imprivata

When AI Agents Enter Healthcare, They Need an Identity

HSCC's guidance for securing agentic AI: treat AI agents like a new class of insider — because in clinical systems, they can act like one.

The new healthcare AI reality

AI agents are moving from recommendations to action. They may help with documentation, triage, care coordination, scheduling, revenue cycle, pharmacy workflows, and EHR-driven tasks. But when an AI agent can access sensitive systems or patient data, traditional application security is not enough.

Healthcare risk lens:

An agent with excessive EHR permissions can expose PHI, disrupt workflows, make unsafe changes, or create downstream patient safety risk. HSCC states that an AI agent with EHR access and excessive permissions is functionally equivalent to a compromised insider with clinical system privileges.

HSCC's five agentic AI control areas

1. IDENTITY MANAGEMENT

Every AI agent needs a unique identity, documented capabilities, and defined boundaries.

2. CREDENTIAL MANAGEMENT

AI credentials should be managed with the same rigor as privileged service accounts.

3. BEHAVIORAL BASELINES

Organizations should establish normal operating patterns for each agent to enable anomaly detection.

4. ROGUE DETECTION

Monitoring should identify deviation from expected behavior, unexpected data access, or privilege escalation.

5. EHR ACCESS CONTROLS

Agents with access to clinical systems must be scoped and constrained appropriately.

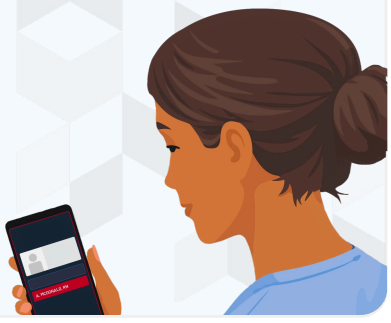

What makes healthcare different?

- **Workflow speed:** Controls cannot slow clinicians down or disrupt care delivery.
- **Patient privacy:** Agent activity may involve PHI and must be visible, auditable, and explainable.
- **Clinical safety:** For sensitive actions, AI should recommend — not autonomously execute — medication changes or order modifications; HSCC calls for least privilege and confirmation for sensitive operations.
- **Legacy systems:** Many clinical systems were not designed for AI agents or modern APIs. Imprivata positions Agentic Identity Management as a way to broker secure access to legacy, on-prem, and modern systems, including systems without native APIs or AI integrations.
- **Continuous monitoring:** AI risk is behavioral, not just code-based; HSCC says continuous behavioral monitoring and periodic adversarial testing should continue post-deployment.

The practical control model

Register the agent → Assign unique identity and approved purpose

Limit the agent → Enforce least privilege and workflow boundaries

Protect credentials → Avoid static credential exposure

Watch behavior → Baseline normal patterns and detect anomalies

Control actions → Require human oversight for sensitive clinical workflows

Revoke fast → Limit, terminate, or revoke access when risk appears

Agentic AI can help healthcare move faster. Identity-first governance helps it move safely.

Learn how Imprivata helps healthcare organizations secure AI agent access.