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.

The use of non-human identities (NHIs) for machine-, service-, and AI-agent communications has quickly grown across a number of industries, and for good reason. They accelerate companies' core capacities for collaboration, innovation, and productivity.

But while 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 as they behave differently from human users and traditional software, often requiring access to highly sensitive data that makes them even more attractive to bad actors yet hard for organizations to secure them.

Without a unified identity-first security approach, organizations leveraging NHIs are dramatically increasing their risk by not securing the quickly expanding attack surface.

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 increasing use of cloud services and cloud platforms like AWS to expand for use for GenAI tools, AI and automation, and digital workflows. Enterprise systems need to interact securely and efficiently without constant human oversight, and NHIs like service accounts used by AI agents make this possible by allowing applications and services to authenticate with one another, enabling automated tasks, data exchange, and complex operational processes that would otherwise be impossible or incredibly time-consuming for humans to manage.

Because unmanaged NHIs are so vulnerable, organizations need to consider strategies that mitigate this new risk. 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 non-human identities (NHIs) are used with generative AI (GenAI) tools like chatbots and digital assistants, a new array of threats can emerge due to the GenAI tool's inability to provide user context:

Prompt injection and data poisoning

Attackers may manipulate prompts or inject malicious content, causing GenAI tools to generate inappropriate or inaccurate information from their data sources.

Shadow access and lateral movement

Bad actors can exploit GenAI tools to access sensitive content from connected systems (e.g., Salesforce, Jira, internal wikis) that the GenAI tool is querying.

Data overexposure

GenAI tools often access vast amounts of sensitive data. If user context isn't properly conveyed to the underlying data sources, users could access information far beyond what is appropriate 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.

This is especially critical when developing GenAI solutions on Amazon Bedrock, where rapid deployment must be matched with secure access control. Okta acts as the identity control plane for your Bedrock environment—ensuring only authorized users and AI agents can access sensitive models and services.

Okta makes it possible

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

Identity Security Posture Management

- Provides continuous monitoring and risk analysis of NHIs
- Automatically detects local non-federated service accounts and flags risky NHIs
- Detects MFA gaps and credential misuse in machine identities
- Provides framework compliance mapping and guided remediation workflows for NHI-related security issues

Okta Privileged Access

- Securely manages service account passwords and enforces policies for who can access what for how long
- Automatically rotates secrets, ensuring no long-lived credentials remain exposed
- Audits and tracks who or what checked out the account

Secure Identity Integrations

Advanced security integrations with SaaS applications that 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:** Helps 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 weave powerful GenAI into the fabric of their core operations without limiting their exposure 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

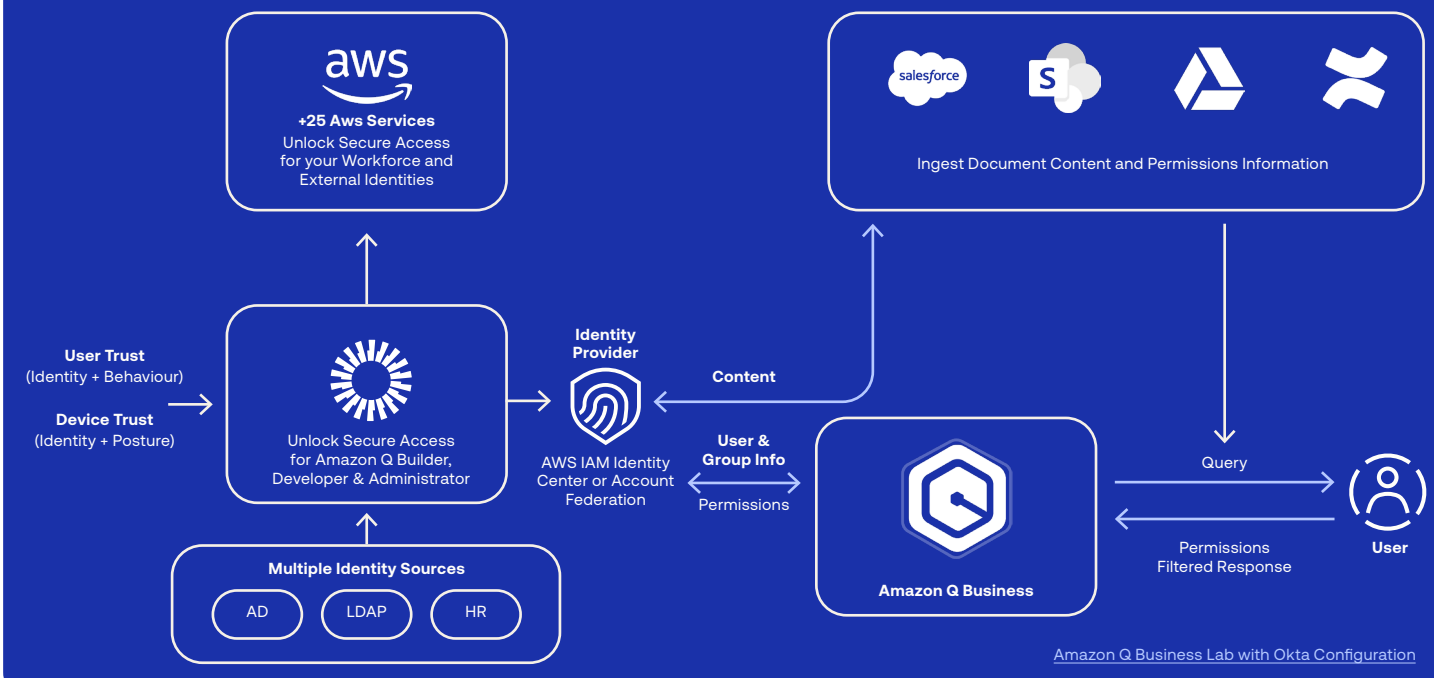

Amazon Q Business Lab with Okta Configuration

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.

With Okta, AI-powered workflows run through a unified identity platform that eliminates risky fragmentation and strengthens the foundation of your security ecosystem — all made possible through Okta's identity security fabric.

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 helps to ensure those identities remain secure, governed, and auditable.

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