imprivata®

Agentic AI access readiness assessment

Real-world risks and security self-assessment

Agentic AI can drive meaningful productivity gains, but it also creates a new class of access risk: autonomous or semi-autonomous agents that can connect to systems, use tools, access sensitive data, and take action. This assessment helps security leaders test their understanding of agentic AI risk and assess how ready their organization is to govern AI agents as identities.

PART 1: REAL-WORLD RISKS

Choose the best answer for each question.

1: What makes AI agents materially riskier than traditional AI chatbots?

- ◯ **A.** They generate longer responses
- ◯ **B.** They can autonomously plan and execute actions across connected systems
- ◯ **C.** They require more expensive infrastructure
- ◯ **D.** They are harder to train

2: What is the biggest access management risk when AI agents are deployed quickly?

- ◯ **A.** Agents may use unmanaged, shared, or embedded credentials
- ◯ **B.** Agents may consume too much storage
- ◯ **C.** Agents may require more user training
- ◯ **D.** Agents may slow down service desk workflows

3: Which control is most important before giving an AI agent access to enterprise systems?

- ◯ **A.** A naming convention
- ◯ **B.** A chatbot usage policy
- ◯ **C.** A governed identity with least privilege, auditable, revocable access
- ◯ **D.** A larger model context window

ANSWER KEY FOR PART 1:

1) b - Agentic systems can plan, use tools, and take action across systems, which expands the risk beyond inaccurate or unsafe outputs.

2) a - Shared, embedded, or long-lived credentials can turn a useful agent into an unmanaged access path.

3) c - AI agents should be governed as identities with controlled, auditable, and revocable access before they connect to enterprise systems.

PART 2: SECURITY SELF-ASSESSMENT

For each question, choose the answer that best describes your current state. Score each response from 1 to 4, where 1 is least mature and 4 is most mature.

4: Which best describes your visibility into AI agents operating in your environment?

- ◯ We do not have a reliable inventory of AI agents ***(1 point)***
- ◯ We know about some approved agents, but not shadow or experimental agents ***(2 points)***
- ◯ We maintain an inventory of approved agents and their business owners ***(3 points)***
- ◯ We maintain a continuously updated registry of agents, owners, access rights, risk level, and lifecycle status ***(4 points)***

5: Which best describes how your organization governs AI agent identities and permissions?

- ◯ Agents use shared or embedded credentials with broad access ***(1 point)***
- ◯ Agents use separate credentials, but access is managed manually ***(2 points)***
- ◯ Agents are assigned governed identities with role-based access and periodic review ***(3 points)***
- ◯ Agents are governed as identities with least privilege, approvals, reviews, and automated lifecycle controls ***(4 points)***

6: How does your organization prevent AI agents from exposing or misusing credentials?

- ◯ Agents may use embedded, shared, or long-lived credentials ***(1 point)***
- ◯ Some credentials are vaulted, but agents may still receive or store secrets ***(2 points)***
- ◯ Agents use controlled access methods with limited credential exposure ***(3 points)***
- ◯ Agents receive brokered, time-bound access without direct access to long-lived credentials ***(4 points)***

7: Which best describes how your organization controls what AI agents can do once deployed?

- ◯ Agents can act freely within the applications they can reach ***(1 point)***
- ◯ Some restrictions exist, but controls vary by use case ***(2 points)***
- ◯ Agent actions are constrained by policy, workflow approvals, and defined task boundaries ***(3 points)***
- ◯ Agent actions are continuously policy-enforced with least privilege, context-aware controls, and human approval for sensitive operations ***(4 points)***

8: How are AI agents governed when accessing legacy, on-premises, or non-agent-ready systems?

- ◯ We do not have a consistent approach ***(1 point)***
- ◯ Access is handled case by case by application or infrastructure teams ***(2 points)***
- ◯ Access is restricted through standard privileged access controls where possible ***(3 points)***
- ◯ Access is brokered, monitored, and revocable across both modern and legacy systems ***(4 points)***

9: Which best describes your organization's response when an AI agent behaves unexpectedly or violates policy?

- ◯ There is no defined process for AI agent misuse or failure ***(1 point)***
- ◯ Issues are handled informally after they are discovered ***(2 points)***
- ◯ Agent incidents follow an established review and escalation process ***(3 points)***
- ◯ Agent incidents trigger automated containment, investigation workflows, policy updates, and post-incident governance review ***(4 points)***

Scoring and results

Add your scores for questions 4-9. Maximum score: 24.

6–10 POINTS:

HIGH RISK / UNMANAGED AGENT ACCESS

AI agents may be moving faster than identity, access, and incident controls can support.

11–15 POINTS:

EMERGING CONTROLS

Some controls exist, but gaps likely remain around inventory, credentials, auditability, and revocation.

16–20 POINTS:

GOVERNED AGENT ACCESS

You have a strong foundation, but you should validate coverage across high-risk, legacy, and non-agent-ready systems.

21–24 POINTS:

AGENTIC-READY ACCESS GOVERANCE

You have mature controls for agent identity, least privilege, monitoring, lifecycle management, and incident response.

AI agents are becoming a new class of privileged identity. Click here to learn how Imprivata Agentic Identity Management helps organizations give AI agents controlled, auditable, revocable access to enterprise systems — without embedding long-lived credentials in the agent.