

Okta for AI Agents

Securely manage your AI agents from a single control plane

AI agents are reshaping the enterprise like the cloud and SaaS did — only faster, far less predictably, and without the identity infrastructure to support them. As agents spread across the enterprise, they're acting on behalf of your users, connecting to your most sensitive systems, and making decisions autonomously. They are often invisible to security and IT teams, have little to no clear owner, managed identity, access controls, and audit trail. 88% of organizations report confirmed or suspected AI agent security incidents. Only 22% have identities tied to their agents.¹

To securely manage AI agents at scale, enterprises must answer three key questions:

- Where are my agents?
- What can they connect to?
- What can they do?


If you can't answer all three today, your organization is exposed.



¹ [The State of AI Agent Security 2026, GraviTEE](#)

Take control of AI agents before they cause a breach

Okta for AI Agents gives your AI agents managed, first-class identities so you can discover, onboard, protect, and govern them from a single control plane. No matter where your agents run, across frameworks, clouds, or SaaS environments, you can scope their access to least privilege, review it over time, and shut it down when something goes wrong.



 Where are my agents?

Discover and onboard your AI agents

Before you can secure AI agents, you need to find them. Okta brings known and unknown agents into view across your entire environment.

- **Manage the identities of homegrown agents:** AI agent registry allows you to register custom-built agents in Universal Directory as first-class identities. Get a single, searchable source of truth to assign human ownership, reduce blind spots, and extend your existing identity controls to your agents.
- **Import agents from external sources:** AI agent import enables you to import agents from leading platforms, including Salesforce Agentforce, Amazon Bedrock AgentCore, and ServiceNow AI Platform. With prebuilt integrations, organizations can reduce manual effort and bring agents under governance faster.
- **Shadow AI agent discovery:** The Okta browser plugin detects shadow AI agents by identifying new OAuth consent grants in managed Chrome browsers and surfacing unmanaged agents in a centralized dashboard with enriched user context, enabling visibility into what agents exist and who connected them.



What can they connect to?

Protect AI agent connections

Once you know where your agents are, you need to map everything they can reach and enforce policies on those connections. Instead of hardcoded credentials and standing access, agents get scoped, short-lived tokens – only for what they need for as long as they need it.

That protection extends across the key resource connection types organizations rely on today:

- **Authorization servers** - Issue tokens with tightly defined scopes so agents receive only the API permissions required for a specific task, limiting overprivileged access.
- **Secrets** - Grant agents access to vaulted secrets on demand rather than embedding credentials in code, reducing secret sprawl and exposure risk.
- **Service accounts** - Govern service account access through a managed system, bringing control, visibility, and accountability to how agents use privileged machine identities.
- **Applications** - Control how agents connect to SaaS apps through managed consent flows, so humans can approve access, handle tokens securely, and maintain a clear record of authorized applications.

- **MCP servers** - Treat MCP servers as governed resources so you can secure and manage how agents connect to the tools and services exposed through them.
- **Other agents** - Agents can now securely call other agents, not just apps and resources, under admin-approved policy – with every handoff captured in a single audit trail.

Together, these capabilities enforce least privilege for AI agents in real time, without slowing down adoption.

Every agent tool call, governed by your access policies

Agent Gateway is a single Okta-secured endpoint that any AI agent can call through to reach enterprise tools, regardless of which vendor built the agent or which tool it needs to reach. Okta is the identity provider in the path: it verifies the agent, controls which tools it can use, holds credentials so the agent never touches them, and logs every call against a managed identity.





What can they do?

Govern AI agents across their lifecycle

Once agent access is in place, you need to review that access over time, and revoke it when needed. Okta for AI Agents enables organizations to govern agents with automated access reviews, helping ensure they have only the access they need, with a complete audit trail. And if an agent goes rogue, security teams can manually deactivate the agent.

- **Request access to AI agents:** User access requests for AI agents replaces ad hoc approvals with a structured governance workflow. Users request access from their dashboard; admins manage approvals, automate actions, and enforce time-bound permissions.
- **Certify user access to AI agents:** Bring AI agents into the same certification workflows used for other enterprise resources. Approve or revoke user-level access on a regular cadence, with full auditability and policy enforcement.
- **Certify resource access for AI Agents:** Certify every agent's connections, revoke what it no longer needs access to, and produce the audit trail compliance requires. Right-sizing agent access reduces stale connections agents no longer need, reducing risk exposure, while reducing costs.
- **Deactivate rogue agents:** Manual agent deactivation gives you an instant kill switch to prevent new token requests and future authorizations when an agent behaves unexpectedly.
- **Capture audit logs & telemetry:** Get access to a complete record of agent activity, including tool calls, access attempts, and authorization decisions. Stream telemetry to your SIEM to monitor behavior, support audits, and accelerate incident response.

Why Okta

As AI agents become a core part of your enterprise, they introduce new identity, access, and governance challenges that traditional approaches were not built to address.

Okta provides the foundation to discover, onboard, protect, and govern AI agents across your environments, providing visibility into where agents exist, control over what they can connect to, and oversight of what they can do.

This is how Okta secures AI.



Getting Started

Ready to learn more about Okta for AI Agents?

Visit okta.com/products/govern-ai-agent-identity/ or [talk to an expert today](#).

About Okta

Okta, Inc. is The World's Identity Company™. We secure Identity, so everyone is free to safely use any technology. Our customer and workforce solutions empower businesses and developers to use the power of Identity to drive security, efficiencies, and success — all while protecting their users, employees, and partners. Learn why the world's leading brands trust Okta for authentication, authorization, and more at okta.com.