Securing AI agents at scale

Identity's role in the agentic enterprise

Okta, in association with Accenture

Executive brief

As the defining platform shift of our generation, AI has already upended the way work gets done across virtually every industry. In particular, AI agents have accelerated workflows and opened up new possibilities through their ability to act on behalf of users, expanding their overall capacity and generally transforming the workforce experience.

But in addition to opportunity, these agentic systems bring risk. The very autonomy that makes them so useful also opens them (and the organizations that deploy them) to the possibility of overreach, non-consented action, and data exposure. Making secure use of AI agents begins with identity — specifically, a modern approach to Identity and Access Management (IAM) that can facilitate agents' efficiency-boosting functionality without granting them broad and/or static access to sensitive information.

This resource covers the essentials of an identity-powered approach to AI agent security:

AI agents possess unique characteristics in contrast to human users and therefore introduce a new set of security risks. Organizations that allow AI innovation to outpace the modernization of their security posture are encouraging avoidable security vulnerabilities.

Modern identity security delivers a cornerstone of AI agent security. By adapting core authentication and authorization protocols to the specifics of agentic systems, modern IAM simultaneously secures and streamlines the effective functioning of AI agents.

Okta and Accenture can support organizations at any stage in their AI agent journey, helping them build a robust identity security fabric that extends all identities, use cases and resources. Built for secure growth throughout the next phases of AI innovation.

AI agents are here

And they're not waiting for instructions

In a few short years, AI transformation has, itself, transformed. In contrast to the simple request-response models of early AI tools, enterprises across industries are finding new uses for AI agents that leverage LLMs and other AI foundations to behave autonomously on behalf of users.

These agents can execute complicated tasks and collaborate with other systems (and other agents) with minimal human involvement, and their flexible applications are already driving transformation across the modern enterprise. For example, a system of AI agents can help employees aggregate and analyze data contained within a CRM or, in a very different vein, arrange the entirety of a business trip — from booking flights and hotels to managing scheduling issues — without step-by-step instructions.

As these examples illustrate, AI agents require a massive amount of data to work efficiently and effectively. Some of that data (e.g., credit card information, travel logistics, CRM data) is highly sensitive, which makes risk management a core priority of AI agent adoption.

Agent adoption is outpacing organizations' ability to secure them:

90% | **of organizations lack a comprehensive strategy for dealing with AI threats**

91% | **of organizations are building and deploying them anyway**

Accenture 2025 State of Cybersecurity Resilience Report

Unique identities. Unique protections.

Like human users, AI agents require authenticated access to sensitive information to work securely.

But unlike human users, AI agents have short lifecycles and authenticate using their own methods, which means they require a novel approach to identity security.

Dynamic and ephemeral lifecycles

AI agents are often spun up and down frequently, which means they need rapid provisioning and de-provisioning

Diverse authentication methods

API tokens, JSON web tokens, mutual TLS, and cryptographic certificates — all very different than human login flows

Granular, time-limited permissions

To minimize exposure, agents require very specific (often temporary) permissions that are closely tailored to their use cases

"If the secure, managed way of creating agents within your organization isn't the easiest way to do it, your teams are going to take a different path. And security gaps are going to emerge."

– Greg Callegari,
Managing Director, Identity Security, Accenture

The risk landscape is changing

Without the right identity controls, agents create and exacerbate risk

Agent manipulation

- **New vulnerabilities:** Bad actors see AI agents as a new entry point. Without the right protections, a simple prompt injection attack can put sensitive information in the wrong hands.

Agent access and behavior

- **Standing privileges:** Agents have ephemeral and dynamic lifecycles, which makes traditional provisioning frameworks obsolete. Agents with static entitlements risk drawing on (and sharing) sensitive data that may not be necessary for their present task.

Regulatory risk

- **Lack of accountability:** Agents often lack traceable ownership and consistent logging, making it difficult to understand or justify their actions. Without consistent, documented approvals of AI agent behavior, organizations risk noncompliance with standards like GDPR, SOX, and DORA.

- **Credential management:** To secure the many credentials that AI agents use for authorization and access, these tokens etc. must be vaulted and rotated.

- **Decisioning and consent:** Agents' value lies in their ability to act on behalf of users. But this independence must include human-in-the-loop consenting protocols to prevent unwanted action.

- **Lack of visibility:** A disjointed and sprawling web of AI agents increases the risk of over-privileged access and long-lived credentials, which heightens the risk of financially and reputationally damaging data exposure.

The rapid rollout of AI tools is creating a warped sense of AI security readiness:

Only **36%** of tech leaders admit generative AI's rapid rollout is outpacing their ability to integrate security measures

Accenture State of Cybersecurity Survey, 2025

Yet **90%** of those same leaders lack the security maturity needed to combat modern threats

Accenture Research analysis based on State of Cybersecurity Data, 2025

"In many organizations, AI is moving faster than security. Most organizations don't even realize how exposed they are."

– Greg Callegari, Managing Director, Identity Security, Accenture

okta | accenture

The identity mandate

Strengthening (and scaling) AI agent security begins with modern identity management

A modern, unified identity and access management (IAM) system is the foundation of strong, enterprise-grade security. This is as true in the age of AI as it has been for years: Strong security begins with a robust means of authenticating who users are and managing what they can access based on granular, context-sensitive, least-privilege permissions.

While they may supply basic protections sufficient for conventional software, legacy identity controls simply cannot apply this level of identity management to the security needs of AI agents. These legacy controls must be modernized to account for the dynamic nature of AI agents, helping to ensure the security of their behaviors and maximizing their contributions to core business functions.

Identity plays a critical role in Accenture's AI Agent Zero Trust Model for Cyber Resilience:

What good looks like

Build autonomous AI agents from the ground up with security as core.

A proactive approach that integrates security as a **fundamental requirement** throughout the entire AI lifecycle, enabling our clients to adopt AI technologies with **confidence** and **resilience** in the face of an ever-evolving landscape, and to **scale with assurance.**

>

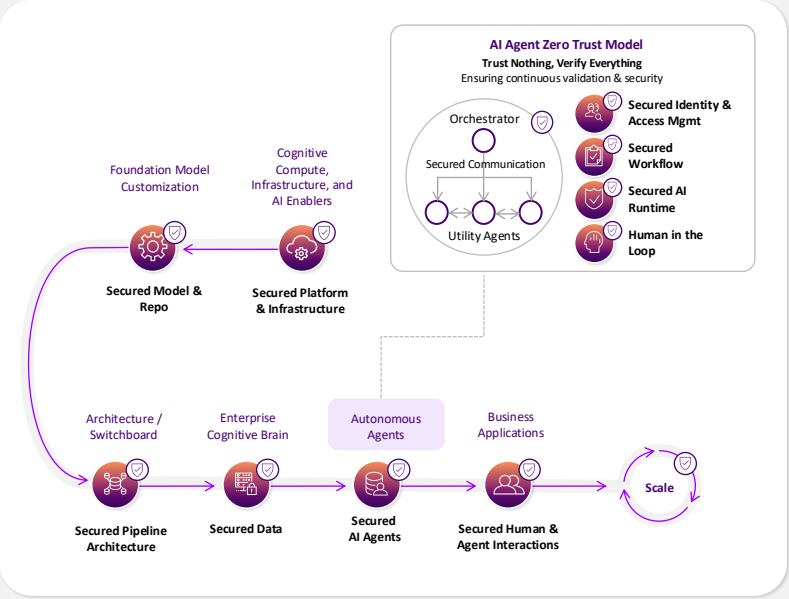

Source: www.accenture.com/us-en/blogs/security/strengthening-ai-agent-security-identity-management

"Agents need their own identity, and that identity must be managed and clearly defined. Everything else flows from that: access, governance, threat detection, everything."

– Harish Peri, SVP + GM, Okta for AI Agents, Okta

Okta secures your organization's identity foundation

For this generation of agentic systems, and the next

We are witnessing the first major wave of AI agent innovation. It will not be the last. As AI agents become even more interconnected, and as the tasks they're able to manage become more complex, the importance of securing these innovations will grow as well. Managing the identity of each agent within these evolving ecosystems will be the key to maintaining security, efficiency, and efficacy at scale.

Okta secures AI agent identities, equipping enterprises with the modern identity infrastructure they need to build agents securely and govern them as first-class identities. By building agentic systems into the fabric of your organization's IAM, Okta helps secure and facilitate the adoption of agentic tools that fit your needs — and your next step.

Okta delivers better visibility into (and control over) AI agents, at scale:

Build

Empower developers to build AI agents that are secure by design with intuitive tools and built-in protections

Connect

Enable agents to securely connect to systems and applications without granting unlimited standing access

Manage

Establish full lifecycle management for agents: visibility, access control, governance, and remediation

Detect

Minimize the attack surface with real-time detection and remediation of AI identity configuration risk

> "It's time for identity to move from the perimeter to the center of AI operations. That means treating agents as first-class identities inyour security model."
>
> **– Harish Peri, SVP + GM, Okta for AI Agents, Okta**

okta | accenture

A closer look

Okta products make strong identity security your baseline

Build

Auth0 for AI Agents

- Easily implement secure login experiences that allow AI agents to log in on behalf of users
- Use secure standards to retrieve, store, and automatically refresh API tokens
- Enable autonomous agents to work independently while maintaining user control by getting explicit user approval approval for critical actions
- Allow granular permissions to help ensure agents only access authorized content

Connect

Cross-App Access (XAA) for AI Agents

- Allow secure AI agent and app connections with an open, standards-based protocol, offered out of the box for Auth0 customers
- Enable enterprise customers using a participating IdP to standardize app and agent access with control unified at the identity layer
- XAA builds on trusted standards (OAuth, JWT assertion grants) and extends them with new innovations (like JAG tokens) to preserve strong user and app context during machine-to-machine interactions

Manage

Agent Identity Lifecycle

- Centralize every aspect of AI lifecycle management: visibility, access control, governance, and remediation
- Detect and remediate AI identity configuration risk to minimize attack surface
- Vault and protect credentials for human-in-the-loop interactions

Detect

Visibility and Governance

- Detect and remediate risky AI identity configurations to minimize attack surface and blast radius
- Quickly deprovision agents using centralized logging, governance, and policy controls in order to contain impact
- Leverage clear audit trails for improved transparency and compliance

Protecting the path forward

Building trust into the DNA of AI adoption

Trust and autonomy are intricately linked. Without a deeply rooted sense of trust in increasingly autonomous agentic systems, employee and customer adoption will be hindered by skepticism — and businesses will struggle to wring maximum value out of the revolutionary efficiencies that AI agents enable. The threat of high-profile security events lingers over every AI adoption journey, which is why it's so critical to secure AI innovations at the outset.

The time to act is now. By securing their agentic systems with modern, identity-based protections, organizations can strengthen their competitive edge now and position themselves for resilience and speed as the AI revolution continues.

This is just the beginning

Securing AI agents now prepares organizations for what's next

Protecting innovation and GTM timelines

Addressing security gaps not only mitigates risk but also enables the secure scaling of AI agent innovation, which prevents slow rollouts and stalled product initiatives

Laying the foundation for a larger agent ecosystem

Establishing a robust human/non-human identity governance strategy prepares organizations for more complicated, interconnected agentic ecosystems

Preparing for a new regulatory landscape

Nailing down the accountability, traceability, and explainability of AI agent behavior prepares organizations for the inevitable intervention of regulatory bodies

Next steps

Mapping your agentic journey with Okta and Accenture

AI agents are already transforming the ways businesses connect with their employees, develop their products, and cater to their customers' changing preferences. The opportunities represented in this agentic sea change come with a wave of new risks and challenges. But an AI-ready security apparatus — complete with modern IAM that can govern and protect agentic identities throughout their lifecycles — lends organizations an unshakable foundation capable of supporting their next big step.

Whether your AI adoption is just beginning or well underway, Okta, in association with Accenture, can help you tailor your organization's identity management to its unique, future-ready needs.

An actionable roadmap to tangible results:

Co-create your AI identity journey

Uncover the opportunities and potential roadblocks of AI agent adoption through generative, personalized discussions

Visionary workshops

Uncover the opportunities and potential roadblocks of AI agent adoption through generative, personalized discussions

Proof of concept

Apply workshop learnings to specific use cases and technology solutions to gauge initial impact

Use case analysis

Deep dive on selected high-impact agentic workflows from a security perspective

End-to-end deployment

Deploy the strategies developed and refined in the previous two steps in a repeatable, scalable pattern

30/60/90 Days plan

Develop a phased plan to start your AI identity journey with concrete steps

Ready to secure and strengthen your AI agents?

Learn more at Okta.com

Resources:

Solution Brief: okta.com/resources/whitepaper-identity-security-agentic-ai