed of detection, response, and recovery, disrupting threat actors' attempts to escalate privilege**
- ② **Lower costs of incident response and recovery**
- ③ **The implementation of a digital immune system that operates without human direction**

Dynamic provisioning and continuous identity validation capabilities provide core components of a digital immune system that operates in real time to prevent attacks without the need for human intervention. The real-time response and recovery steps enabled enables an enterprise through a continuous protection capability that is not dependent on humans to consume information, wrap it with context, and then initiate actions. This is similar to how your body's immune system springs into action by producing white blood cells and antibodies to attack bacteria or viruses that you are exposed to. The year over year growth of agentic AI web traffic is 7,851%, according to [The 2026 State of AI Traffic & Cyberthreat Benchmark Report](#).

The digital immune system enables enterprises to operate faster than threat actors as technology usage continues to evolve. Enterprise use of agents at scale requires attributes of the digital immune system. It requires a way to identify and track agent activity. It will ultimately require the use of agents designed for the enforcement of guardrails necessary to prevent bad business outcomes with agents. Our body is a comprehensive system dealing with billions of entities at once that generally works effectively, and our immune systems do not take direction from our brains. Agents deployed for governance represent the next generation of controls. The adoption of agentic solutions will accelerate the need for agents deployed for governance and control. They too, will be part of the digital immune system for enterprises.

IV. The Economic Engine: Self-Funding Transformation

CISOs in the past often deferred to their identity security team leaders to define the requirements for new identity management capabilities. This resulted in a long list of features and functions to improve on current processes and the limitations of the infrastructure in place.

An alternative approach worth considering is for the CISO to define a set of high-level economic targets for identity security capabilities going forward that the identity team can add to and share with experts from within and external to the enterprise. This becomes a blueprint for identity transformation. The CISO is in a position to force the identity leaders to consider redesigning core processes, moving from record keeping to a data lake architecture of entitlement attributes.

The requirements from the CISO need to include the support of non-human identity discovery, registration, and policy management. The CISO can determine the order of priorities for dynamic provisioning and continuous validation to use the benefits of lower operating costs and productivity gains to offset the implementation costs of the transformation.

A well-designed Identity Security transformation program will deliver substantially greater benefits at a lower operating cost and fewer resources with higher skills. These benefits are not dependent on AI; however, AI capabilities will improve the benefits.

The operational cost savings in this approach are substantial for larger enterprises (tens of millions of dollars in annual savings) with large existing identity teams. The requirements can drive decisions on the choice of professional service providers and identity security vendor platforms, while also providing a road map for a new architecture. Enabling identity professionals to both select and realize their professional development objectives will improve the probability of a successful outcome. Professionals at all levels are learning how to use AI to improve their capabilities and output at work and at home. This approach accelerates the learning opportunities for identity professionals.

Data science skills and AI experience are essential to both acquire and develop the necessary talent for a successful transformation. Educational content should be up to date and available to identity professionals, accelerating the transformation and skill upgrade.

CISOs should be careful and discourage a detailed discussion of features and functions from the available identity platform vendors and professional service specialists in identity. Discussion with external providers should start with the CISO's requirements for economic benefit, focused on the desired economic outcomes. Vendor partners and service providers should understand and contribute to the outcomes desired by the CISO and CIO.

V. Conclusion: A Call to Action for the CISO

Past practices of the pursuit of technology enhancement and technical upgrades for identity security will not meet the dynamic needs for transformation and AI enablement at enterprise scale. Desired technical upgrades do not represent a compelling business case for an identity security program. The pursuit of a digital immune system that responds to threats in milliseconds while improving stakeholder satisfaction and significantly lowering operating costs is a more compelling business case.

This approach gives the CISO more flexibility to fund some of the investment costs with an offset in some of the operating cost savings and productivity gains. One of the techniques for a CISO is to define a set of principles for the program based on the pursuit of lower operating costs. The principles can be shared with the identity security team and evolve into more specific requirements. Principles are easier to create and provide a direction for the transformation that improves the probability of harvesting the desired benefits. Here are examples of principles for identity security:

- ① **We will transition to identity transaction approvals using individual risk scores, based on activity patterns**
- ② **We will eliminate the need for people to review access decisions in the majority of access requests and make better decisions based on activity data and established risk scores**
- ③ **Total operating costs for identity security will be reduced by a minimum of 20%**
- ④ **The new capabilities will support the effective launch of AI agents at enterprise scale with identification, registration, and policy management**
- ⑤ **The architecture must include a data lake for attribute capture and management for all identities (human and non-human)**
- ⑥ **Entitlement information will be easy to understand and manage for all users**
- ⑦ **We will manage patterns of behavior for all registered identities across many attributes to improve risk decision-making**
- ⑧ **We will revoke all entitlements not used within 90 days**
- ⑨ **We will attract and develop data scientists and AI experts for identity security**
- ⑩ **Operating cost savings for identity security operations will be used to offset the investment in new platforms and capabilities**

Identity security is on the critical path to agentic AI to effectively manage non-human identities at enterprise scale from both a registration and policy enforcement perspective. Identity security implies the redesign of core identity management practices and platforms to lower costs, improve capacity and quality, and increase cyber resilience. Identity security includes the management of AI agents for the enterprise. Gartner calls “agent management platforms” the most valuable real estate in AI.

This represents a break from the past use of recordkeeping platforms to satisfy compliance needs. CISOs play a critical role in defining a set of audacious goals to pursue and make identity security a core competency for the enterprise’s pursuit of benefits from AI.



Jim Routh

Fellow, Institute for Critical Infrastructure Technology

Jim Routh serves on the Boards of Savvy Security, Accountable Digital Identity Association, and the Global Resiliency Federation. He is the former Board Chair for the Health Information Sharing & Analysis Center (H-ISAC) and former Board member for the Financial Services Information Sharing & Analysis Center (FS-ISAC). Jim is the Chief Trust Officer for Saviynt. Jim is a former CSO/CISO for American Express, DTCC, KPMG, Aetna, CVS, and MassMutual. Jim brings a vast business and technology background to the boards and senior executives and is considered a digital and cyber security industry expert and thought leader. Jim is an advisor for Wiz, Netskope, Armis, Transmit Security, Security Scorecard, Gurucul, Data Theorem, Panaseer, Legit Security, CodeZero, Picnic, and Rekin. He serves in an advisory capacity and is an investor for cyber-specific venture funds including Syn Ventures, CyberStarts, Security Leadership Capital, Ballistic Ventures, and Rain Capital. Jim is an ICIT Fellow and an adjunct faculty member, and he teaches cybersecurity at the NYU Tandon School of Engineering. Jim also mentors over 90 cybersecurity professionals and students.



ICIT

The Institute for Critical Infrastructure Technology (ICIT) is a nonprofit, nonpartisan, 501(c)3 think tank with the mission of modernizing, securing, and making resilient critical infrastructure that provides for people's foundational needs. ICIT takes no institutional positions on policy matters. Rather than advocate, ICIT is dedicated to being a resource for the organizations and communities that share our mission. The views and opinions expressed in this essay are solely those of the author(s) and do not necessarily reflect the official policy or position of ICIT. Any assumptions made within the analysis are not reflective of the position of any entity other than the author(s).

To learn more, please visit www.icitech.org



CyberRisk Alliance

CyberRisk Alliance provides business intelligence that helps the cybersecurity ecosystem connect, share knowledge, accelerate careers, and make smarter and faster decisions. Through its trusted information brands, network of experts, and innovative events it provides cybersecurity professionals with actionable insights and act as a powerful extension of cybersecurity marketing teams. CyberRisk Alliance brands include SC Media, the Official Cybersecurity Summits, InfoSec World, Identiverse, Cybersecurity Collaboration Forum, Cybersecurity Collaborative, Security Weekly, Channel Pro Networks, ChannelE2E, MSSP Alert, and LaunchTech Communications.



ICIT

www.icitech.org