



AI Agent Identity Security Readiness Cheatsheet

Enterprises today find themselves in a race to deploy AI agents across systems, cloud environments, and customer-facing applications. But deployment eagerness is not always matched by readiness — **91% of organizations** are already using AI agents, and **nearly 50% lack formal oversight**, creating new challenges for governance and control.

While these non-human identities (NHIs) enable powerful new capabilities, they introduce risk in equal parts. Without a strong identity framework, AI agents expand the attack surface and create security gaps that traditional controls were simply not designed to manage.

To govern AI agents safely and at scale, organizations need an **identity-first security foundation that secures the entire agent lifecycle, from initial pilot to full scale production.**

Use this AI readiness cheat sheet as a quick reference for the identity security practices that help you **maintain visibility, enforce control, and verify compliance at scale** across AI agent ecosystems.



Most organizations are already using AI agents, and nearly 50% lack formal oversight.



Core identity security practices for AI agent ecosystems



Govern all agents



Establish centralized visibility and control over AI agents across the organization

01.

Discover and track AI agents and NHIs

Continuously discover and track AI agents and associated service accounts operating across environments. This prevents unmanaged identities from operating outside established governance or monitoring controls.

02.

Establish a centralized agent registry

Register discovered agents, assign ownership, and permissions. A centralized registry improves accountability, auditability, and visibility across AI systems.

03.

Centralize the governance of AI agent identities

Manage provisioning, authentication, authorization, and deprovisioning through a centralized identity system. A unified platform eliminates fragmentation, improves policy consistency, and supports scalability.

04.

Continuously monitor AI agent activity in real time

Monitor authentication activity, API calls, and behavioral patterns associated with AI agents in real-time. Real-time visibility helps detect misuse, anomalies, or compromised credentials quickly.

05.

Log and audit AI agent actions across systems

Record and audit authentication events, access attempts and agent-initiated actions. Strong logging supports incident response, compliance validation, and governance enforcement.



Secure every agent



Deploy AI agents that act safely, predictably, and in line with organizational intent.

06.

Bind agent sessions to verified human identity

Before any agent action occurs establish a verified human identity so actions are attributable, accountable and trusted.

07.

Secure tokens, credentials, and secrets in a dedicated vault

Store long lived tokens, credentials and secrets in a secure, isolated vault rather than embedding them in agents. Issue short-lived tokens at runtime to reduce blast radius and prevent leakage.

08.

Enforce authorization checks and human-in-the-loop controls

Grant agents permission to act only when allowed, with human approval required for higher-risk actions.

09.

Apply least-privilege access to AI agents

Limit agents to authorized resources and authorized actions, evaluated at request time to keep behavior scoped and predictable.

Turn readiness into action

Ready to take control of your AI agent ecosystem?

Access the complete [AI Identity Security Compliance Checklist](#) to guide your next steps.