Healthcare Cybersecurity & AI Benchmarking Study 2026

Executive Summary

From Cybersecurity Fundamentals to AI Governance Readiness

A collaborative health industry study sponsored by Censinet, the American Hospital Association (AHA), Health-ISAC, Health Sector Coordinating Council (HSCC), the Scottsdale Institute (SI), and The University of Texas at Austin.

March 2026

Introduction

Healthcare cybersecurity has stabilized, but governance is not keeping pace with AI adoption.

For the fifth consecutive year, the Healthcare Cybersecurity Benchmarking Study benchmarks U.S. healthcare providers against the NIST Cybersecurity Framework (CSF). This year, we are expanding the study to measure AI governance maturity against the NIST Artificial Intelligence Risk Management Framework (AI RMF) and a supplemental AI governance questionnaire examining governance structures, approval processes, inventory practices, oversight mechanisms, emerging AI risk concerns, and agentic AI uses.

In the 2026 study, 54 organizations participated through structured self-assessments, representing a cross-section of the healthcare industry: small community hospitals, large academic medical centers, and multi-state health systems across both rural and urban settings.

Four Key Takeaways

1 Healthcare cybersecurity is maturing.

Most healthcare cybersecurity programs operate at the NIST CSF Risk-Informed tier, with many approaching the threshold for the Repeatable tier. The clustering of scores near this boundary suggests the industry is transitioning from risk-aware practices toward more formalized, consistently implemented cybersecurity governance.

2 Reactive cybersecurity strengths still mask foundational gaps.

For the third consecutive year, Response and Recovery functions lead performance while Govern and Identify functions trail. Asset Management and Supply Chain Risk Management have ranked lowest for three consecutive years. This is also reflected in the AI inventory visibility gap: only 30% of organizations maintain an enterprise-wide AI inventory.

3 AI adoption is outpacing governance.

AI risk management capabilities remain limited across organizations, yet more than one in four organizations are already running agentic AI in production. Meanwhile, nearly one in three organizations have no formal AI governance committee or approval process. The infrastructure to govern this technology is not keeping pace with their deployment.

4 AI governance challenges are structural, not technical.

Cybersecurity performance is consistent across organizations, but AI governance varies widely. The differentiator is governance structure, specifically the presence of committees, formal approval processes, and defined accountabilities, not technical capability.

NIST CSF 2.0: A Strong Cybersecurity Foundation with an Uneven Profile

Based on the NIST CSF four-tier model, 74% of organizations operate at Tier 2 (Risk-Informed), with an additional 17% reaching Tier 3 (Repeatable). The industry has converged on response capabilities: the Respond function reaches Tier 3.00 and the Detect function 2.97.

Function-level performance reveals an industry pattern: reactive capabilities such as Respond and Recover functions lead (median 2.96), while proactive capabilities such as Govern and Identify functions trail (median 2.58). The data shows that 78% of organizations score higher on these reactive functions.

Within the NIST CSF four-tier methodology, this gap reflects the difference between risk-informed practices and formalized, repeatable processes.

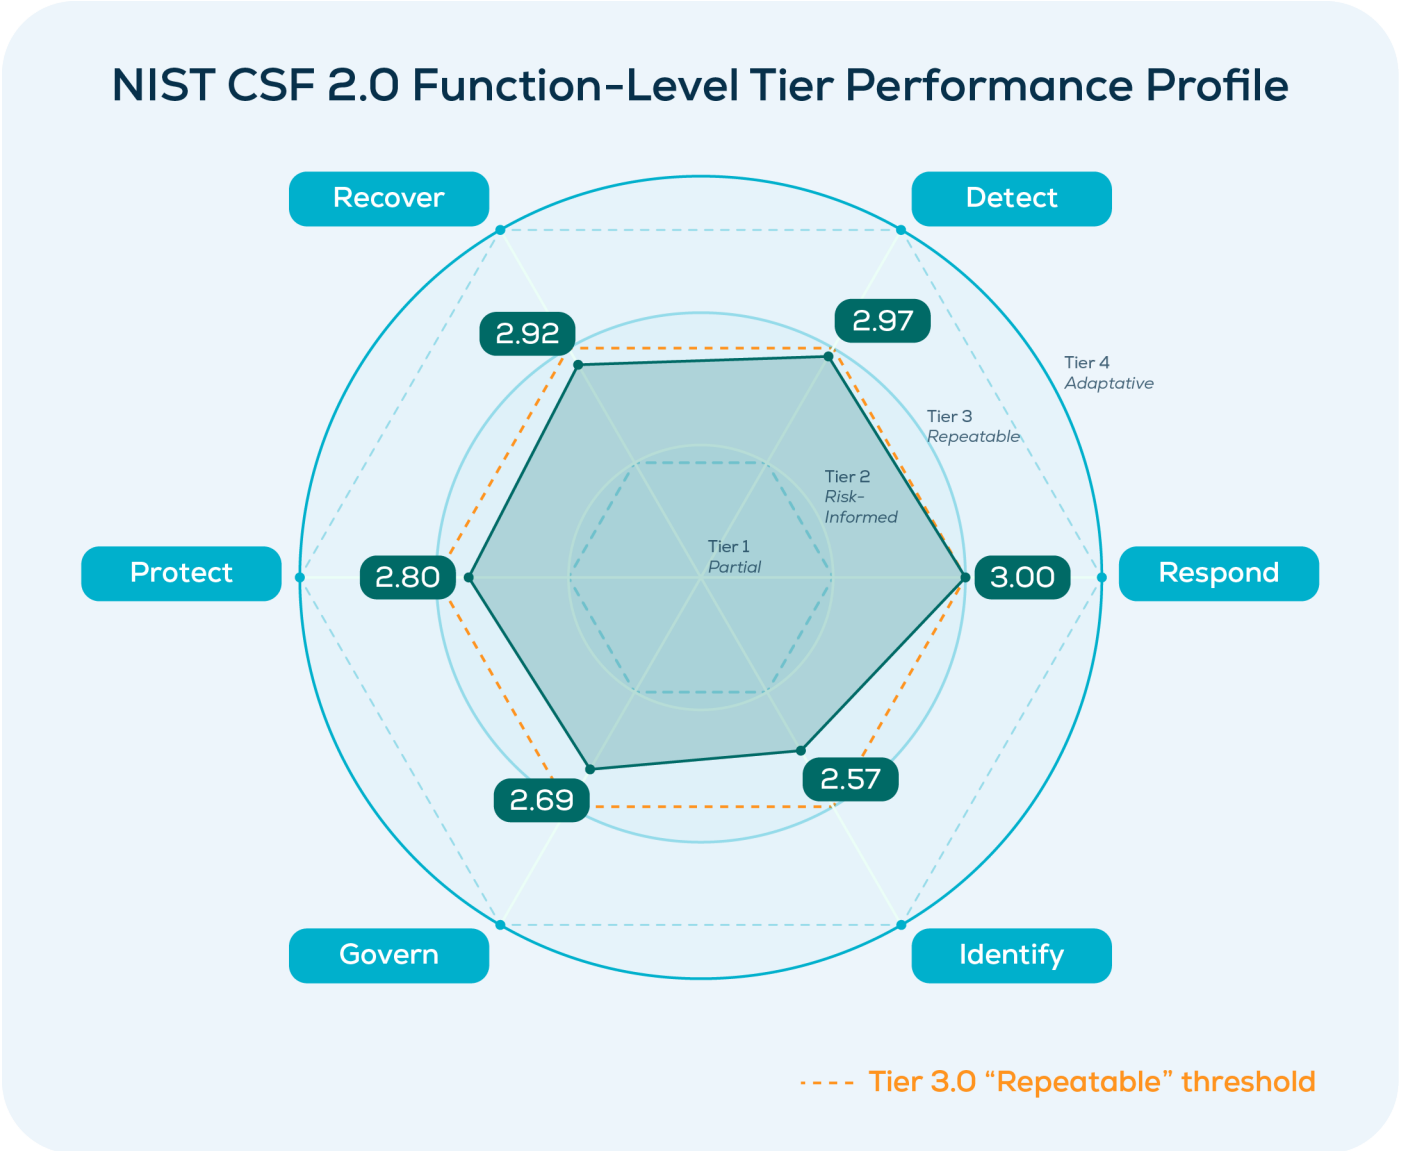

Figure 1: Respond and Detect lead NIST CSF performance.

Two foundational NIST CSF categories have ranked lowest for three consecutive years: Asset Management (Tier 2.43) and Supply Chain Risk Management (Tier 2.50). These are core disciplines needed for governance. Asset Management determines whether an organization has visibility into its technology environment. Supply Chain Risk Management determines whether it understands and oversees risks introduced by third parties. Both are structural prerequisites for effective AI risk management. Organizations cannot maintain accurate AI inventories without mature asset visibility, and they cannot govern vendor-embedded AI without disciplined supply chain oversight.

Board-ready Insight:
The two cybersecurity categories with the most persistent gaps, Asset Management and Supply Chain Risk Management, are also the structural prerequisites for AI governance.

NIST AI RMF: Early-Stage Maturity Across the Board

The NIST AI RMF provides a structured model for identifying and managing AI risk across four core functions: Govern, Map, Measure, and Manage. The 2026 study is the first to benchmark healthcare provider organizations against this framework at scale.

As seen in Figure 2, results show early-stage maturity across all functions. No AI RMF function exceeds 38% coverage. Govern leads at 38%, followed by Map at 28%, Manage at 22%, and Measure at 18%. None reaches “Substantial” coverage under the study’s definitions.

The ordering reflects a predictable maturity pattern observed in the data. Organizations typically begin by establishing governance structures such as committees, policies, and defined accountabilities. Progress declines as capabilities require operational execution. Risk identification (Map) follows governance formation, while risk treatment (Manage) and systemic measurement and monitoring (Measure) lag further behind.

At the AI RMF category level, Policies, processes, procedures, and practices (Govern) show comparatively stronger adoption at 45% coverage. Operational capabilities, such as Mechanisms for Tracking Identified AI Risks (Measure) at 12%, Evaluation of AI Systems (Measure) at 24%, and Risk Treatments (Manage) at 27%, remain limited. This indicates that many organizations have established oversight structures but have not yet operationalized systematic AI risk measurement, monitoring, and lifecycle management.

The industry has begun building AI risk management structures. The next phase is operationalizing them.

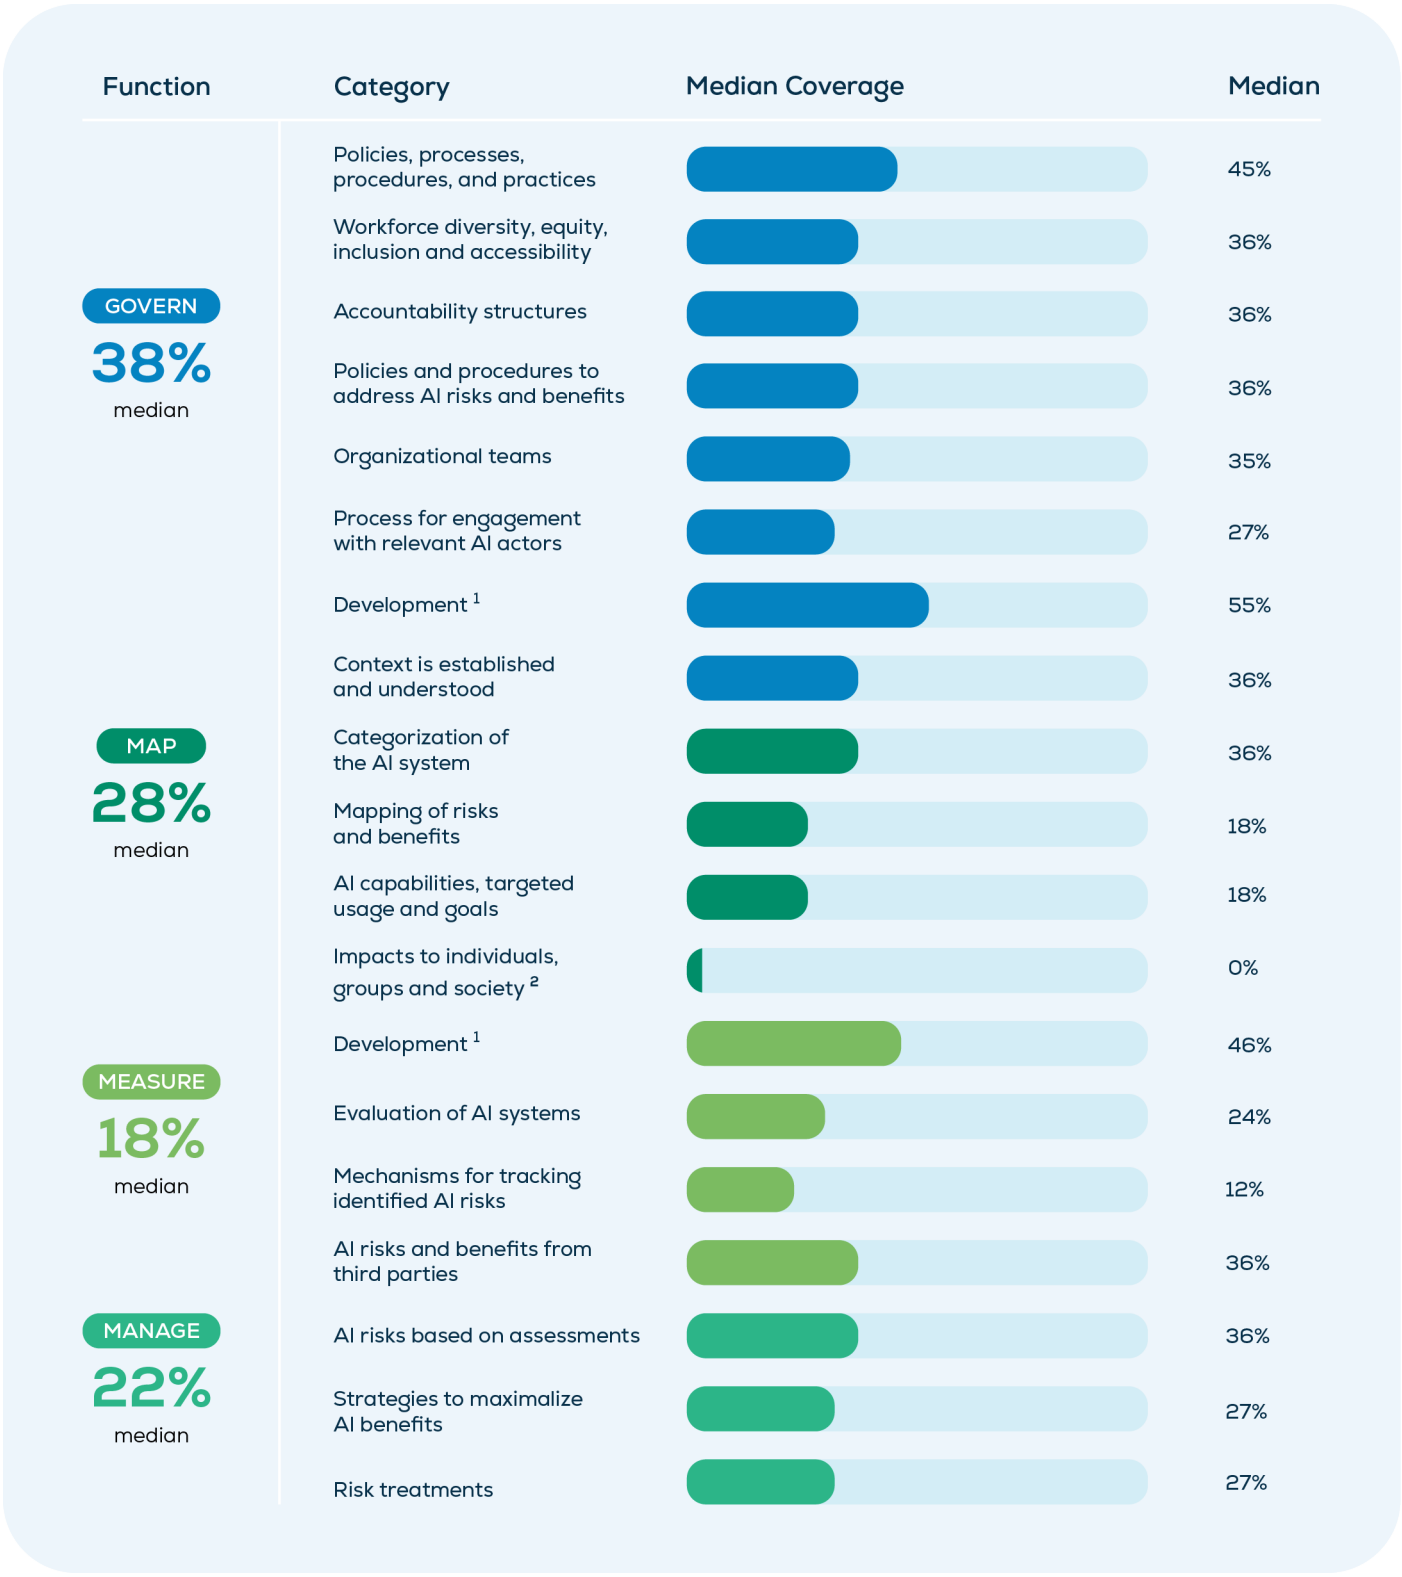

Figure 2: NIST AI RMF adoption remains early stage.

[1] Optional question; n=18; limited sample size, results should be interpreted with caution. Survey questions for NIST AI RMF were separated into two tracks: one for organizations that develop AI and one for organizations that do not. Survey questions related to AI development were relocated under Development, a category not original to NIST AI RMF. This methodology ensures that each organization benchmarks against peers with a similar approach to AI use or development.

[2] The median coverage for "Impacts to individuals, groups, and society" is 0% because fewer than half of organizations demonstrate documented practices aligned with this AI RMF category. The mean coverage is 22%, reflecting a small number of organizations with more developed practices in this area.

The AI Governance Gap: Intent Without Operational Visibility

The most actionable finding in this study is the disconnect between governance intent and operational visibility.

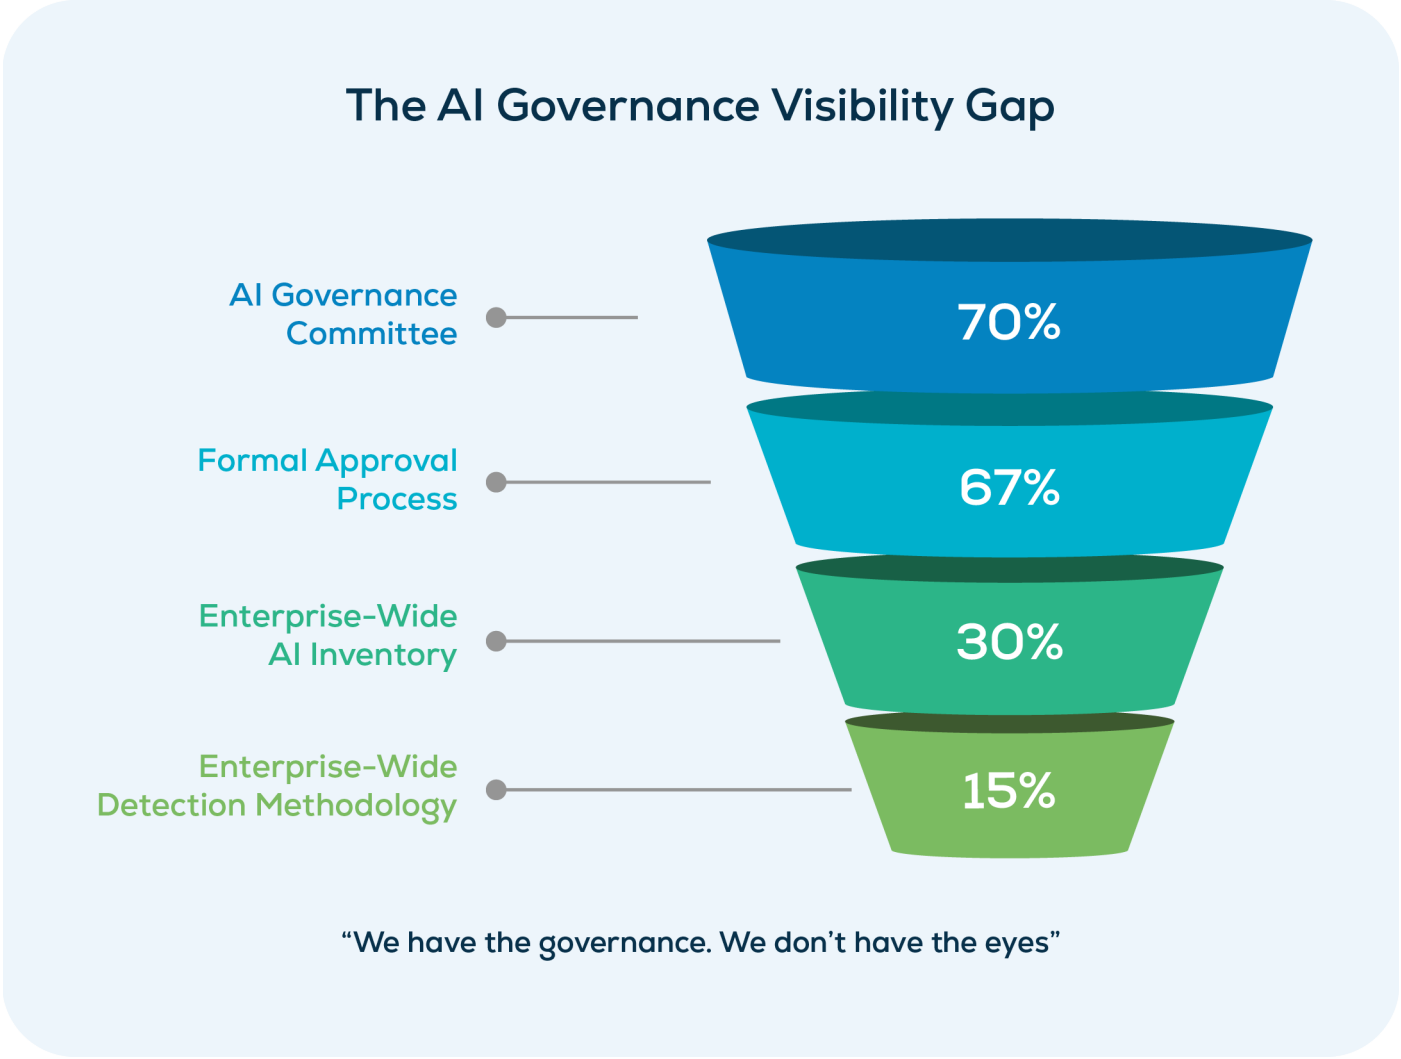

Figure 3: Governance adoption outpaces mechanisms for enterprise AI visibility.

Governance structures are increasingly common, with 70% of organizations reporting an AI governance committee and 67% reporting a formal approval process. Yet only 30% maintain an enterprise-wide AI inventory, and just 15% have an enterprise-wide detection methodology. Governance structures exist. Operational mechanisms to identify and track AI across the enterprise do not.

This gap is material because AI adoption is not waiting for governance to mature. Vendors are activating AI capabilities within products that organizations already own. One-third of organizations report having no documented methodology for detecting vendor-embedded AI. Among those that do, the most common approach remains informal or ad hoc discovery (see Figure 4).

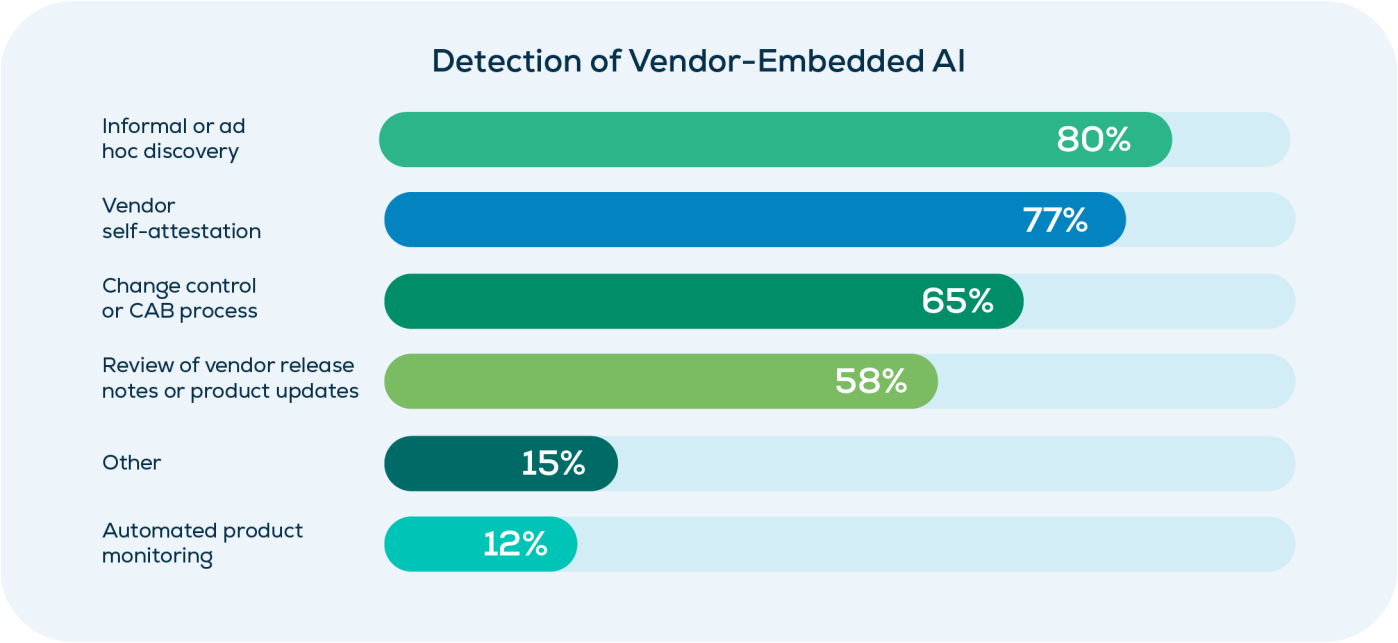

Figure 4: Informal discovery dominates methods for detecting vendor-embedded AI.

The pace problem compounds this visibility gap. Nearly a quarter of organizations (23%) report that AI adoption is moving faster than governance can keep up with. Counterintuitively, this pressure is reported most often by organizations with established governance structures. Half of organizations with governance committees say AI adoption is moving faster than desired, while only 7% of Pre-Governance organizations report similar concerns.

Organizations without governance structures may not perceive pace pressure because they lack the mechanisms to detect it. At the same time, more than one in four organizations report running agentic AI in production, including 27% of organizations with no governance committee, formal approval process, or defined oversight structure.

Board-ready Insight:
Governance without enterprise-wide inventory and detection remains incomplete. The persistent weakness in supply chain risk management is not only a cybersecurity concern, it is also the primary pathway through which undetected AI enters the organization.

What the Industry Is Watching– And What It Is Missing

Healthcare organizations are most concerned with visible, immediate AI data risks: output quality or inconsistency (34%), hallucinations (25%), and biased or unfair outcomes (15%). These are real risks that warrant attention.

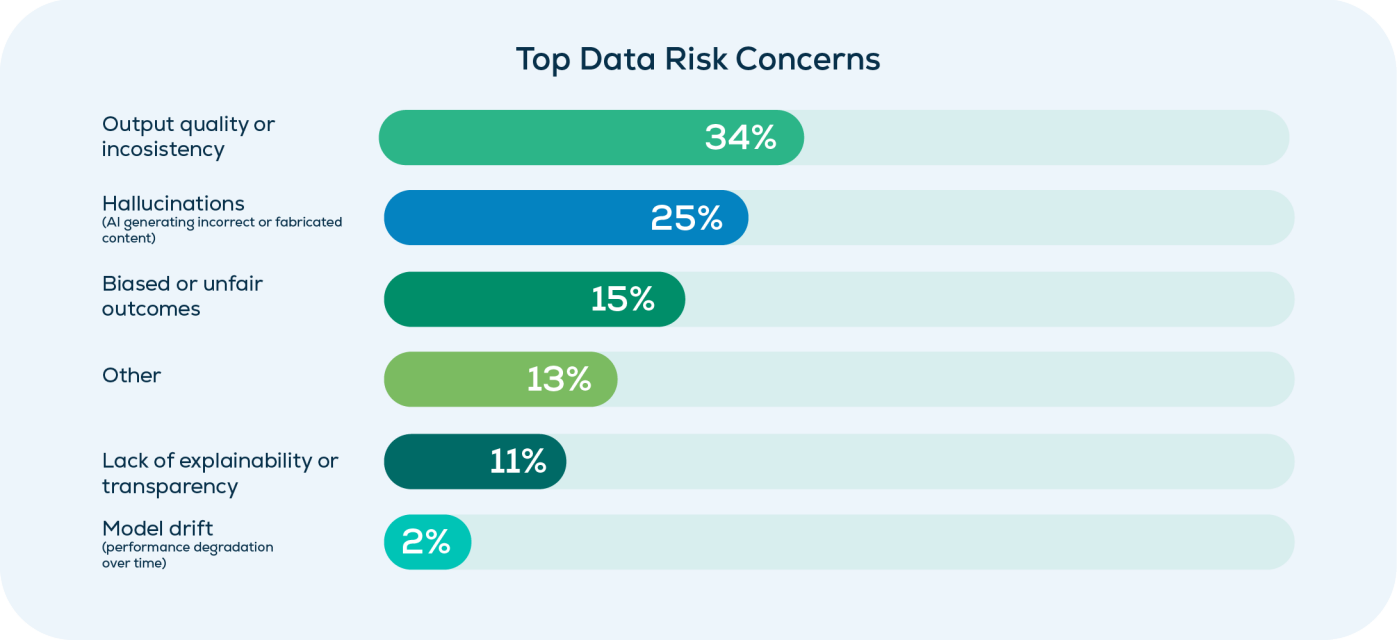

Figure 5: Output reliability concerns overshadow model drift risks.

Model drift, the gradual degradation of AI performance over time, ranked as the top data risk concern for just 2% of respondents. In healthcare environments where clinical protocols evolve, patient populations shift, and documentation practices change, unmanaged drift can quietly erode model reliability, affecting diagnostic accuracy, treatment decisions, and clinician trust. As governance fundamentals mature, lifecycle risks like model drift may demand greater attention.

Four Governance Archetypes– And Why They Matter

This section represents one of the study's most novel contributions. Organizations cluster into four distinct AI governance archetypes based on committee presence, formal approval processes, and observable governance practices. These archetypes provide healthcare leaders a practical new lens for understanding where their organization stands, and what must happen next to advance.

Archetype	% of Study	Key Characteristics
Pre-Governance	28%	No AI governance committee; no formal approval process; AI adoption is slower than desired
Committee-Led, Early-Stage	8%	AI governance committee formed; formal approval process not yet established; low confidence in identifying AI risks
Committee-Led, Developing	15%	AI governance committee and formal approval process in place; implementation practices still maturing
Operationalized Governance	49%	AI governance committee; formal approval process; established oversight and execution practices; measurement and lifecycle management capabilities still maturing

The critical insight is that cybersecurity performance remains consistent across archetypes. Cybersecurity tier performance is consistent, ranging from 2.62-2.83. AI RMF coverage, by contrast, varies significantly. Pre-Governance organizations perform comparably on cybersecurity, but trail on AI risk management practices. The differentiator is governance infrastructure: committees, defined accountabilities, and formal approval processes. The gap is structural, not technical.

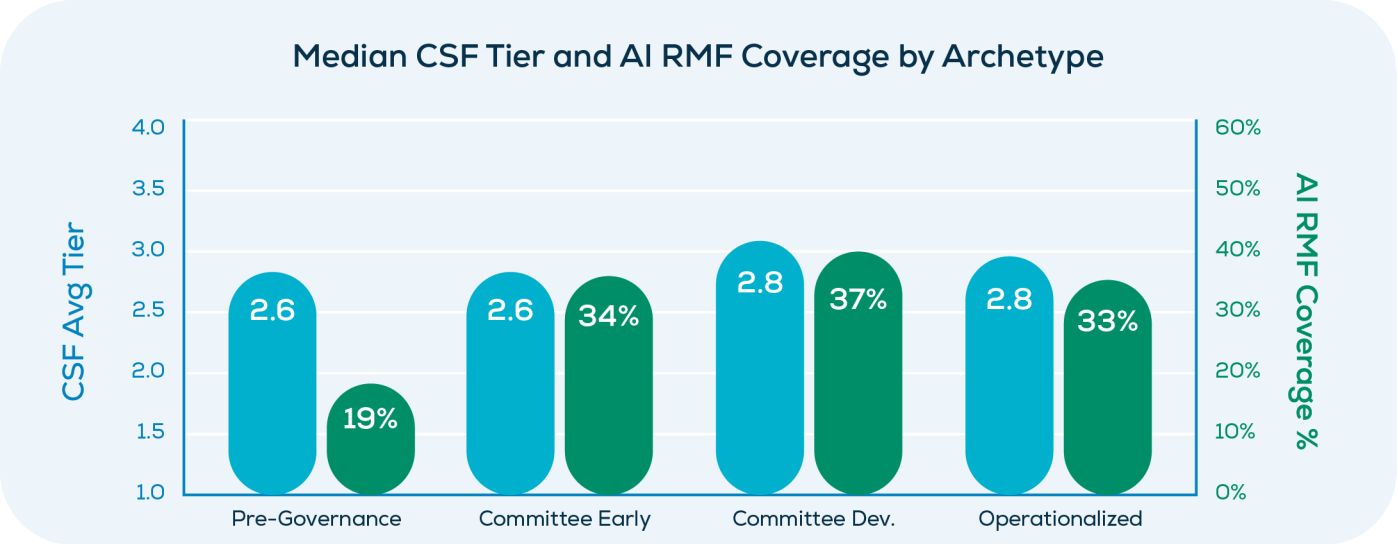

Figure 6: CSF tiers similar; AI RMF coverage diverges.

This distinction is operationally significant. For example, a CISO at a Committee-Led, Early-Stage organization and a CISO at a Pre-Governance organization may have comparable cybersecurity capabilities. However, one has visibility, accountability, and decision rights around AI while the other does not. Archetype classification clarifies immediate priorities. For some, the next step is establishing a committee and approval gates. For others, it is translating structure into enterprise-wide inventory, detection, and lifecycle oversight.

Organization size influences AI readiness but not cybersecurity performance (see Figure 7). CSF tier performance remains consistent across workforce size. AI RMF coverage, however, increases from 15% among small organizations to over 50% among very large health systems. Rural-only organizations show comparable cybersecurity performance but are twice as likely to fall into the Pre-Governance archetype. The constraint is governance infrastructure, not technical capability.

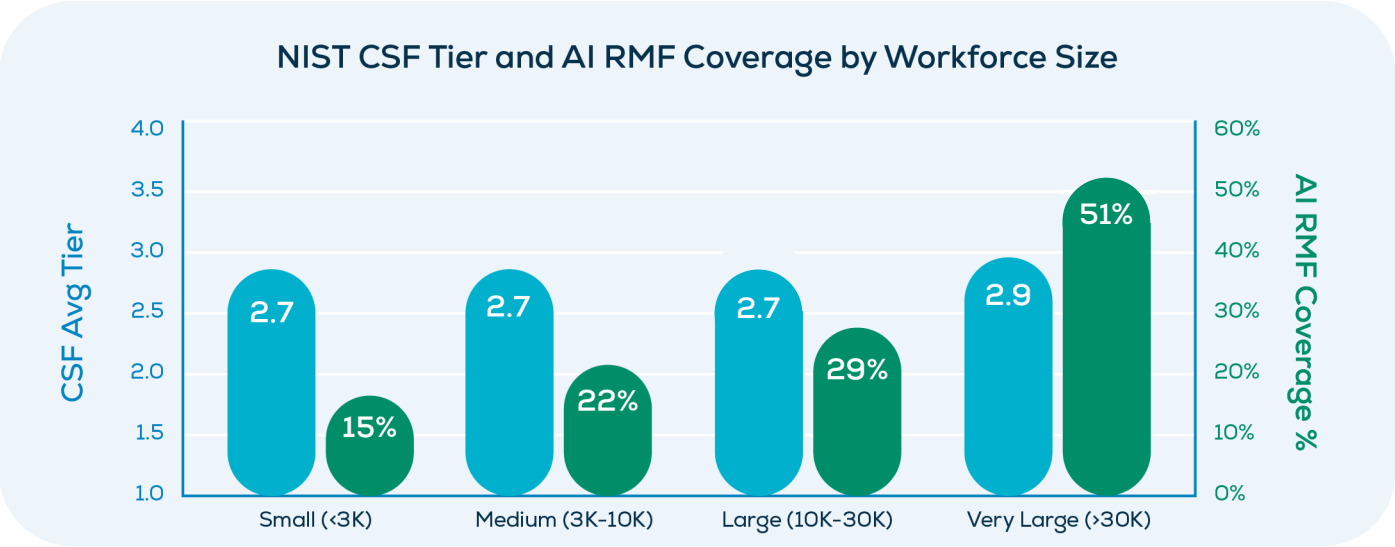

Figure 7: Workforce size predicts AI RMF coverage, not CSF tier.

Across all archetypes, four universal priorities apply:

1. Treat AI governance as a cross-functional mandate spanning clinical operations, privacy, ethics, compliance, IT, and cybersecurity.
2. Close persistent asset management and supply chain risk management gaps that directly constrain AI inventory visibility.
3. Strengthen enterprise-wide inventory and detection capabilities. Governance without visibility is incomplete.
4. Use benchmarking data to support executive and board-level investment decisions.

For Study Participants

In addition to this summary report, study participants have access to a full report which includes detailed analysis across all sections, archetype-specific maturity roadmaps with concrete timelines, category-level benchmarks, and committee composition data. Within the Censinet platform, participants can also see how their organization performs against the study cohort, providing ready-made data for board and executive reporting.

For questions on how to access the full report and in-product benchmarks, contact **benchmarks@censinet.com**.

Report Information

This study is based on structured self-assessments from 54 healthcare provider organizations across a range of sizes, geographic settings, and care delivery models in the United States. Data were collected between September and November 2025. Assessments cover NIST CSF 2.0, NIST AI RMF, and a supplemental AI governance module. Self-attestation is standard practice in benchmarking studies of this type. Findings should be interpreted in that context, as responses reflect each organization's own assessment of its maturity rather than independently verified performance.

About CENSINET

Censinet provides the healthcare industry's leading risk intelligence platform, designed specifically to help organizations manage systemic risk across their digital and third-party ecosystems. Powered by its RiskOps™ platform, Censinet gives healthcare leaders enterprise-wide visibility into vendor, product, cyber, and AI risk, helping them prioritize mitigation efforts that strengthen operational resilience and protect patient care.
Learn more at censinet.com.

Ed Gaudet
CEO & Founder
egaudet@censinet.com

Paul Russell
Chief Product Officer
prussell@censinet.com

David Woska, PhD
Chief Information Security Officer
dwoska@censinet.com

Meredith Miller
Senior Product Manager
mmiller@censinet.com

About Scottsdale Institute

A not-for-profit health system membership organization, the Scottsdale Institute (SI) advances healthcare's digital transformation within an equitable, consumer-centered, community health framework via collaboration, education, and networking. Comprising over 60 not-for-profit health systems and academic medical centers, Members connect through intentionally small, authentic, and informal forums. For today's most pressing issues, SI brings the right audiences together to address healthcare's urgent and long-term challenges.
Learn more at scottsdaleinstitute.org.

Janet Guptill, FACHE, CPHIMS
President & CEO, Scottsdale Institute
jguptill@scottsdaleinstitute.org

John Hendricks
VP, Benchmarking and Research.
jhendricks@scottsdaleinstitute.org

A collaborative study by Censinet, the Scottsdale Institute (SI), the American Hospital Association (AHA), Health-ISAC, The University of Texas at Austin, and Health Sector Coordinating Council (HSCC).