Saviynt™

2026

CISO AI Risk Report

Research by

Cybersecurity INSIDERS

Introduction

Many security leaders didn't authorize AI expansion. It happened around them. Someone plugged in a copilot in a SaaS tool or an engineering team tested an agent or a business unit installed an assistant without waiting for approval. None of these choices feel significant in isolation, but together they create systems acting on behalf of people, without the structures we rely on to govern human access.

In our survey of more than 200 CISOs and security leaders, the same concerns surfaced repeatedly. AI systems already have meaningful access, often with privilege levels no one explicitly granted. They generate activity that can be difficult to trace, behave in ways that don't match human patterns, and sometimes leave behind incomplete or temporary records. None of this is catastrophic on its own, but it complicates the basic questions security teams rely on, namely: "Who did this?" and "Should this action have been allowed?"

Leadership teams are worried because AI is already reading customer data, modifying configurations, invoking APIs, and chaining actions together in ways that are difficult to trace back to a single owner. AI identities don't behave like human users or traditional service accounts.

Security leaders are clear-eyed about the challenge. They want workable visibility, a way to understand how these systems operate, and a practical path to keep privileges from quietly expanding beyond what anyone intended. This report focuses on what leaders are dealing with right now. AI is active in production environments, and most organizations can't clearly explain the scope of its access.

KEY FINDINGS:

- **71% of CISOs say AI has access to core business systems,** but only 16% govern that access effectively. These agents have access without governance.
- **92% of organizations lack full visibility into AI identities,** and 95% doubt they could detect misuse if it happened. AI is already acting in environments most security teams can't monitor.
- **86% don't enforce access policies for AI identities.** Only 17% govern even half of their AI identities like human users, and just 5% feel confident they could contain a compromised agent.
- **75% have discovered unsanctioned AI tools** currently running in their environments, often with embedded credentials or elevated system access that no one is monitoring.
- **Only 25% use AI-specific identity or monitoring controls.** Most organizations are trying to manage machine-speed risk with fragmented tools designed for manual workflows and human users, not autonomous agents.

The detailed findings that follow outline where security gaps are emerging and highlight the areas leaders say they need to strengthen most. Each section includes practical steps teams can take today to improve visibility, limit unnecessary privilege, and build more dependable oversight as AI activity continues to expand.

*Throughout this report, 'AI' or 'AI identities' refers to GenAI or LLM entities, including copilots, agents, and LLM-based apps.

AI Identities Are Moving Faster Than Security

AI identities don't need a security sign-off to get into your environments. As teams continue to adopt copilots, agents, and AI-driven apps, this new class of identities is entering business-critical systems unchecked. Sometimes through sanctioned projects, sometimes without approval. Once inside, these agents can act independently or even create additional identities, which means they often begin expanding before security teams can define how to control them.

Seventy-one percent of organizations say AI tools now have access to core systems like Salesforce and SAP. But only 16% say that access is governed effectively. The result is a level of access and autonomy we would never grant a human user.

83%

of CISOs are concerned about AI Access

The problem isn't just access, but also scope. Many AI tools have write access when read-only would suffice, broad API permissions when narrow scopes would work, and standing privileges when just-in-time would be safer.

Eighty-three percent of security leaders say they're concerned about AI access. And nearly half (47%) have already observed AI agents exhibit unintended or unauthorized behavior. A third of organizations dealt with an actual security incident or near-miss in the past year. These aren't hypothetical risks, they're operational failures CISOs are already containing. But as AI moves faster, the risk increases. These agents act with system-level access, often with no clear owner and outside traditional enforcement models. Each AI integration becomes its own identity that needs to be governed like any other high-risk user.

Security Leaders are Concerned About AI Risk

83%

are very or somewhat concerned about **GenAI identities accessing critical systems**

71%

say **AI tools now have access to core systems** like Salesforce and SAP

47%

have already seen their **AI agents exhibit unintended or unauthorized behavior**

33%

have experienced a security incident or contained near-miss in the past 12 months

ACTION: Start thinking about AI identities differently.

AI identities don't behave like human users or even traditional machine accounts. Governing them requires a paradigm shift. This is as much a change management problem as it is a tech innovation one. Start by assessing and planning for what level of change is needed to get your AI governance in order. The teams that make progress in this new reality will develop a system for applying rigor and context to continuously validate intent, privilege, and behavior of AI identities.

The Evolution of Autonomous AI Identities

Each AI identity needs tailored governance, yet all operate far faster and more autonomously than traditional identity controls can manage.

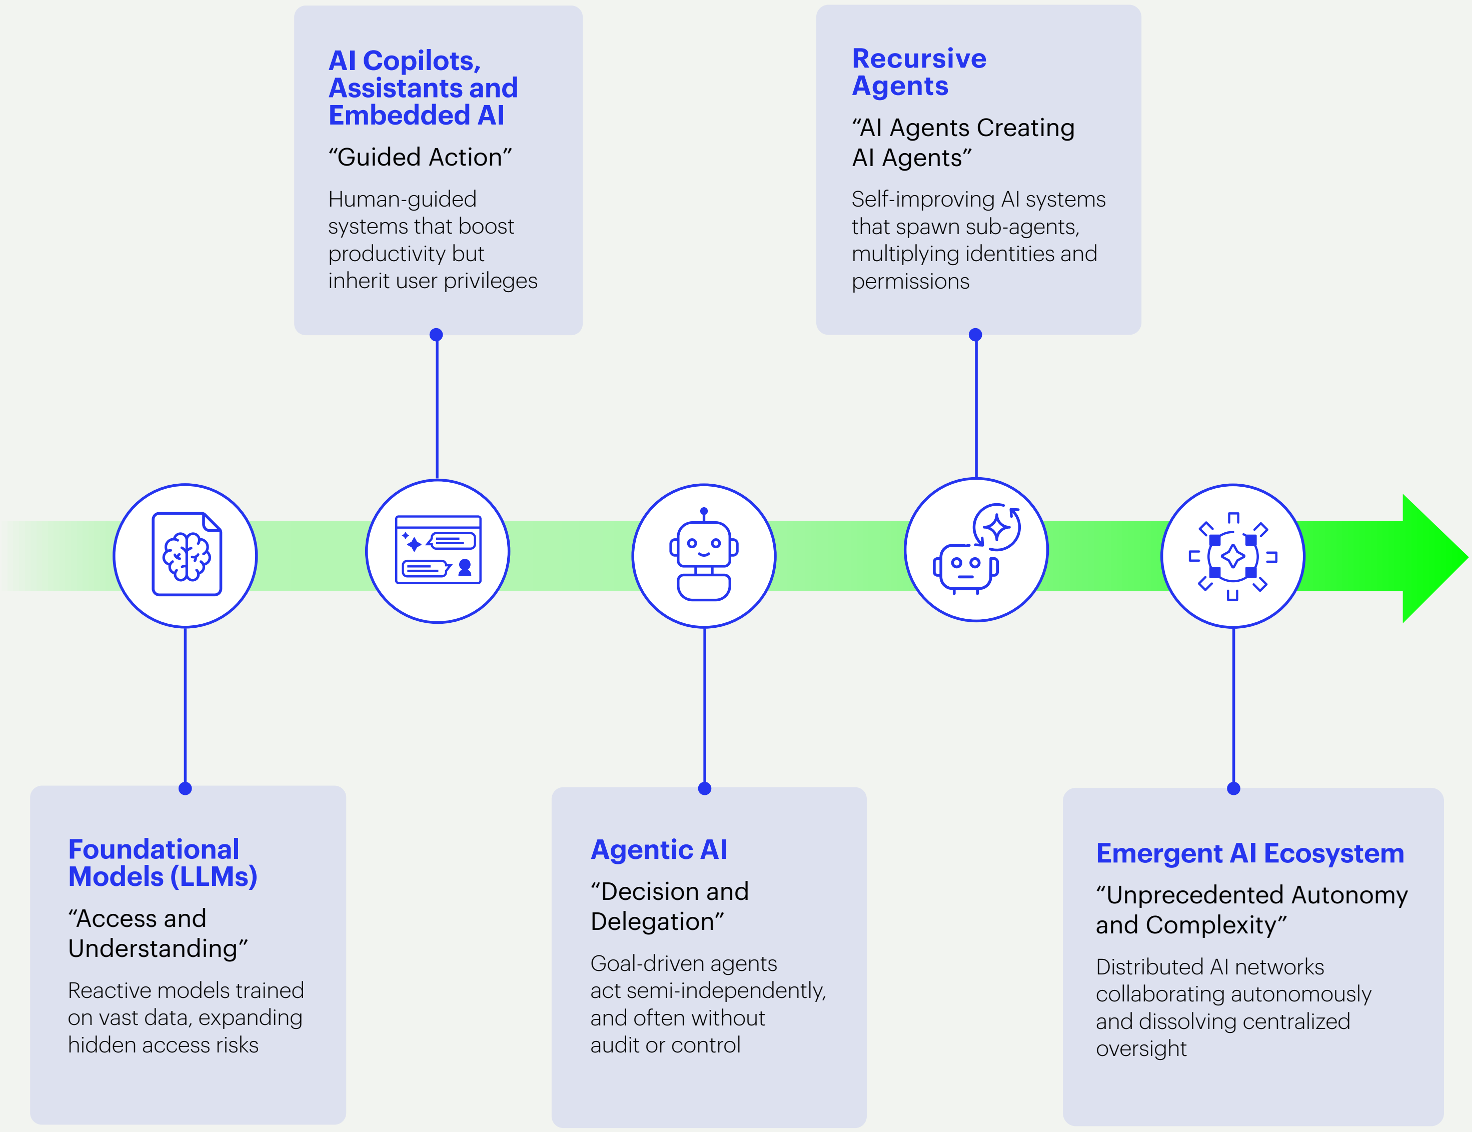

Foundational Models (LLMs)

"Access and Understanding"

Reactive models trained on vast data, expanding hidden access risks

AI Copilots, Assistants and Embedded AI

"Guided Action"

Human-guided systems that boost productivity but inherit user privileges

Agentic AI

"Decision and Delegation"

Goal-driven agents act semi-independently, and often without audit or control

Recursive Agents

"AI Agents Creating AI Agents"

Self-improving AI systems that spawn sub-agents, multiplying identities and permissions

Emergent AI Ecosystem

"Unprecedented Autonomy and Complexity"

Distributed AI networks collaborating autonomously and dissolving centralized oversight

The Visibility Crisis: You Can’t Secure What You Can’t See

Before you can control AI access, you need to know where those agents are operating. For most security teams, that’s still unclear.

The report shows 92% percent of organizations lack full visibility into their AI identities. And 95% percent doubt they could detect or contain misuse. These agents aren’t just connected to SaaS tools; they’re invoking APIs, writing to systems, and operating across environments with no clear owner.

Part of the problem is architectural. Traditional IAM tools were built for people. They assume login events, human workflows, and static roles. They weren’t designed to track systems that create their own accounts, act autonomously, or escalate privileges without asking. What’s more, AI workflows are non-deterministic, making “normal” behavior difficult to baseline.

Most teams don’t have a complete picture of what AI is doing, or where.

That’s why visibility and monitoring are now top investment priorities. According to the data, if budget wasn’t an issue, 73% of CISOs would be focusing on API and workload identity discovery and inventory. Another 68% would prioritize posture analytics and continuous monitoring. The goal is not just detection, but to answer simple questions in real time: What does this AI identity have access to? Who authorized it? What is it doing? Does it still need this access? Has it created other AI identities?

AI Visibility Crisis in Numbers

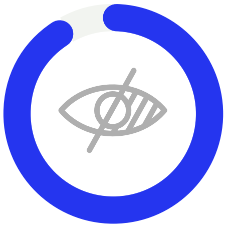

92%

lack full visibility into AI identities

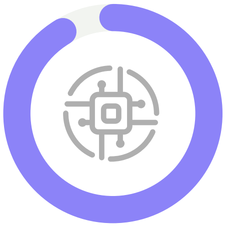

95%

doubt they could detect or contain misuse

ACTION: Get the full picture of who and what is operating inside your organization.

Start by finding, identifying, and labeling every identity you have. This includes separating AI identities from human users, and distinguishing between different types of AI (LLMs, agents, bots, and so on). Make sure this classification happens inside every application, not just at the system level. Shadow AI will spread quickly, and missing even one app or identity can trigger a chain of problems that’s hard to unwind later. From there, ongoing visibility requires a process for continuous discovery and proactive posture management.

AI Governance Is the Weakest Link

AI tools are now acting with real authority and privilege. But in most organizations, they're doing it without governance. Eighty-six percent of security leaders lack or don't enforce access policies for AI identities. Only 19% govern even half of their GenAI accounts with the same rigor they apply to human users.

A mere 5% feel prepared to contain a compromised AI agent. That's not surprising when most of those identities aren't being provisioned, certified, or deprovisioned through structured workflows.

A growing number of teams are bringing AI into their existing governance models by applying identity lifecycle controls. But parity with human users is just a starting point. AI identities act independently, often across systems, on behalf of users or other agents. When one AI invokes another, attribution becomes unclear and privilege boundaries blur. AI identities need controls built for autonomy and policies that account for delegated actions, escalation paths, and behavior at machine speed.

Privilege without policy is becoming the norm.

Does your organization have formal policies specifically governing GenAI/LLM agent identities and access?

86% of security leaders lack or don't enforce access policies for AI identities

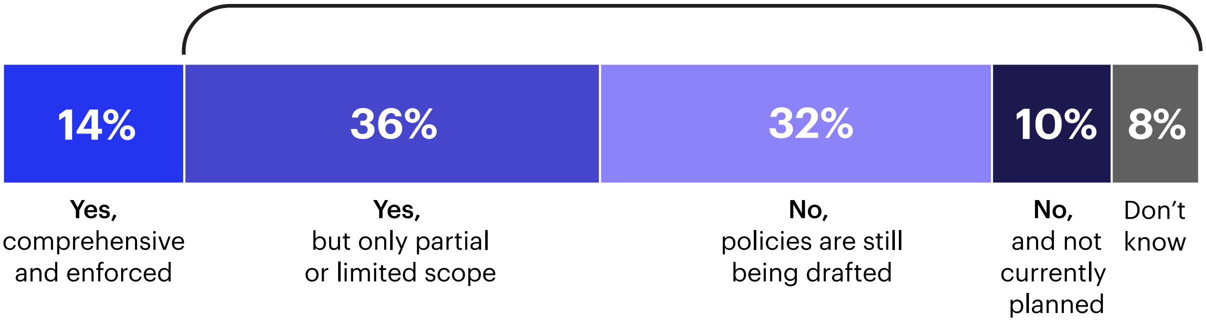

ACTION: Understand how AI is actually being used inside your organization.

Find out which apps AI agents are already touching, why they were introduced, and who put them there. Those answers will help you define the right rules and guardrails for each AI identity. And make sure your governance keeps access limited to only what's truly needed and nothing more. Ultimately, the goal is to apply a unified system of lifecycle, least-privilege, and certification controls to both human and AI identities.

Shadow AI Is Already Inside

In many organizations, AI isn't just being rolled out by IT through sanctioned projects. Instead, "Shadow AI" is showing up through unapproved AI tools that teams bring in on their own. Three out of four CISOs have discovered unsanctioned GenAI tools already running in their environments. Another 16% aren't sure, which likely means they have the same problem.

These AI tools aren't confined to browser-based assistants. They often come with embedded credentials, API integrations, or OAuth tokens that plug directly into enterprise systems. In many cases, they operate with elevated permissions and completely outside the standard provisioning workflows that exist for human users. Each AI tool also represents a trust relationship with an external, third-party provider whose security practices and data handling are outside your control.

Unsanctioned AI tools are spreading.

Security teams are beginning to treat Shadow AI like any other unmanaged identity. Scanning for unknown accounts and tokens, building inventories, and applying policies after the fact. But in most organizations, AI tools are already deeply embedded, and spreading faster than current controls can catch up.

Has your organization identified shadow or unsanctioned GenAI/LLM tools in use?

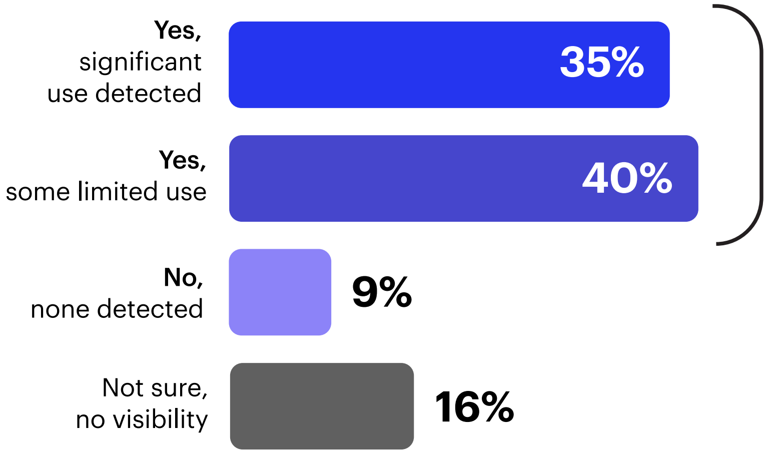

75%

of CISOs discovered unsanctioned GenAI tools in their environments

ACTION: Help stakeholders implement AI tools.

People are already using AI, whether it's approved or not. Instead of fighting that reality, give them a clear, safe way to bring new tools into the company. Set up a small governance group, define a straightforward review process, and you'll prevent issues long before they show up. From there, evaluate each tool for access scope, embedded credentials, and data exposure so you can manage risk before it spreads across systems.

CISOs Can’t Rely on Legacy Tools Anymore

Most identity tools were designed for human users, not autonomous AI systems. Many organizations are still trying to manage AI risk with tools designed for a different era and for environments comprising on-premises systems, human users, and static access.

Sixty percent still use traditional login-based authentication patterns like session management or password policies for AI identities that instead need API-first controls, like token lifecycle management, scope-limited authorization, and runtime enforcement. Only one in four organizations use AI-specific monitoring or controls today.

In practice, this means AI agents act at machine speed while security enforcement still happens manually, if at all. This results in disconnected point solutions for access, privilege, and governance, creating gaps in visibility, policy alignment, and response time.

Only
1 in 4
organizations use AI-specific monitoring or controls today

To close those gaps, security leaders are starting to unify identity oversight by pulling access logs, privilege usage, and account behavior into a single layer of context. For most teams, this oversight is still scattered across tools that don’t talk to each other. So, by the time they assemble the full picture, the AI identity has already moved.

What governance solutions are you currently using to manage GenAI/LLM identities?

60%	Access Management (SSO/MFA)
44%	Privileged Access Management (PAM)
34%	Identity Governance & Administration (IGA)
25%	GenAI-specific access monitoring tools
22%	No governance solutions in place for GenAI/LLM identities
17%	Identity Security Posture Management (ISPM)

ACTION: Take an honest look at the systems you already rely on.

Most organizations are juggling too many point solutions that barely keep up with today’s human users and applications. Add AI identities on top of that, and the cracks in those older tools become impossible to ignore. Let this be the moment to straighten out your entire identity security program, retiring the outdated IGA tech and moving toward something built for how your organization actually works today. The report reflects a move (31%) toward unified identity platforms that converge IGA, PAM, and access analytics, consolidating what’s currently spread across disconnected tools. (See page 11, “How CISOs Are Responding.”)

The Path Forward: Identity as the Enforcement Layer

Perimeter controls don't follow AI into cloud platforms, and device policies don't apply to headless agents. Identity is the most consistent enforcement layer that remains where access decisions, privilege boundaries, and audit trails converge across environments.

Security teams are already shifting their priorities accordingly. Seventy-three percent of CISOs are investing in identity discovery and inventory. Sixty-eight percent are focused on continuous monitoring and posture analytics. Another 58% are investing in AI-driven detection. The common thread is automation and the need for controls that move as quickly as the identities they're trying to govern.

In more advanced environments, these investments are being embedded in day-to-day operations. Teams are automating AI account deprovisioning, revoking access when thresholds are breached, and enforcing least privilege in real time. The goal isn't just more visibility or reporting. It's shortening the time from detection to action while each new identity is still within a manageable window for control.

As traditional boundaries break down, identity becomes the first line of control.

What CISOs Want to Invest in for Securing AI

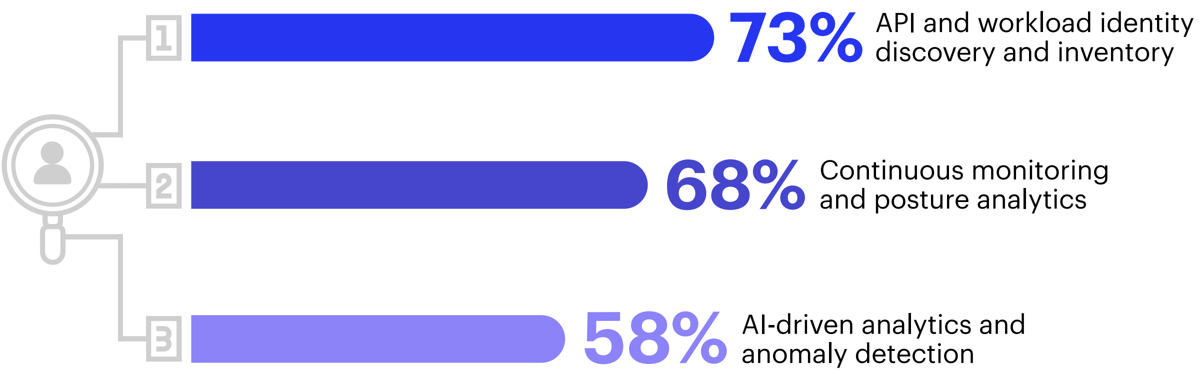

ACTION: Lay the groundwork for an identity-first approach.

Again, start with visibility. That's discovery, identification, and classification across your entire ecosystem. You can't fix what you can't see. Once you have that baseline, you can start identifying the biggest gaps like missing systems, unmanaged privileged users, or inconsistent governance. Then solve the high-risk issues first, with a clear view of where your organization needs to go. Identity can be your consistent enforcement point for AI, verifying who or what is acting, with what privilege, and for how long.

How CISOs Are Responding to AI Risk

CHALLENGE	WHAT CISOS ARE STARTING TO DO	WHY IT MATTERS
1 **AI Agents Moving Faster Than Security Control**	44% are implementing automated lifecycle governance to provision, recertify, and revoke AI access on defined schedules, treating agents like high-risk users rather than infrastructure.	Nearly half (47%) have already seen unintended AI behavior, and a third dealt with incidents in the past year. Reactive controls can't keep pace with systems that act autonomously.
2 **Limited Visibility Into AI Behavior**	73% are building comprehensive identity inventories and 68% are deploying continuous monitoring, shifting from periodic audits to real-time telemetry across AI actions.	You can’t govern what you can’t see. With 92% lacking full visibility and 95% doubting they could detect misuse, discovery and monitoring are table stakes.
3 **Governance Gaps**	Teams are extending existing IGA workflows to AI identities. Applying provisioning, access reviews, and deprovisioning to agents that previously operated outside policy frameworks.	When 86% lack enforceable policies and only 5% feel ready to contain a compromise, bringing AI under structured governance is the most direct path to reducing exposure.
4 **Shadow AI**	Security teams are scanning for unsanctioned tools and embedded credentials, then onboarding discovered AI identities into centralized control frameworks rather than trying to block adoption.	With 75% already finding shadow AI in production, the goal isn’t prevention, it’s visibility and control after the fact. Unmanaged tools become embedded risk.
5 **Fragmented Identity Tooling**	31% are moving toward unified identity platforms that converge IGA, PAM, and access analytics, consolidating what's currently spread across disconnected tools.	When 60% rely on basic SSO/MFA and only 25% use AI-specific controls, fragmentation makes it impossible to answer "Who owns this?" or "Should this access exist?" in real time.

Conclusion

For years, organizations have been in the midst of addressing the challenges of digital transformation, continuing their Zero Trust journey, and increasing compliance mandates. But AI has completely upended the cybersecurity status quo. New entities like autonomous AI-powered agents, LLMs, and MCP servers must now all work in concert with humans.

This report shows that the AI era requires a different approach to security. Traditional tools and processes can't keep pace, and organizations need a new system that can evolve and scale just as quickly as AI. The results confirm AI identities already have a concerning level of access, often privileged access. Meanwhile, most organizations don't yet have the right guardrails in place. Identity security, when done correctly, can become the one layer that provides consistent enforcement of AI identities and systems across environments.

Enterprises that move first on evaluating their current identity security posture, identifying the gaps in their current programs, and accelerating their journeys toward intelligent, unified identity security stand to gain a competitive advantage. They can enable their organization to leverage AI securely and use it as a strategic growth lever. This shift is happening now across the enterprise, just in time.

CISOs are moving toward always-on identity governance of AI operations, detecting privilege drift, enforcing policy in real time, and remediating risk automatically.

Here's what's possible with this new breed of identity intelligence and automation:

Automated de-provisioning of dormant AI service accounts

Just-in-time privilege elevation with time-bound access

Policy-driven revocation when risk thresholds are breached

Self-healing workflows that restore least-privilege baselines

These controls should match AI's speed and close the loop between detection, decision, and action.

Methodology & Demographics

The 2026 CISO AI Risk Report is based on a structured survey of 235 CISOs, CIOs, and senior security leaders across the United States and United Kingdom. All respondents represent large enterprises (5,000+ employees) across major industries, including technology, financial services, healthcare, and manufacturing. Responses were self-reported through structured multiple-choice questions examining how AI identities access core enterprise systems, where privilege drift or unintended behavior occurs, the depth of visibility and shadow AI gaps, the maturity of governance policies, and how well existing IAM tools control autonomous access.

With 235 respondents, the survey carries a margin of error of ±6.4% at a 95% confidence level, providing a reliable view of how enterprises are confronting AI-driven identity risk.

PRIMARY JOB FUNCTION

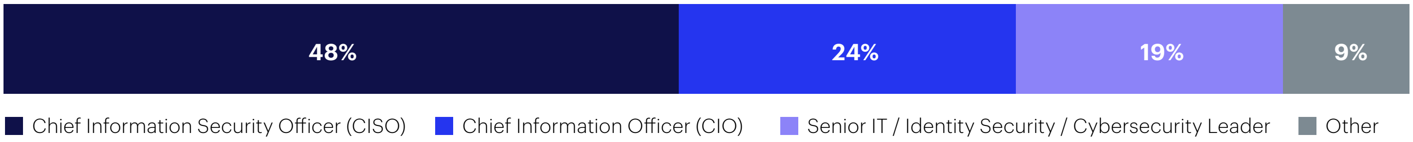

COMPANY SIZE

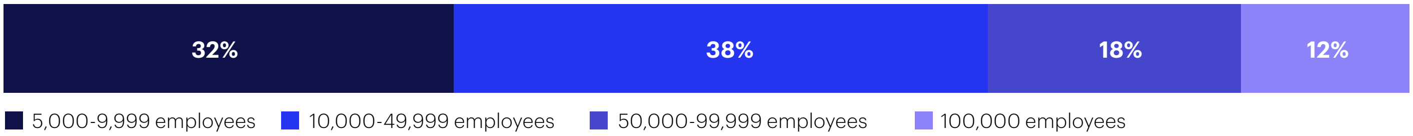

INDUSTRY

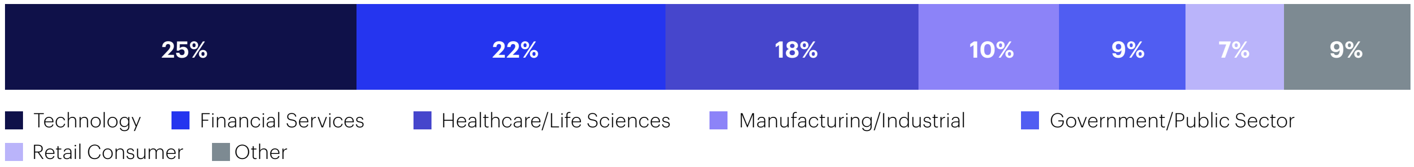

Saviynt's AI-powered identity platform manages and governs human and non-human access to all of an organization's applications, data, and business processes. Customers trust Saviynt to safeguard their digital assets, drive operational efficiency, and reduce compliance costs.

Built for the AI age, Saviynt is today helping organizations safely accelerate their deployment and usage of AI. Saviynt is recognized as the leader in identity security, with solutions that protect and empower the world's leading brands, Fortune 500 companies and government institutions.

For more information, please visit **www.saviynt.com.**

Cybersecurity

INSIDERS

STRATEGIC INSIGHT FOR CYBERSECURITY LEADERS

Cybersecurity Insiders provides independent research and analysis focused on the operational reality of enterprise cybersecurity. We gather insights from senior security and IT leaders to examine how high-level strategies translate into day-to-day execution. Our analysis identifies the measurable gaps between intended strategy and actual risk exposure, offering a credible, data-driven foundation for security decision-making and industry benchmarking.

For more information, visit

cybersecurity-insiders.com

Jan 2026 // Version 1.1