How to control AI agents and other non-human identities

With modern, identity-first security

Non-human identities drive better business. They also invite higher risk.

Non-human identities (NHIs) like machine-, service-, and AI-agents have quickly become commonplace across a number of industries, and for good reason. They accelerate companies' core capacities for collaboration, innovation, and productivity.

But insofar as NHIs prepare businesses for an AI-enabled future, they also create another attack vector for bad actors to exploit. AI agents, in particular, create a unique vulnerability for organizations looking to carve out a competitive advantage using the power of AI.

Without a modern, unified approach to identity-first security, organizations leveraging NHIs risk weakening their overall security posture by allowing their attack surface to balloon out of control.

In some enterprises,

NHIs now outnumber human identities **50 to 1**

Forbes

46% of organizations have experienced compromised NHI accounts or credentials in the past year (and another **26%** suspect they have)

TechTarget ESG

A new generation of risk

The use of NHIs has risen in tandem with the increasing use of cloud services, AI and automation, and digital workflows. Enterprise systems need to interact securely and efficiently without constant human oversight, and NHIs, like service accounts and AI agents, make this possible by allowing apps to authenticate to one another.

But these NHI authentications are just as vulnerable as authentications initiated by humans — sometimes even more so. If the secrets, keys, and/or tokens that NHIs use to authenticate wind up in the wrong hands, adversaries can gain deep and wide-ranging access to sensitive applications and data.

Especially when these NHIs make use of generative AI (as AI agents like chatbots and digital assistants often do), they are vulnerable to an array of new threats:

Prompt injection and data poisoning

Attackers may manipulate prompts or inject malicious content to trick GenAI tools into generating inappropriate or inaccurate information.

Shadow access and lateral movement

Bad actors can use GenAI tools as a gateway for pulling sensitive content from connected systems (e.g., Salesforce, Jira, internal wikis).

Data overexposure

GenAI tools require access to a vast array of sensitive data to function. If identity is mismanaged, users could access information far beyond the appropriate level of access for their role.

51% of companies have deployed AI agents

Pager Duty

Only **15%** of security teams feel confident in preventing NHI-related breaches

Cloud Security Alliance

The unique risk of AI agents

Nowhere is the potential risk and reward of NHIs more apparent than in the case of ai agents. By acting autonomously on behalf of people and organizations, AI agents enable levels of operational efficiency and personalized, automated customer service that were previously unimaginable.

But this power is a double-edged sword. AI agents depend on data, resources, and feedback in order to continually improve, all of which depend on (authorized, authenticated) access. This extensive access is a golden goose for bad actors looking to carry out identity-based attacks on the NHIs behind them.

Within many organizations, non-human and machine identities lack sufficient monitoring, if they are monitored at all. Too frequently, NHIs are over-permissioned, never rotated, or still active long after their purpose ends, creating a critical vulnerability for bad actors to exploit. In it and security environments where identity functions are scattered over different systems and applications, these vulnerabilities tend to slip through the cracks and go unnoticed — until it's too late.

Identity is the vulnerability.
Identity security is the solution.

Bottom line:

The best defense against NHI-related vulnerabilities begins with eliminating the fragmented identity systems that make gaps in visibility and enforcement possible. By unifying identity systems on a single platform, organizations can gain better control over their NHIs while also driving better administrative efficiency. Modern identity platforms help you achieve this unified approach to security.

Okta makes it possible

The Okta platform enables a robust and vastly simplified approach to managing non-human and machine identities. By unifying the management of your organization's identities in one centralized platform, Okta eliminates blind spots and gives you a comprehensive view into where your NHIs exist and what they can access.

Identity Security Posture Management

- Gain real-time visibility into threats across all systems, devices, and user types, ensuring a proactive security posture
- Leverage third-party signals alongside first-party data from Okta for deeper insights and faster threat detection
- Quickly mitigate threats with customizable automated actions, such as triggering MFA or logging out compromised users

Privileged Access Management

- Enable passwordless, phishing-resistant authentication for a seamless and secure user experience
- Verify device security posture during authentication to enforce compliance
- Alert users and admins of phishing attempts and log attacks for greater visibility
- Block untrustworthy apps before they can exploit authentication processes

Secure Identity Integrations

Advanced security integrations with SaaS applications prevent overlong NHI access and protect NHIs across your ecosystem.

- **Lifecycle and Entitlement Management:** Automate identity provisioning and deprovisioning for NHIs, ensuring just-in-time access
- **Unified Single Sign-On (SSO) & Policy Enforcement:** Extend SSO and security policies to service accounts and machine identities
- **Workflow Automation and Session Termination:** Prevent orphaned NHIs by enforcing automated offboarding and session revocation

Okta + Amazon: the key to securing GenAI

By delivering GenAI-powered assistance to anyone building with AWS, Amazon Q is transforming how software developers, business intelligence analysts, contact center employees, and other core teams get work done.

Now, by combining the efficiency, productivity, and CX benefits of Amazon Q with Okta's suite of AI-ready security tools, organizations can securely weave powerful GenAI into the fabric of their core operations without exposing themselves to new sources of risk.

Securely Implement GenAI

Protect data stored in Amazon Q from unauthorized access with Okta's enterprise-grade identity security, all while streamlining identity management with Okta's automation-powered lifecycle management tools.

Supercharge Productivity

Empower teams to be more creative and productive with Amazon Q's assistance. Generate meeting summaries, create content, and complete tasks — without exposing data to new risks.

Provide Seamless Access

Enable streamlined, secure access with role-based access to specific data sets. Get a better handle on costs by granting subscriptions on request or through automation-based criteria.

Unlock Amazon Q Business with Okta

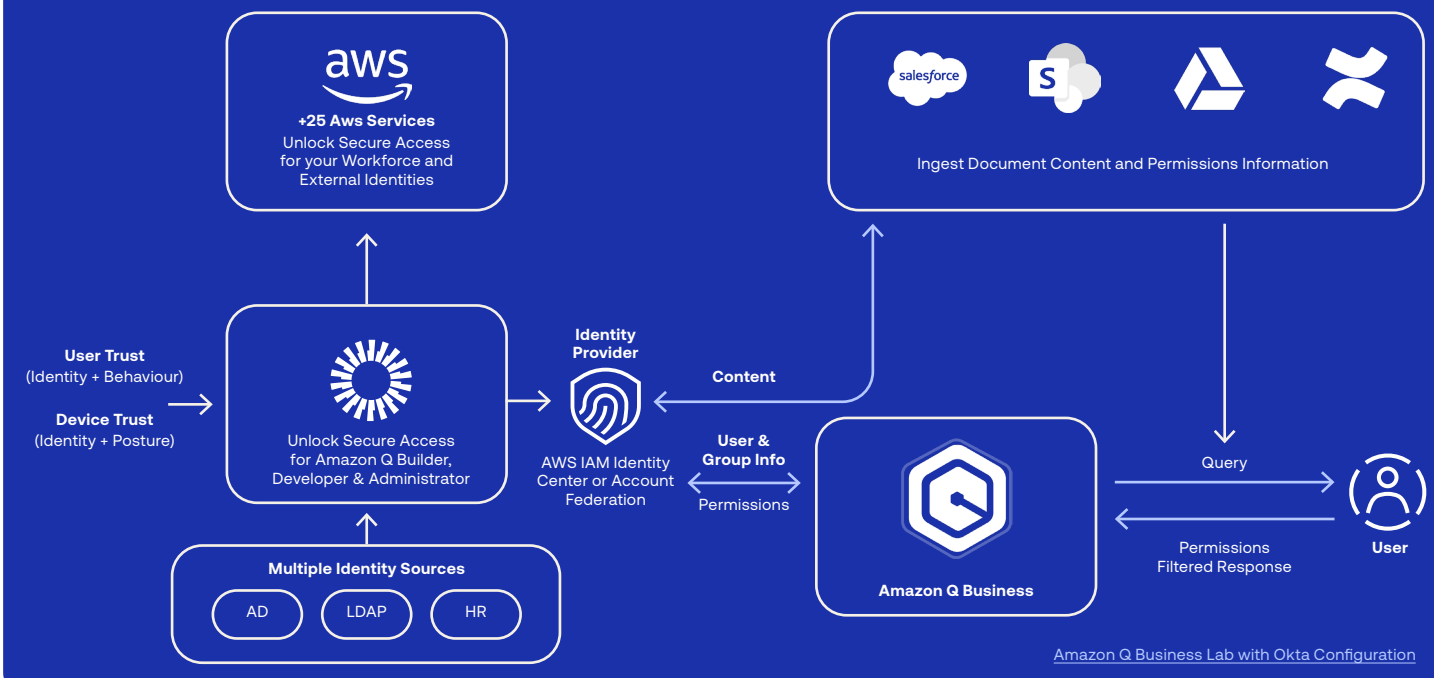

Amazon Q Business Lab with Okta Configuration

Put your non-human identities to work — securely

In record time, AI has gone from a hypothetical future advantage to a present necessity. Making full use of modern tools like AI agents is table stakes for organizations looking to remain competitive. But the effective use of NHIs hinges on your organization's ability to secure them.

The surest path to secure, AI-powered workflows runs through a unified identity platform that eliminates risky fragmentation and strengthens the foundation of your security ecosystem. Okta makes that possible.

That's why AWS and Okta are working together to provide a secure foundation for the next generation of intelligent automation. Whether you're deploying AI agents in AWS or leveraging Amazon Q across teams, the Okta Identity Platform ensures those identities remain secure, governed, and auditable.

Ready to learn more about unifying your security strategy with modern identity solutions? Reach out to our team and see the Okta Platform in action.