REPORT

2026 Identity Security Trends & Predictions

Savıynt™

Contents

2026 Identity Security Trends & Predictions

Enterprises rolled out AI faster than any major technology in history, driven by an urgent need for efficiency. But this speed created a dangerous gap between employees adopting autonomous agents and copilots, and security teams that lacked the frameworks to govern them.

In 2026, the experiment ends, and the reality of risk begins. This report outlines why identity security is no longer just a support function, and is instead the essential foundation for AI growth. We are moving toward a world where non-human identities (NHIs) and autonomous workflows redefine the security stack.

The traditional boundaries of the network have dissolved. We are now managing a “triple threat” of complexity:

- **Agentic Risk:** AI agents act with administrative privileges that often exceed those of their human creators.
- **The Governance Deficit:** Organizations are struggle to govern machine-speed identities using human-speed manual processes."
- **The Visibility Gap:** Most leaders cannot identify how many autonomous agents are currently active or what data they are accessing.

Navigating this triple threat requires a fundamental re-engineering of how we verify and monitor access. These defining trends outline the evolution of identity from a perimeter tool into the primary operating system for modern security and resilience.

Insights from Saviynt's Identity Security Experts

Simon Gooch

Field Chief Information Officer

> A complete top-to-bottom re-engineering of business processes is required to fully capitalize on the new paradigms AI enables.

David Lee

Field Chief Technology Officer

> AI is hands down the fastest and most pervasive piece of technology in modern history.

Jeff Margolies

Chief Product & Strategy Officer

> We're moving from AI hype to applications with real business value.

Jim Routh

Chief Trust Officer

> The extended attack surface created by agents is emerging as a significant risk for enterprises large and small.

Vibhuti Sinha

Chief Product Officer

> Identity Security will emerge as the biggest threat as businesses prepare and adopt Agentic workflows with AI agents as the new identities powering them.

Henrique Teixeira

SVP, Strategy

> The notion of progress is evolving, transitioning from mere optimization to complete reinvention.

Trend 1: Bad Actors Will Target AI Identities

We used to build security around the human person, but a new resident has moved into the network: the AI agent. These non-human identities often operate with elevated administrative privileges - sometimes exceeding the authority of the individuals who created them. Yet, they frequently do the oversight and lifecycle controls applied to human accounts.

The challenge is that human identity governance itself remains a work in progress. Many teams are still managing manual access reviews, incomplete lifecycle processes, and entitlement sprawl. Now they are being asked to extend governance to identities that move faster, act autonomously, and do not fit frameworks built for people.

Key Insights:

- Most organizations can't say how many agents are running or what decisions they are making.
- Attackers use prompt injection and model manipulation to turn agents into insider threats.
- Dormant machine credentials and unsecured agents give attackers easy access points.
- Teams are being asked to extend governance to autonomous identities that move at machine speed, even as human identity governance remains a work in progress, plagued by manual reviews and entitlement sprawl.
- AI identities require the same governance rigor that organizations have spent decades building for human users.

Priority Actions:

- Treat identity as infrastructure by integrating non-human identities (NHI) into the core security stack.
- Centralize lifecycle management to provide a single view of both human identities and NHIs.
- Implement dynamic privilege enforcement and continuous monitoring as baseline requirements for autonomous workflows.

> **AI is driving automation, modernization, and stronger security postures, but it's also a threat to cybersecurity and IT ecosystems. Organizations are already facing Identity targeting, automated attacks, and sophisticated AI-driven kill chains, and this will only amplify in 2026.**
>
> - **Vibhuti Sinha,**
> Chief Product Officer

Trend 2: MCP Will Accelerate and Secure AI Innovation

MCP creates a standard way for AI agents to connect directly to applications, tools, and data sources across the enterprise. It plays a role for autonomous systems similar to what APIs once did for cloud platforms. Instead of routing work through a human, MCP allows machines to work with machines.

An MCP connection carries real authority. It allows an agent to retrieve data, trigger workflows, and act inside critical systems without a person in the middle. When those connections are poorly governed, they become high-value access paths. If compromised, they offer attackers a way to influence trusted systems at machine speed and largely out of sight.

In 2026, machine-to-machine dialogues will become the new frontier of risk.

Key Insights:

- By removing the person in the middle, we remove the pause where judgment once lived.
- MCP tokens, credentials, and access rules are becoming primary targets because they enable agents to operate within critical systems at machine speed.
- Most organizations lack visibility into this layer, unable to see which agents are speaking or what permissions they carry.

> **The emergence of MCP as a standard protocol in the agentification of work offers hope for enterprises to better manage the extended attack surface created by AI. Agents and MCPs must be part of identity security to control how agents interact with systems and what controls are required.**
>
> - **Jim Routh,**
> Chief Trust Officer

Priority Actions:

- Treat MCPs as part of the identity surface, not as an application detail.
- Bring MCP under the same governance disciplines as any privileged access path, including strong authentication, clearly defined scopes, least-privilege enforcement, and continuous monitoring of how agents are using it.

> **MCP is the equivalent of HTTPS for our agents. It enables our agents to access any application we want them to without going from AI to human and back to AI. It's just AI to AI.**
>
> - **Jeff Margolies,**
> Chief Product & Strategy Officer

Trend 3: Data Security Will Return as a Frontline Challenge

For years, organizations deferred data classification and cleanup because it was slow and easy to deprioritize. AI has changed that calculus by making buried data immediately findable and actionable.

"Remember those data security and classification projects we abandoned ten years ago? Well, guess what, they will be back with a vengeance," says David Lee. "Why? AI tools are excellent at correlating large amounts of data in a very short amount of time. So, those sensitive files you created four years ago on SharePoint that you forgot about are now accessible by everyone in the company with a simple prompt."

Key Insights:

- AI does not distinguish between what it can access and what it should access, making data security a frontline challenge.
- AI is a master of correlation. A file abandoned years ago is no longer buried; it is a single prompt away from being seen by the entire company.
- When an agent acts, it inherits the permissions of its creator, including both intentional and accidental permissions. This turns every instance of excess privilege into an instant exposure.
- Cleaning the environment is no longer optional. This requires revisiting identity metadata, classifying access patterns, and enforcing least privilege with actual rigor.

Priority Actions:

Clean your data. That means:

- Improve identity data hygiene.
- Tighten access controls
- Enforce least privilege
- Establish clear ownership over who and what can access sensitive information.

> **“Beyond the data itself, the real challenge is that our old processes were never designed for this level of speed or exposure. We can’t keep patching the past. We have to rebuild our foundations for an AI-powered future.**
>
> - **Simon Gooch,** Field Chief Information Officer

Trend 4: Breaking Down Siloes and Moving Toward Zero Trust

Fragmented security was always a liability, but the speed gap created by AI makes it a crisis. Most organizations still rely on a collection of tools that lack sufficient context sharing. That may have been manageable when threats moved at human speed, but it becomes a problem when attackers use automation and AI to probe, pivot, and escalate faster than teams can respond. Fragmented security transitions from a manageable liability to a critical risk.

Key Insights:

- Important signals are often delayed or lost because a risky access pattern appearing in one system may never reach another.
- As organizations adopt agentic AI, the ability to enforce just-in-time access and least privilege becomes essential to maintaining control.
- Security platforms must work together to consistently govern access decisions and automated workflows, fulfilling the original goal of zero trust: verify everything and assume nothing.

Priority Actions:

- Identity is the common layer that connects disparate systems. Shift toward a unified control plane where identity context flows freely into detection and response systems.
- To manage agentic AI, implement just-in-time access controls to ensure permissions are active only when necessary.
- Prioritize interoperability between security platforms to fulfill the zero-trust requirement of verifying every action across the entire environment.

> **Platformization and interoperability help teams achieve more precision, more productivity, and effectiveness because they're working together. I see the security industry as a whole making a lot of progress around this in 2026.**
>
> - **Jeff Margolies,** Chief Product & Strategy Officer

> **AI will serve as a driving force for zero-trust architectures, with AI as the connective tissue. The more organizations lean into agentic AI use cases, the more the need to have concepts like just-in-time access and least privilege will become critically important.**
>
> - **David Lee,** Field Chief Technology Officer

Trend 5: AI Brings Identity to the Center of Cybersecurity

Identity is the domain that determines who or what is taking an action and whether that action should be allowed. While AI creates new risks, it is also the only way to provide the scale required to solve the governance problems it has accelerated.

Key Insights:

- Identity is no longer just a gatekeeper; it is the strategy that enables safe, efficient AI adoption.
- AI's ability to correlate massive amounts of data can finally automate the discovery of orphaned accounts and unowned access that have plagued programs for years.
- Organizations that treat identity as infrastructure will scale AI safely; those that treat it as a compliance exercise will struggle to maintain control.

Priority Actions:

- Leverage AI tools to address long-standing governance gaps, such as identifying orphaned accounts and classifying complex access patterns.
- Move identity management from a compliance-focused exercise to a central architectural role that governs all automated and human workflows. Building this foundation is necessary to scale AI innovation without increasing organizational risk.

> **Identity has the potential, and the duty, to become the operating system and resilience mechanism of the AI era. While its primary function is security, identity also acts as a catalyst for efficiency, accelerating how individuals can be productive and access more information and systems, especially in the age of AI.**
>
> - **Henrique Teixeira,** SVP, Strategy

Conclusion: From Experimentation to Resilience

The rapid adoption of AI has ended the era of passive identity management. The triple threat of agentic risk, governance deficits, and visibility gaps cannot be solved with legacy tools or human-speed processes.

In the AI era, identity isn't just supporting your security strategy - it is the strategy. By integrating non-human identities, unifying visibility, and enforcing adaptive security, organizations can build a system of trust that scales with the speed of innovation.

The Identity Security Platform for the AI era.

Platform | IGA | ISPM | Application Identity Security | PAM | External Identity Management

Saviynt provides the converged platform enterprises use to govern human and non-human identities, protect critical systems, and bring structure to increasingly complex digital environments.

By unifying identity governance, privileged access, and identity security posture management, Saviynt gives security, IT, and business leaders a single system of trust for how access is granted, monitored, and enforced across people, machines, and AI.

Trusted by leading global enterprises and public sector organizations, Saviynt helps organizations scale innovation safely, reduce risk, and meet regulatory demands in a world where identity defines security.

Contact Us

Headquarters
1301 E El Segundo Blvd,
Suite D El Segundo, CA 90245
United States

310.641.1664
info@saviynt.com
www.saviynt.com