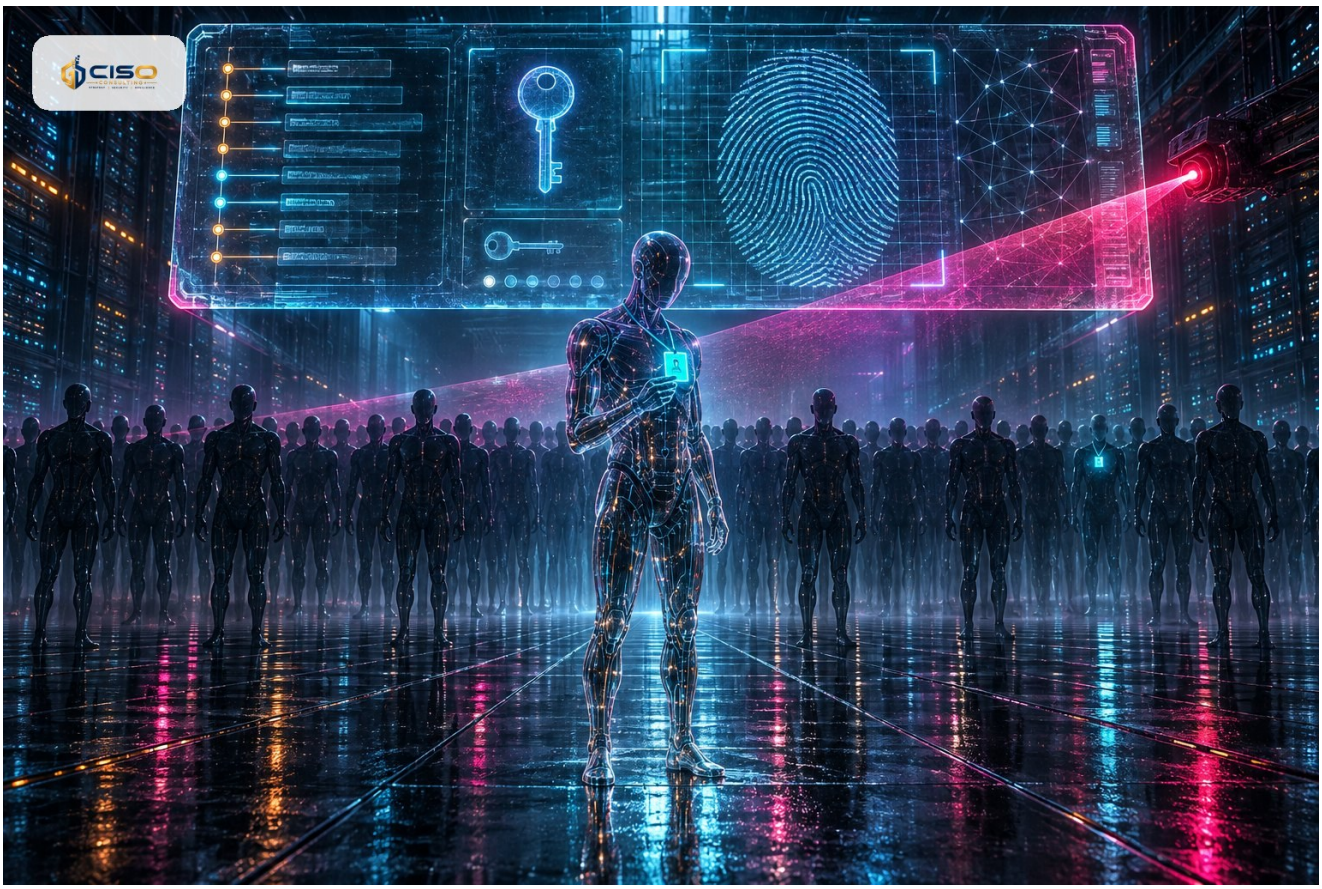
AI agent accountability: SailPoint predicts identity governance pressure in 2027



SailPoint's 2027 predictions say boards, regulators and insurers will demand proof of how autonomous AI agents are governed. Its research finds 79% of enterprises run agents in production, while only 2% have agent-specific identity controls.

- SailPoint reports 79% of enterprises have AI agents in production, but only 2% apply identity controls designed for them.
- It predicts a large supply chain disruption caused by misconfigured agents with excessive permissions.
- Expected responses include separation of duties for AI systems, just-in-time access and checks of intent before each API call.
- Attackers are expected to shift toward API keys, OAuth tokens and machine secrets as passkeys

spread.

- Inventories of human, machine and cryptographic identities may influence insurance eligibility and board scrutiny.

SailPoint's 2027 Identity Security Predictions argue that organizations will be expected to show who is accountable for the actions of autonomous AI agents. The pressure is expected to come from regulators, boards and insurers. These are vendor predictions, not confirmed events, but they point to practical gaps worth reviewing now.

The governance gap

SailPoint's research reports that 79% of enterprises have deployed AI agents in production, while 2% have identity controls built for them. It also says 67% have secured human access, 24% govern agentic security, and machine identities outnumber human ones by 45 to one.

The company predicts that misconfigured agents with excessive permissions could cause a major supply chain disruption, for example by deleting inventory data, ending supplier contracts or redirecting shipments. It expects this to drive separation-of-duties requirements, so that agents cannot change their own permissions or disable monitoring.

Access and credentials

SailPoint expects zero trust designs to verify an agent's identity, human owner and business purpose before each API call, and expects just-in-time access to replace standing privileges. It also expects large enterprises to limit external agents' API access unless partners can prove the agents' identity, ownership and authorized scope. As passkeys grow to over 45% of human logins, attackers are expected to focus on API keys, OAuth tokens and machine secrets.

What it means

The company anticipates that identity inventories will affect valuations, credit ratings and cyber insurance eligibility, and that CIOs and CFOs may require documented controls before funding agent projects. For regulated institutions, the practical step is to treat agents as accountable identities with named owners.

What to do now

1. Inventory all AI agents in use, including unsanctioned ones running on employee credentials, and assign a named human owner to each.
2. Review agent permissions against least privilege and remove the ability to alter their own access or disable monitoring.
3. Move agents toward just-in-time access instead of standing administrative privileges.

4. Inventory and rotate API keys, OAuth tokens and machine secrets, and monitor their use.
5. Require third parties that connect agents to your systems to document agent identity, ownership and access scope.
6. Brief the board and risk committee on agent accountability and document the controls expected before funding new deployments.