

Powered by Okta | Delivered by BeyondID

AI Agents - copilots, autonomous workflows, RPA bots, and API-driven services are going from ideation to production in record pace using more AI driven automation and less human intervention than ever before.

Meanwhile, viibe coded applications built by technical non-developers are making their way into production.

Agentic AI (MCP Clients and Servers) are acting on behalf of a user, with high privilege, autonomy, and speed with or without the user present. Yet most organizations lack visibility, governance, and security controls over them.

Traditional IAM and security models were not designed for AI. The need to ensure that user context is present from agent to the resources requested is paramount. Without identity-first controls, AI adoption introduces significant security, compliance, and operational risk.

Okta's Priority: Securing AI

Okta's #1 strategic priority is Securing AI by extending identity as the control plane beyond humans to non-human and AI identities.


Identity becomes the foundation for authentication, authorization, policy enforcement, and governance across all AI interactions.



Why Okta + BeyondID

Okta defines how AI should be secured → identity as the control plane

BeyondID ensures AI is secured in production → strategy, architecture, and execution



Our Solution

BeyondID's AI Agent Identity & Security Blueprint operationalizes Okta Secures AI by helping organizations discover, classify, and govern AI agents using an identity-first, Zero Trust approach, built directly on Okta.

This is not theoretical advisory. It is an implementation-ready blueprint that enables secure AI adoption at enterprise scale.

What the Blueprint Delivers

- **AI Agent Discovery & Inventory:** Identify AI agents across the environment, including copilots, autonomous agents, RPA workflows, and third-party AI integrations.
- **Risk & Privilege Classification:** Classify AI agents by autonomy, access level, and blast radius to establish a security baseline.
- **Identity-First AI Security Architecture:** Establish Okta as the central identity and authorization authority for AI agents and non-human identities.
- **Zero Trust Enforcement Model:** Define secure authentication and authorization patterns for: Agent-to-LLM, Agent-to-API and Agent-to-agent interactions
- **MCP-Centric Control Layer:** Design an MCP-based enforcement and traXic control architecture for all LLM interactions.
- **Governance & Operating Model:** Define ownership, lifecycle management, policy enforcement, and incident response for AI agents.
- **AI Security Adoption Roadmap:** A prioritized, board-ready roadmap to guide secure AI implementation at scale.

Business Outcomes

- Secure AI adoption without slowing innovation
- Reduced AI security and compliance risk
- Clear visibility into AI agent behavior and privilege
- Elimination of future AI security debt
- Board-ready insights and actionable next steps

How We Deliver

- Up to 12 structured working sessions with security, identity, and AI stakeholders
- Duration: 6–8 weeks
- Delivered by BeyondID's Okta-certified identity and security experts