imprivata®

EXECUTIVE BRIEF

The Agentic AI Trust Gap

Why healthcare needs identity-led governance

Bottom line: Healthcare is moving from AI experimentation to live deployments of agents that can access systems, make decisions, and act with varying degrees of autonomy. Leaders need more than visibility into the agents operating in their environment. They need to ensure every agent is known, authorized, constrained, monitored, and accountable.

What the research shows

A Vanson Bourne survey of 250 U.S. healthcare leaders responsible for identity security and/or AI strategy indicates the industry is approaching an inflection point. Agentic AI is entering live healthcare workflows, but the governance models needed to control it are still maturing.

- 28% of survey respondents have agentic AI in production and use, 44% are piloting or conducting proof-of-concept projects, and 21% expect implementation within 12 months
- 88% of respondents expect agents to operate with at least some degree of autonomy across clinical and operational workflows
- Agentic AI is already involved to some degree in 33% of clinical workflows and 37% of operational workflows; respondents expect this to reach 41% and 46% within 12–18 months

The trust gap: Confidence versus control

The findings reveal a gap between leadership confidence and operational reality: 86% of respondents are confident they have or will have visibility into AI agent activity, and 86% are confident they can fully control and govern agent actions today. Yet 72% also say that AI tools or agents are deployed without formal IT approval at least occasionally.

One reason for this is fragmented provisioning. Respondents cite centralized IT management (71%), security management (49%), self-service through approved tools (47%), decentralized departmental provisioning (41%), and ad hoc or unapproved deployment (37%). Multiple approaches may exist within the same organization, making consistent attribution and oversight difficult.

The stakes are higher in healthcare

In healthcare, the stakes are high. AI agents may interact with EHRs, identity systems, clinical applications, medical devices, APIs, and other infrastructure that directly and indirectly supports patient care. As agents gain access to sensitive data and privileged workflows, the risk is not limited to a single compromised tool or task. An agent may retrieve information, invoke downstream tools, or take unintended or unauthorized actions that are difficult to reconstruct after the fact.

Security is a top-three concern for 57% of respondents; 79% cite security, identity, or governance as a major barrier to scaling

57% rank excessive or unnecessary access permissions among their top concerns

Nearly half cite unauthorized access to PHI (47.6%) or compliance and regulatory violations (49.6%) as top concerns

Autonomy must be risk-based: 53% require human review for high-risk clinical actions, compared with 40% for high-risk operational actions

Implication for senior leaders

Identity should be the foundation for governing AI agents. Each agent should have a defined owner, purpose, authority, access scope, autonomy level, and lifecycle. Existing IAM practices provide a starting point — including least privilege, permission management, monitoring, and auditability. But agentic AI requires organizations to extend these practices to govern delegated authority, dynamically adjust permissions based on risk and context, and attribute each action to the agent and the authority under which it acted.

Governing AI agents is an enterprise responsibility. IT and security are central, but leaders in clinical, compliance, legal, and operational departments must help classify use cases, set approval thresholds, define accountability, and decide where human oversight is mandatory.

Five actions to take now

1. Discover and assign ownership
Inventory every agent, including shadow AI; document its purpose, owner, sponsor, deployment status, connected systems, data, and downstream tools.

2. Give every agent a unique, governable identity
Record whose authority it operates under and manage provisioning, changes, suspension, delegated authority, and decommissioning.

3. Limit access and authority
Apply least privilege based on workflow risk, data sensitivity, and autonomy. Separate permission to access information from authority to take action; use time-bound or context-aware controls for higher-risk activity.

4. Set risk-based guardrails
Define which actions an agent may perform independently and which require human review, approval, or step-up authentication. Test before expanding autonomy or entering higher-risk clinical workflows.

5. Monitor agents and alert on anomalous behavior
Monitor agent activity and alert on anomalous behavior. Ensure access can be constrained or revoked quickly.

Leadership test

Before scaling agentic AI, leadership should be able to answer seven questions with evidence:

1. Which agents are operating?

2. Who owns each one and under whose authority does it act?

3. What can it access and do?

4. When is human oversight required?

5. Can every action be monitored, attributed, reconstructed, and audited?

6. Can authority be quickly constrained or revoked?

7. Is agentic network access limited only to the resources and protocols absolutely necessary to perform its functions?

Healthcare-specific risks require stronger controls

Agentic AI can improve efficiency, productivity, and aspects of patient care, but its value depends on trust that agents will operate within approved boundaries. Healthcare organizations that build identity-led governance now will be better positioned to scale increasingly autonomous AI agents while preserving security, accountability, compliance, and patient trust.

Source: The Agentic AI trust gap: Why healthcare needs identity-led governance. Independent research conducted by Vanson Bourne for Imprivata; survey of 250 U.S. healthcare leaders responsible for identity security and/or AI strategy.

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

7695-2026_WP-Agentic AI Trust Gap