auth0

AI and identity readiness checklist

How ready is your retail identity stack for AI-driven agents, automation, and machine-to-machine interactions?

AI agents are already acting across your retail stack. Most teams don't have a clear way to control what they can access or do.

Use this checklist to assess whether your identity layer can securely support AI systems acting on behalf of customers, employees, and internal services.

Check each item that applies to your current setup. If you're unsure, treat it as a gap.

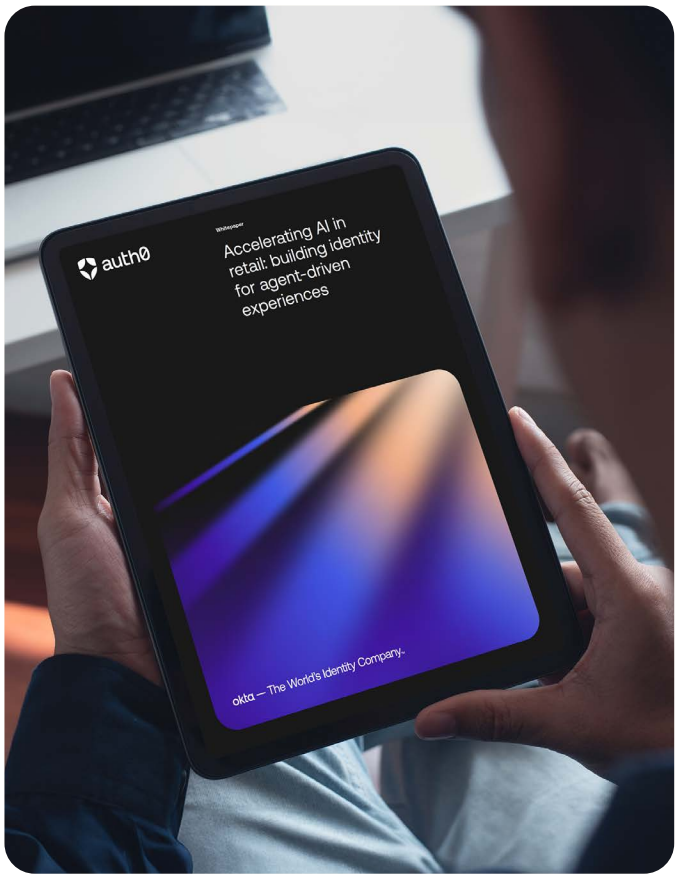

1. Agent identity and ownership

- ☐ Do you have visibility into all AI agents, copilots, and automated services in your environment?
- ☐ Is each agent tied to a specific user, system, or business function?
- ☐ Can you clearly attribute every agent action to a source of intent (user, workflow, or system)?
- ☐ Do you maintain an inventory or registry of active non-human identities?

2. Authentication and delegation

- ☐ Do AI agents authenticate using modern standards like OAuth 2.0 or OIDC instead of shared credentials?
- ☐ Are agents acting on behalf of users using secure token exchange or delegated identity flows?
- ☐ Is user context preserved when an agent performs actions across systems?
- ☐ Can you revoke or re-authenticate agents without disrupting the entire workflow?

3. Token and credential management

- ☐ Are API tokens short-lived, scoped, and automatically rotated?
- ☐ Are secrets removed from code, logs, and pipelines?
- ☐ Do agents avoid using hardcoded API keys or static credentials?
- ☐ Can you centrally manage and audit all tokens used by AI systems?

4. Authorization and policy enforcement

- ☐ Do you enforce least-privilege access for every AI agent?
- ☐ Are permissions evaluated at request time, not hardcoded in application logic?
- ☐ Can you define what actions an agent can take based on context (user, transaction, risk)?
- ☐ Are high-risk actions (e.g., refunds, account changes) gated by additional approval or policy checks?

5. Data access and RAG controls

- ☐ Do you enforce fine-grained access controls on data retrieved by AI systems?
- ☐ Can you restrict which documents, records, or fields an agent can access?
- ☐ Are AI responses limited to data the user is authorized to see?
- ☐ Do you prevent sensitive data exposure in prompts, logs, or model outputs?

6. Visibility and monitoring

- ☐ Can you track what each AI agent is doing across systems and APIs?
- ☐ Do you log which identity, token, or policy enabled each action?
- ☐ Can you detect unusual or out-of-pattern agent behavior in real time?
- ☐ Do you have a way to quickly disable or contain a misbehaving agent?

7. Integration and control plane

- ☐ Is your identity layer centralized across apps, APIs, and AI systems?
- ☐ Can you apply consistent authentication and authorization policies across all agents?
- ☐ Does your identity system integrate with AI platforms, APIs, and data sources without custom workarounds?
- ☐ Can developers extend identity logic (policies, hooks, workflows) without rewriting core systems?

Score:

- **24–28 boxes checked:** Your identity layer is well-positioned to support AI agents securely and at scale.
- **16–23 boxes:** You have a solid foundation, but gaps may appear as AI adoption grows.
- **Below 16 boxes:** Your identity layer may not be ready to support AI systems safely in production.

Want to dig deeper?

Explore how Auth0 helps teams secure AI agents with authentication, token management, and fine-grained authorization.