Datasheet

September 2026

Okta for AI Agents: What You Need to Know

What capabilities does O4AA provide?



Introduction

AI agents now have access to customer data, source code, and the SaaS apps your business runs on. Adoption is moving fast: 92% of executives say autonomous agents are already in use at their organization.[1] Most of that access was never formally granted, reviewed, or logged. Most security teams cannot say how many agents they have, who owns them, or what they touched yesterday. That gap is the real risk, because you cannot govern access you cannot see.

Okta for AI Agents™ treats every agent as an identity, held to the same standards Okta® already applies to your people and applications. That gives security and AI teams one place to answer four questions:

1. Where are my agents?
2. What can they do?
3. What are they doing?
4. How do I respond?

Here's what Okta for AI Agents does at each step.

1. AI Agents at Work 2026: Securing the agentic enterprise

© 2026 Okta, Inc. and/or its affiliates.

Discover and Onboard Your AI Agents

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

Manage the Identities of Homegrown Agents (AI Agent Registry): 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): 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.

Note: Okta's capabilities to import and/or discover agents from Microsoft Azure AI Foundry and Microsoft Copilot Studio are currently available in beta.

Discover Shadow AI Agents and MCP Servers: Detect shadow AI agents by identifying new OAuth consent grants in managed Chrome and Microsoft Edge browsers, enriching them with user context, and surfacing unmanaged agents in a centralized view. Okta's integration with CrowdStrike Falcon EDR extends discovery to the endpoint, surfacing coding agents on employee laptops (e.g., Claude Desktop, Cursor, Claude Code, Windsurf), mapping their connections, and attributing each one to a named owner.

Next steps:

- Take a deeper look at why agent identity is different from human and app identity.
- Evaluate your AI agent security program to see where the gaps are.



© 2026 Okta, Inc. and/or its affiliates.

Govern AI Agents Across Their Lifecycle

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, helping ensure agents get scoped, short-lived tokens — only for what they need, for as long as they need it. That 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, bring 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 help secure and manage how agents connect to the tools and services exposed through them.

Other agents: Agents can securely call other agents, not just apps and resources, under admin-approved policy, with each handoff captured in a single audit trail.

Once agent access is in place, you need to review and certify that access over time, and revoke it when needed:

Request Access to AI Agents (User Access Requests for AI Agents): Replace 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 Access to AI Agents (User Access Certifications for AI Agents): Bring agents into the same certification workflows used for other enterprise resources. Review, approve, or revoke access on a regular cadence with auditability and policy enforcement.

Certify Resource Access for AI Agents (Resource Access Certifications for AI Agents): Certify your agent's connections, revoke what it no longer needs access to, and produce an audit trail. Right-sizing agent access reduces stale connections agents no longer need, reducing risk exposure, while reducing costs.

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

Next steps:

- Explore the fundamentals of AI agent lifecycle management.
- See AI agent lifecycle governance in action with an interactive demo..

© 2026 Okta, Inc. and/or its affiliates.

Control Agent Access at Runtime

Even with access scoped correctly, you need visibility into what agents actually do at runtime. Agent Gateway is a single Okta-secured endpoint that your 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 doesn't reach them, and logs each call against a managed identity — monitoring runtime access in real time to help prevent data leakage and defend against prompt injection.

Next steps:

- Learn how runtime governance works with Okta's Agent Gateway.
- See how scoping tool access can reduce token spend.

Act Fast When an Agent Behaves Unexpectedly

Detecting a problem is only useful if you can act on it.

Deactivate Rogue Agents (Agent Deactivation): A rapid kill switch to prevent new token requests and future authorizations the moment an agent deviates from its intended mission or accesses sensitive data it shouldn't, helping contain risk in real time before it escalates, with an audit trail for compliance.

Audit Logs & Telemetry: Capture a comprehensive 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.

Next steps:

- Discover the top security risks AI agents introduce and how to close them.
- Assess the readiness of your agentic security architecture with simulated attack scenarios.

© 2026 Okta, Inc. and/or its affiliates.

Conclusion

Securing AI agents doesn't stay finished. New agents show up every week, their access drifts, and the one that misbehaves won't wait for a ticket. Okta for AI Agents gives you a single identity layer to help you find your agents, give each one only the access it needs, see what it does, and manually deactivate them when something goes wrong. It works across the platforms your teams already build on, so security keeps up with AI adoption instead of slowing it down.

See it in action. Watch the on-demand technical demo to see how Okta finds hidden agents, assigns each one a human owner, replaces standing credentials with short-lived tokens, and deactivates an agent the moment it goes off-script.



About Okta

Okta 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.

Any mention of future products, features, functionalities, or certifications in this blog is for informational purposes only. These items are not commitments to deliver and should not be relied upon to make purchasing decisions.These materials are intended for general informational purposes only and are not intended to be legal, privacy, security, compliance, or business advice.

© 2026 Okta, Inc. and/or its affiliates.