Saviynt

The 4 Identity Risks of AI Agents

And Why Most Enterprises Can't See Them

AI agents are the fastest-growing actors in the enterprise. They're also the least governed. It's predicted that 40% of enterprise applications will include task-specific AI agents by the end of 2026, up from less than 5% in 2025.[1] Meanwhile, **91% of organizations lack visibility into AI identities** already running in their environments.

Traditional identity security assumes a person logs in, does their work, and logs out. AI agents don't work that way. They inherit permissions, act at machine speed, and outlive the people who built them.

Here are the four risks expanding your attack surface right now.

Risk 1 Shadow AI you haven't inventoried

Signal: Developers, business teams, and contractors are spinning up agents without IT involvement.

What's happening: MCP servers, LLMs, and connected tools are spreading across environments that security teams can't see.

Why it matters → Visibility gap. Without discovery, governance doesn't start. **Identity Security Posture Management** is where it begins.

Risk 2 Orphaned agents with nobody accountable

Signal: Agents built for projects that ended, by employees who left, are still running with active credentials.

What's happening: Ownership is either never assigned or handled inconsistently. When the builder moves on, the agent goes untracked.

Why it matters → Accountability gap. If something goes wrong, there's no one to call. This leads to an increasingly uncontrolled attack surface.

[1] *https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025*

Risk 3 Governance that can't scale with AI

Signal: Identity programs built for human users are being asked to govern tens of thousands of AI agents or non-human identities.

What's happening: Review cycles and approval models designed for predictable behavior don't scale to continuous, autonomous activity.

Why it matters → Scalability gap. AI adoption stalls when the governance model can't grow with it Retrofitting human-centric systems won't close the distance.

Risk 4 Permissions without judgment

Signal: Agents inherit the permissions of their creators or other agents. But they don't inherit the judgment.

What's happening: Agents execute actions without pause or context. They don't question intent—they simply act, often chaining actions in ways no static policy anticipated.

Why it matters → Intent awareness gap. The fastest eCrime breakout is 27 seconds.[2] Weekly access reviews can't keep up with that.

Reality check

If you can't answer these four questions, you don't have control of your AI agents:

- How many agents are running in your environment right now, and on which platforms?
- Who owns each one, and what happens when that owner leaves?
- Can your governance solution absorb another 10,000 agents this year without breaking?
- Can you stop an agent mid-action if it steps outside its intended scope?

Start by finding what's already running.

Saviynt

The Identity Security Platform for the AI era.

Platform | IGA | ISPM | Application Identity Security | PAM | AI Agent Security

Saviynt is identity security. We secure every identity and empower global enterprises to run their businesses at the speed of AI.

Automate access, enforce least privilege, and eliminate standing risk across cloud, on-prem, and hybrid environments with Saviynt.

[2] *https://cyberscoop.com/crowdstrike-annual-global-threat-report-attack-breakout-time/*