imprivata

Governing AI Agents at Scale

A readiness checklist for IT and security leaders

Agentic AI is moving from experimentation into live healthcare workflows. Unlike traditional AI tools that primarily provide information or recommendations, AI agents can plan tasks, make decisions, access systems and data, and act with varying degrees of autonomy.

Research indicates that healthcare organizations are adopting these capabilities quickly...but a trust gap exists:

86% are confident they have or will soon have visibility into AI agent activity

44% are piloting or conducting proof-of-concept projects

28% have agentic AI in production and use

21% expect to implement agentic AI within the next 12 months

And yet...

72% report that AI tools or agents are deployed without formal IT approval at least occasionally

The following cross-functional checklist can help healthcare organizations identify governance issues that should be addressed before expanding agentic AI across clinical and operational workflows.

1. Discover and assign ownership

- [ ] Maintain an inventory of all AI agents
- [ ] Establish a process for detecting shadow or unapproved AI agents
- [ ] Document each agent's purpose, owner, sponsor, and deployment status
- [ ] Identify connected systems, applications, data, and downstream tools or agents
- [ ] Assign clear business and technical accountability for every agent

2. Establish a governable identity

- [ ] Assign every AI agent a unique, governable identity that supports independent attribution and auditability
- [ ] Document whose authority each agent is acting under
- [ ] Define how delegated authority should be granted, time-limited, and revoked
- [ ] Establish processes for provisioning, modifying, suspending, and decommissioning AI agents

3. Control access and authority

- [ ] Apply least-privilege access based on the AI agent's specific role
- [ ] Limit access according to the agent's workflow risk, autonomy level, and data sensitivity
- [ ] Define which systems, records, tools, and APIs each agent is authorized to use
- [ ] Separate an agent's permission to access systems or data from its authority to take action
- [ ] Use time-bound and context-aware permissions for higher-risk activity

4. Set clinical and operational guardrails

- Classify each AI agent's use case according to clinical, operational, privacy, security, compliance, and patient-safety risk
- Define which actions an agent may perform independently and which require human review, approval, or step-up authentication
- Establish escalation procedures for unexpected, unsafe, or unauthorized agent behavior
- Test agents before expanding autonomy or moving them into higher-risk workflows

5. Monitor, audit, and respond

- Continuously monitor AI agent access and behavior
- Alert on excessive permissions, anomalous behavior, and policy violations
- Maintain an audit trail showing what each agent did, when it acted, and under whose authority
- Ensure the organization can swiftly constrain or revoke an agent's access when necessary
- Incorporate AI agent incidents into existing response processes

6. Govern throughout the lifecycle

- Establish a cross-functional AI governance body
- Review agent ownership, access, risk classification, and activity regularly
- Reassess controls whenever workflows, integrations, or autonomy change
- Retire inactive agents and promptly revoke associated credentials, permissions, and delegated authority
- Report adoption, exceptions, and material risks to senior leadership

Six key governance readiness questions

Leadership should be able to answer the following six questions with confidence and evidence:

1. Which AI agents are operating across the organization?
2. Who owns each agent and under whose authority does it act?
3. What are they permitted to access and do?
4. When is human oversight, approval, or step-up authentication required?
5. Can every AI agent action be monitored, attributed, reconstructed, and audited?
6. Can AI agent access be quickly constrained or revoked if needed?

Visit the Imprivata Agentic Identity Management page to learn more about securing and governing agentic AI at enterprise scale.

imprivata®

Imprivata delivers simple and secure access management solutions for healthcare and other mission-critical industries to ensure every second of crucial work is both frictionless and secure. The Imprivata platform of innovative, interoperable access management and privileged access security solutions enables organizations to fully manage and secure all enterprise and third-party identities to facilitate seamless user access, protect against internal and external security threats, and reduce total cost of ownership.

For more information, please contact us at:

Global headquarters USA
Waltham, MA
Phone: +1 877 663 7446
www.imprivata.com

European headquarters
Uxbridge, England
Phone: +44 (0) 208 744 6500
www.imprivata.com/uk

Germany
Langenfeld
Phone: +49 (0) 2173 99 385 0
www.imprivata.com/de

Australia
Melbourne
Phone: +61 3 8844 5533