AI Governance Before Enablement

This article is part of a series expanding on the themes explored in **Executive Brief: The Meditology 2026 Healthcare Security Outlook**. Each installment takes a deeper look at the strategic priorities shaping healthcare cybersecurity in 2026, drawn from executive interviews with CISOs and senior security leaders across the industry.

MEDITOLOGY SERVICES

EXECUTIVE SUMMARY

- **Metric Shift:** AI adoption shifts from speed of deployment and feature enablement to governed use, data readiness, and measurable risk reduction.
- **The Goal:** Moving from "How quickly can we deploy AI tools?" to "Do we have the governance, controls, and data maturity required to use AI safely and sustainably?"
- **The Bottom Line:** Organizations that govern AI before enabling it reduce "shadow AI" risk, protect sensitive data, and create the foundation for step-wise AI adoption.

AI arrived in healthcare like a flash flood: Clinical documentation tools, diagnostic decision support, ambient scribes, and administrative automation appeared almost overnight. For many health systems, the clinical demand for these tools outpaced the ability of security teams to build guardrails. By the time the scale of adoption was recognized, the priority had shifted from enablement to containment.

Industry observers have called 2026 the year of AI governance in healthcare,[1] and the framing is accurate, but it understates the urgency. Health system leadership is not adopting governance because the technology has matured. They are adopting it because the risks of not governing have become impossible to ignore.

Across our interviews with CISOs and senior security leaders for the *2026 Healthcare Security Outlook*, AI emerged as a universal priority, and a universally unresolved one. The consistent message: **governance must precede enablement.**[2]

“Shadow AI”: The Risk That Is Already Inside

The most immediate AI risk facing healthcare is not adversarial. It is internal. “Shadow AI” (the unauthorized use of consumer AI tools by employees) represents the primary data disclosure concern across the industry.

Clinicians, administrators, and operational staff are adopting AI tools faster than security teams can evaluate and approve them, creating uncontrolled pathways for sensitive data to flow into public models. A physician pastes a clinical summary into ChatGPT to draft a patient letter. A billing coordinator uploads claims data to an AI tool to accelerate coding. A researcher feeds de-identified (but not truly anonymized) datasets into a public model for analysis. Each instance creates an unmonitored channel through which protected health information can leave the organization.

The data confirms the scale of the problem. A 2025 IBM report found that 97 percent of organizations with AI-related security incidents lacked proper AI access controls, and that organizations with high levels of shadow AI reported higher breach costs, contributing an additional $200,000 to the global average breach cost.[3] A Wolters Kluwer survey of over 500 healthcare administrators and providers found that patient safety was the top concern around AI tools, yet

only 18 percent of healthcare professionals were aware of formal organizational policies governing generative AI use, and only 20 percent had received structured training on authorized AI tools.[4] Meanwhile, 86 percent of healthcare IT executives reported instances of shadow IT in their health systems, and shadow AI represents a deeper, more dangerous layer of the same problem.

> "AI is moving faster than our ability to secure it."
>
> — Healthcare CISO

As one CISO presenting at RSAC 2026 noted, clinicians and nurses are not using shadow AI to be evasive. They want to be more efficient. The tools promise relief from the administrative burden that consumes one to two hours of EHR documentation for every hour of direct patient care.[5] The motivation is understandable. The risk is not.

This challenge mirrors the shadow IT problem of the previous decade, but with materially higher stakes. The data at risk includes clinical records. The tools in question are designed to retain and learn from their inputs. And unlike a personal Dropbox account, a consumer AI model cannot be wiped once PHI has been ingested.

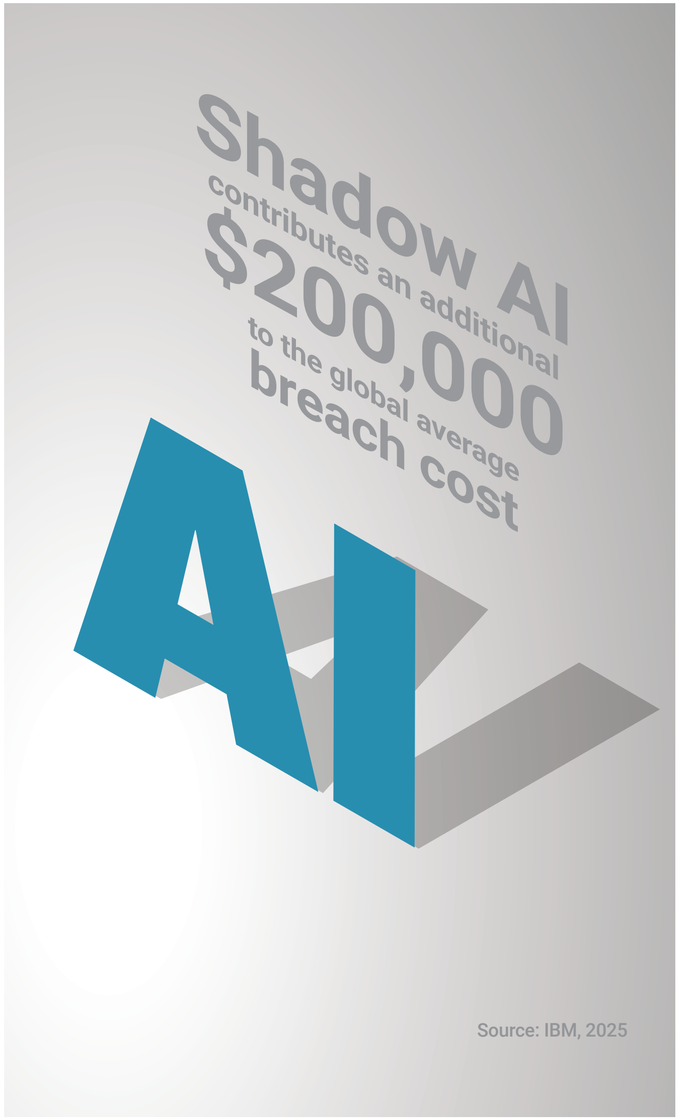

The Data Maturity Ladder: Why Most Organizations Aren't Ready

Even where governance structures are taking shape, many organizations are discovering a more fundamental problem: their data isn't ready for AI.

One of the most instructive frameworks to emerge from our interviews was a maturity progression that places AI enablement at the top of a five-stage ladder: standardized operational processes, structured data capture, visualization and analytics, risk-based decision making, and finally AI enablement. Organizations that attempt to leap directly to AI without building the preceding layers will stall.

This insight reframes the AI conversation from a technology problem to an operational discipline problem. Before investing in AI-driven automation, organizations must first centralize their issue tracking, standardize intake processes, normalize their data, and build the visualization layer that turns raw security telemetry into decision-ready intelligence. The organizations moving fastest toward AI readiness are those investing most heavily in these unglamorous foundational capabilities, not in AI tools themselves.

The implication is sobering for an industry eager to capture AI's efficiency gains. If your security data is scattered across dozens of unintegrated tools, each identifying assets differently and telling only one part of the risk story, an AI layer on top of that fragmentation will not produce insight. It will produce noise with a veneer of confidence. That may be worse than no AI at all.

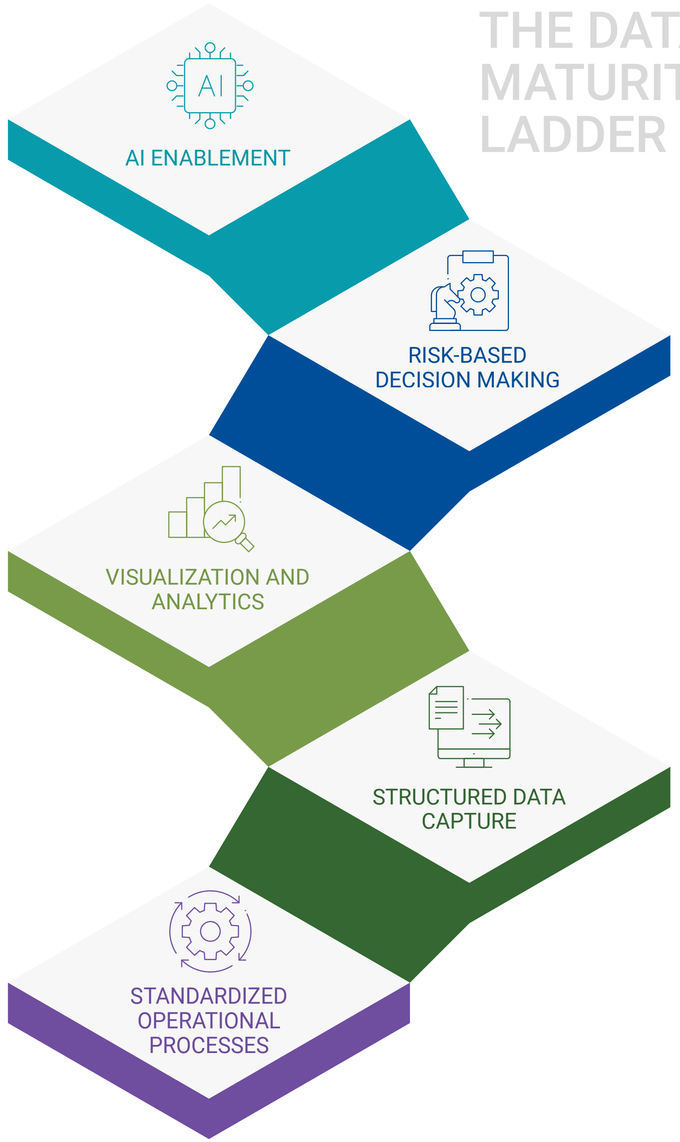

Building the Governance Structure

The CISOs we interviewed described governance structures that are emerging along a recognizable pattern. The organizations furthest ahead share several common elements.

Cross-Functional AI Steering Committees

Effective AI governance cannot live in the CISO's office alone. The most mature models bring together clinical, legal, compliance, and security leadership into a unified governance body that evaluates AI tools before they are deployed in clinical or administrative settings. One organization described having this structure in place for approximately 18 months, with volume and maturity continuing to increase as AI adoption requests multiply across the enterprise.

These committees treat AI as another technology product that must pass through existing risk review gates, with additional considerations for data handling, model transparency, and patient safety implications. The regulatory landscape reinforces this approach: more than 250 AI-related bills were introduced across 46 states in the past year, with Colorado, Utah, Texas, and California all enacting laws that impose governance, disclosure, and oversight requirements on AI systems in healthcare.[6] The organizations establishing cross-functional

governance now will be far better positioned when enforcement begins than those scrambling to build policies retroactively.

Acceptable Use Policies with Teeth

Policy without enforcement is aspiration. Effective organizations are pairing acceptable-use policies with technical controls that prevent PHI from being ingested by unapproved platforms. This includes DLP tools tuned to detect AI-bound data flows, network-level controls that restrict access to consumer AI endpoints from clinical devices, and endpoint monitoring that flags unauthorized AI tool installation.

Equally important is providing approved alternatives. The most effective way to eliminate shadow AI is to give clinicians and staff sanctioned tools that actually reduce their burden. When an organization deploys enterprise-grade ambient scribes, EHR-integrated documentation assistants, and compliant AI-powered coding tools, the motivation to reach for ChatGPT diminishes. As one CISO framed it: the goal is to enable the business to reap the benefits of AI in a thoughtful and controlled way, not to ban the technology that clinicians are telling you they desperately need.

Third-Party AI Risk

Shadow AI is not the only governance challenge. Organizations must also evaluate the AI capabilities embedded in the vendor products they already use. EHR platforms, clinical decision support tools, revenue cycle systems, and TPRM platforms are all integrating AI features, sometimes without explicit disclosure to the customer. One CISO noted that a governance discussion around AI policy raised an immediate follow-on question: we need to know the AI policy and practice of all of our third parties.

This extends the TPRM challenge explored in Article 2 of this series. Vendor AI usage (what data is being fed into models, whether outputs are used in clinical decisions, and how model behavior is monitored) is rapidly becoming a required element of vendor risk assessment. Premier Inc.'s 2026 report found that 80 percent of health systems lack internal governance standards to guide AI adoption, and that most are moving forward with siloed point solutions rather than a thoughtful AI ecosystem.[7] Without governance that extends across the vendor landscape, organizations are building AI strategies on foundations they do not control.

Where AI Is Already Delivering Value

Despite the governance challenges, AI is not a future promise in healthcare security; it is delivering measurable value today, particularly in the security operations center.

Organizations report that AI-driven agents can reduce threat triage time from 20 minutes to under five minutes, enabling constrained security teams to process significantly higher alert volumes without additional headcount. By automating repetitive triage, AI enables analysts to focus on the judgment calls that require human context: distinguishing true threats from noise, escalating appropriately, and making risk decisions informed by clinical workflows.

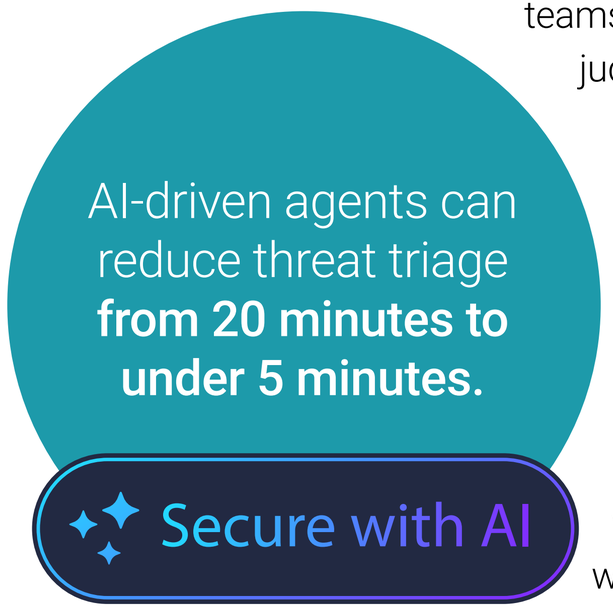

The operational value is clear, and the cost efficiency is real. But leaders stress that realizing this value requires keeping skilled analysts in the loop. The organizations seeing the best results treat AI as a way to extend the reach and impact of their existing teams, accelerating the work without removing the human judgment that makes security decisions meaningful in a clinical context.

The irony is instructive: the area where AI is delivering the most value today, security operations, is also the area where the data maturity prerequisites are most commonly met. SOC telemetry is structured, normalized, and high-volume. These are exactly the conditions under which AI excels. Organizations looking for AI quick wins would do well to start where their data is already ready, rather than forcing AI into domains where the foundational data work has not been done.

Govern First, Enable Second

The pressure to adopt AI in healthcare is enormous and legitimate. Clinician burnout, staffing shortages, rising alert volumes, and the administrative weight of compliance all create genuine demand for automation. The organizations that will capture AI's benefits are not the ones moving fastest; they are the ones building the governance, data foundations, and policy frameworks that make sustainable adoption possible.

That means standing up the cross-functional steering committee. Defining the acceptable use policy. Solving the shadow AI problem before it becomes a breach headline. Mapping vendor AI usage into the TPRM framework. And, perhaps most importantly, investing in the data normalization and operational discipline that determines whether AI produces insight or illusion.

AI governance is not a barrier to innovation. It is the precondition for innovation that lasts.

About Meditology Services and this Research

Meditology Services is a healthcare-exclusive cybersecurity advisory firm. As the facilitator of this research and a trusted partner to hundreds of healthcare organizations nationwide, Meditology translates peer insights into actionable programs that accelerate security program maturity.

To execute on the strategic priorities outlined in this report, Meditology partners with healthcare organizations across critical areas:

Operationalize Resilience & Third-Party Risk (TPRM)

Transition from reactive vendor assessment to proactive supply chain resilience through architectural risk reviews, continuous monitoring, and vendor integration into business continuity planning.

Achieve Industry-Standard Certifications

Pursue HITRUST certifications and SOC 2 attestations to provide validated proof of security maturity to partners, patients, and regulators.

Validate Security Through Continuous Technical Testing

Replace point-in-time assessments with continuous testing and ethical hacking to ensure controls remain effective against an evolving threat landscape.

Establish AI & Emerging Technology Governance

Proactively manage the risks of AI adoption through AI risk program and governance frameworks that ensure innovation does not outpace security.

This executive report is based on in-depth executive interviews conducted during January and February 2026 with CISOs and senior cybersecurity leaders representing multi-state integrated delivery networks, academic medical centers, regional and community health systems, health plans, and health technology companies.

Discussion areas included enterprise risk prioritization, budget strategy, vendor governance, security operations, regulatory readiness, and emerging technology risk. All findings are reinforced by Meditology's proprietary data derived from hundreds of healthcare cybersecurity engagements conducted annually.

Participants included:

- Gillian Lockwood, Director, Information Security, Mass General Brigham
- Matthew Webb, AVP, Cyber Risk Management, HCA
- Zishan Siddiqui, CISO, Sentara Health
- David Finklestein, CISO, St. Luke's University Health Network
- Raj Raju, CISO, Integra Connect
- Denis Tanguay, CIO, Sturdy Health
- Paul Curylo, CISO, Inova Health System
- Among others who requested anonymity.

All insights are presented in aggregate to reflect industry-wide themes rather than the position of any single organization.

Citations

1 Wolters Kluwer Health, “2026 Healthcare AI Trends: Insights from Experts,” December 2025. https://www.wolterskluwer.com/en/expert-insights/2026-healthcare-ai-trends-insights-from-experts

2 Meditology Services, 2026 Healthcare Security Outlook: The Shift to Operational Resilience, March 2026. https://meditologyservices.com

3 IBM, “Cost of a Data Breach Report 2025.” Referenced in TechTarget, “Shadow AI in Healthcare: The Hidden Risk to Data Security,” September 2025. https://www.techtarget.com/healthtechsecurity/feature/Shadow-AI-in-healthcare-The-hidden-risk-to-data-security

4 Wolters Kluwer Health, “Health System Size Impacts AI Privacy and Security Concerns,” January 2026. https://www.wolterskluwer.com/en/expert-insights/health-system-size-impacts-ai-privacy-and-security-concerns

5 Dark Reading, “Shadow AI in Healthcare Is Here to Stay,” April 2026. https://www.darkreading.com/cyber-risk/shadow-ai-in-healthcare-is-here-to-stay

6 Forvis Mazars, “How the Rise of Artificial Intelligence Affects Patient Data,” January 2026. https://www.forvismazars.us/forsights/2026/01/how-the-rise-of-artificial-intelligence-affects-patient-data

7 Premier Inc., “The AI Wild West Is Over: Why 2026 Is the Year Health Systems Must Take Control,” February 2026. https://premierinc.com/newsroom/blog/the-ai-wild-west-is-over-why-2026-is-the-year-health-systems-must-take-control

About This Series

This article is the third in a series of deep dives from the **Executive Summary Brief: The Meditology 2026 Healthcare Security Outlook**. *Previous installments include Article 1: The Rise of Operational Resilience and Article 2: Third-Party Risk as Operational Accountability. A forthcoming piece will explore operationalizing GRC programs for healthcare.*

MEDITOLOGY SERVICES

To access the full report or to speak with our team, visit **www.meditologyservices.com**.