DATA SHEET | FEDERAL

June 2026

Okta for AI Agents - Core

Govern AI agents inside the same regulated boundary you already trust Okta to run.

Where things stand

Public sector AI adoption is under a clear mandate to promote innovation and security together, and autonomous agents are moving quickly from read-only assistants toward high-impact, transactional roles across civilian agencies, defense, and the integrators who build for them. That pace exposes critical gaps in identity infrastructure. As agents take on consequential work, agencies lack the centralized governance required to manage non-human identities securely, and the friction of balancing innovation, compliance, and security create a paralysis; stalling mission velocity.

Underlying it all is one hard rule: running AI agents against regulated data outside a FedRAMP-authorized environment is more than just a compliance violation, more importantly, a significant security threat. The same principle extends to other regulated boundaries, including HIPAA-covered environments for agencies and institutions that handle protected health information.

The challenge across federal

US Federal Civilian Agencies. As autonomous agents move toward high-impact, transactional roles, agencies lack the centralized governance to manage non-human identities. Without it, manual compliance friction threatens efficiency and public-purpose missions of an AI-augmented government.

US Defense. The rollout of generative models and advanced LLMs to roughly 3 million Department of War personnel has created a governance-speed gap. Scaling this advantage requires moving beyond manual oversight to identity-driven authority, with the dynamic visibility and lifecycle governance needed to keep unmanaged risk from stalling mission velocity.

US Federal System Integrators. As architects of government AI, FSIs must ensure their AI Software Factories are secure by design. That requires verifiable agent identity and accountability so unmanaged systems do not become new attack vectors, plus a unified identity fabric for granular oversight across complex, cross-boundary agency environments.

100x

— AI agents are the fastest-growing identity type, capable of expanding an organization's attack surface by **100x**, not by doing more but by doing it faster. Yet, most currently have **no identity, no owner, and no consistent security controls.**

The solution

Okta for AI Agents - Core manages the full agent lifecycle inside the same regulated environments you already trust Okta to run for human identity - whether that be FedRAMP, HIPAA, or others. Same boundary. Same assurances. New coverage for a new identity class.

Agents are registered and managed in Universal Directory as first-class identities, each with a unique identity, an assigned human owner, and a complete audit trail. Access is scoped to least privilege, replacing long-lived static keys, and can be reviewed, certified, or revoked at any time. Because agents operate inside Okta's existing identity fabric, there is no parallel infrastructure to deploy and no separate compliance boundary to manage. Teams govern non-human identities using the same identity platform they already apply to people.

DISCOVER & ONBOARD

Where are my agents?

Your agents become a known, owned, first-class identity inside your environment, whether it came from a third-party platform or your own developers.

- A**I Agent Import.** Import known agents from leading supported third-party platforms using pre-built OIN integrations, reducing manual setup and accelerating onboarding.
- **AI agent registry.** Bring custom-built, homegrown agents into Universal Directory as first-class identities. Assign human owners, apply the identity controls you already use, and manage your agents from one source of truth.

PROTECT

What can they connect to?

Once agents are known and owned, you control exactly what they can reach, with scoped, short-lived credentials in place of long-lived static keys. Protection extends across a variety of resource types:

- **Authorization Servers**, with defined and enforced access for homegrown agents.
- **Applications**, connected through brokered consent for secure, user-authorized access, enforced at runtime.
- **MCP Servers**, connected directly as managed resources with admin-defined controls and runtime enforcement.

GOVERN

What can they do?

Agents enter the same governance lifecycle you already use for your workforce, so oversight is continuous.

- U**ser Access Requests.** Agents request access while admins manage approvals and enforce time-bound permissions.
- **User Access Certifications.** Agents enter the same certification workflows used for SaaS apps, with full auditability and automatic enforcement.
- **Agent Deactivation.** Manually deactivate the agent with a kill switch to prevent new token requests and future authorization when an agent deviates from its mission.
- **Audit logs and telemetry.** Log all agent activity and optionally stream it to your SIEM, fulfilling the rigorous reporting requirements of oversight bodies like the GAO.

Value-by-segment table

Segment	What Okta delivers
US Federal Civilian Agencies	The identity control plane to securely scale high-impact AI without sacrificing velocity, anchoring every AI-driven decision to a verified human owner and keeping every action authorized, auditable, and aligned with mandates.
US Defense	Proof of oversight without sacrificing mission tempo, linking agentic behavior to verified human intent in real time to reduce risk while operationalizing an AI-first force at the speed of the frontier.
US Federal System Integrators	Secure-by-design as a competitive advantage, anchoring agent actions to verified identities for granular oversight across multi-agency environments and protecting high-impact use cases and Pace-Setting Projects from unmanaged risk.

Use cases

- **Public Benefits Eligibility & Fraud Prevention (Civilian).** Anchors benefit-triage agents to a verified human owner and enforces control, mitigating identity hijacking and bulk data exfiltration.
- **Automated Correspondence & Records Processing (Civilian).** Replaces long-lived static keys with scoped, governed access, closing the backdoor into agency document repositories.
- P**redictive Logistics & Readiness (Defense).** Enforces least-privilege access and provides a kill switch to prevent a hijacked or over-scoped agent from sabotaging readiness
- **Secure-by-Design AI Software Factories (FSI).** Provides verifiable agent identity and audit evidence, giving integrators the compliance-by-design proof needed to win high-impact contracts.

About Okta

Okta, Inc. is The World's Identity Company™. We secure AI, machine, and human identity so everyone is free to safely use any technology. Our customer and workforce solutions empower businesses and developers to protect their AI agents, users, employees, and partners while driving security, efficiencies, and innovation. Learn why the world's leading brands trust Okta for authentication, authorization, and more at okta.com.