Clearwater

Healthcare – Secure, Compliant, Resilient

Closing the AI Governance Gap

The 2026 Healthcare AI Security and Governance Benchmark



FOREWORD

A message from our president

Three years ago, healthcare's AI conversation was about permission: which tools could be used, by whom, and under what policy. That conversation has largely been settled. Boards approved AI committees. Legal and compliance teams wrote acceptable-use policies. Vendors added AI clauses to contracts. On paper, the industry did the work.

The question in front of healthcare now is different, and harder: Does any of it work?

This report is Clearwater's attempt to answer that question with evidence rather than assumption. We surveyed hospitals, health systems, physician groups, health plans, and the technology companies that build for them, and we asked not just what governance exists, but what it can prove.

The results describe an industry that has taken governance seriously as a documentation exercise and has not yet built the technical muscle to back it up. Only 8.6% of organizations can verify their AI inventory through automated discovery. Only 14.3% fully validate the AI controls of the vendors they depend on. Nearly half rely on a data protection policy with no technical enforcement behind it.

None of this reflects a lack of effort. It reflects a mismatch in pace. AI adoption in healthcare moved from pilot to production in a few budget cycles, and appears to be accelerating. Governance, by its nature, moves more slowly — through committees, policy cycles, and procurement. That is not a flaw in governance; it is what governance is supposed to do. But it means the distance between what organizations have approved and what is actually running in their environment will keep widening unless someone closes it deliberately.



The industry needs to act now, before the gaps are large enough to lead to potential disruptions, security breaches, and compliance issues.

That is the posture we think the industry needs to take in the year ahead: Not more policy, but more proof. Discovery that runs continuously rather than annually. Enforcement that acts on data as it moves rather than describing how it should move. Vendor assessments that produce evidence rather than attestations. Agent governance built before agents are given standing access to clinical systems, not after.

We built this study to give the industry a shared, honest baseline. Next year, we will ask the same questions again, and we expect the numbers to move. Where they move, we will say so plainly. Where they do not, we will say that too.

Healthcare organizations are at different points in their AI governance journey, but the next phase is about demonstrating that policies, controls, and oversight are working in practice. We hope this report helps organizations understand where they stand today, and we look forward to measuring that progress again in 2027.

Baxter Lee
President
Clearwater Security & Compliance

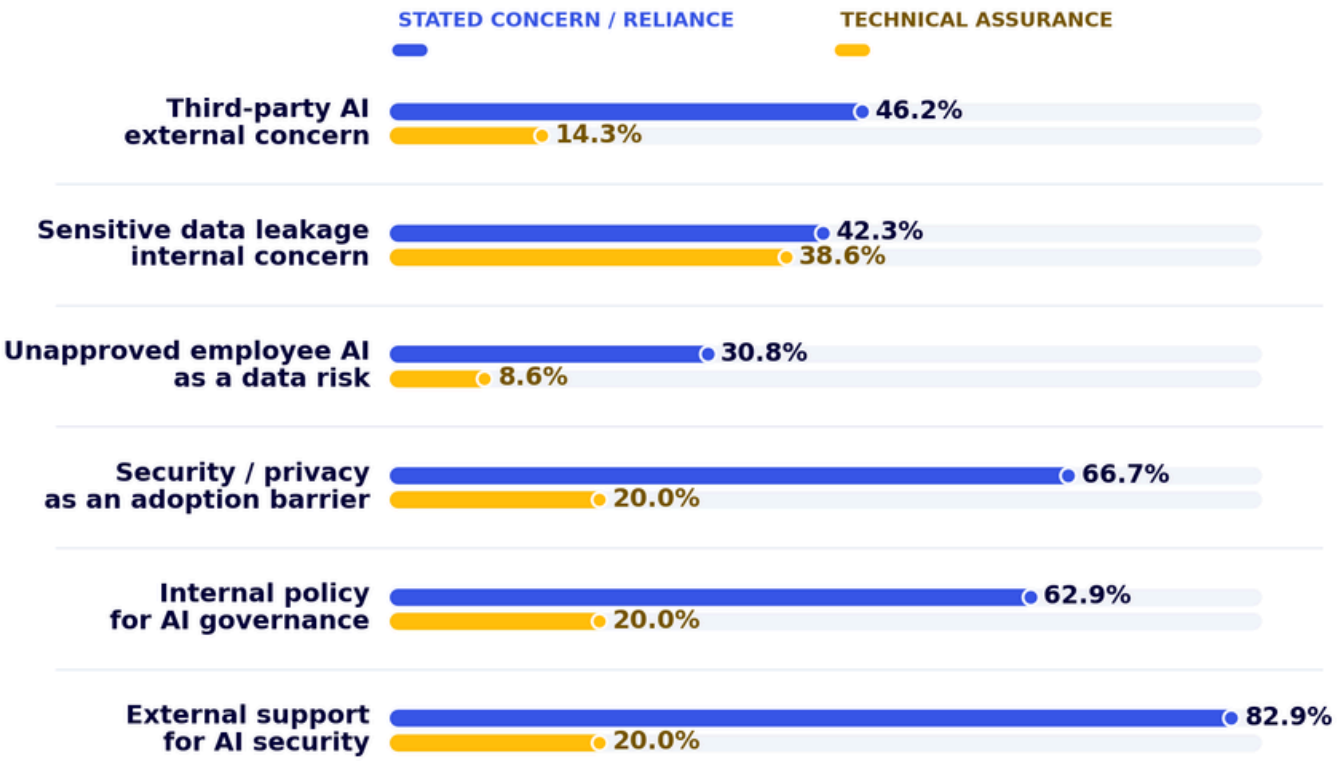
EXECUTIVE SUMMARY

What the data tells us

Healthcare has built the foundations for AI governance. The data shows a widening distance between what organizations are concerned about, what they have documented, and what they can technically verify.

The distance between concern and assurance

Across the study, concern and policy adoption consistently outpace technical verification.



EXECUTIVE SUMMARY

Five findings

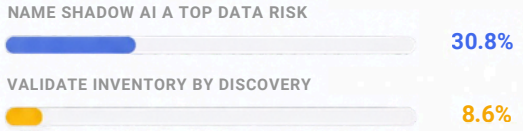
Across every question, the pattern held. Organizations can state what concerns them and point to the policy that addresses it. Far fewer can produce the technical evidence that the control is operating. **Each finding below pairs what the industry reports with what it can back up.**



1

Organizations are beginning to inventory AI, but only 8.6% validate their inventory through automated discovery, leaving 91.4% without a view they can verify.

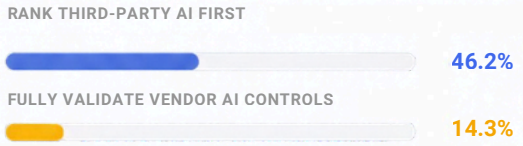
Shadow AI is already visible: 30.8% name unapproved employee AI uses their greatest data risk, second only to sensitive data leaking through AI tools at 42.3%.



2

Third-party AI is the leading external concern but the least verified.

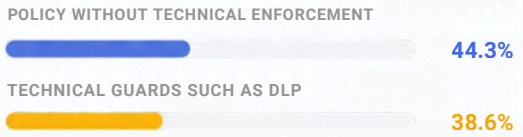
AI introduced through third parties ranks first hand 46%, ahead of AI-powered attacks at 34.6% and reputational damage at 11.5%. Only 14.3% perform full technical validation of third-party AI controls.



3

Organizations are documenting AI controls, but they still lack a reliable way to establish assurance.

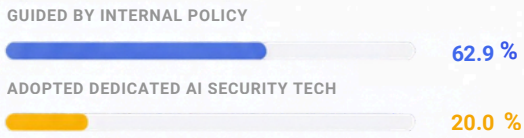
Inventory, third-party, and data protection controls are commonly designed on paper but rarely deployed and tested at the speed of adoption.



4

AI governance awareness is expanding, but AI security frameworks and controls remain thin.

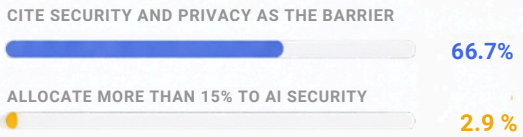
AI in policy guides 62.9%, and 45.7% cite the NIST AI Risk Management Framework, which is broader than security. 42.9% are extending existing tools, while 37.1% remain undecided.



5

Security is the leading adoption barrier but is not not funded accordingly.

51.4% spend less than 5% of annual security budgets on AI security and 28.6% do not know their allocation. The gap appears to reflect limited execution guidance and investment concentrated on adoption and new capability.



TAKEAWAY

Healthcare has taken its first steps toward governing AI, but accelerating adoption continues to outpace governance and verification trails further behind.

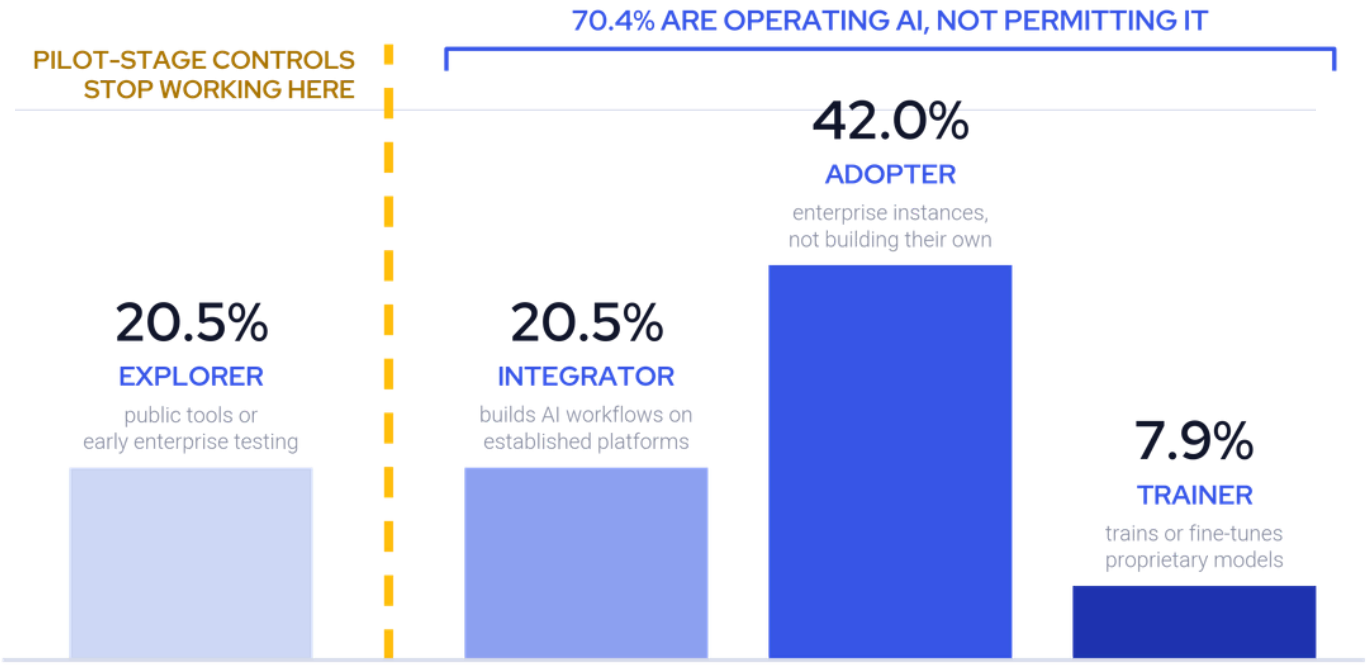
THE LANDSCAPE

AI is already embedded in healthcare's IT environments

Seven in ten respondents have moved beyond permitting AI to operating it.

The industry's AI risk profile is no longer defined by isolated experiments. More than four in ten organizations in this study have enterprise instances running, and one in five is building workflows on development platforms. Eight percent train or fine-tune their own models. Only one in five is still at the exploration stage. The controls appropriate to a pilot were never designed for adoption at scale, and that mismatch is where most of this report's findings originate.

HOW ORGANIZATIONS INTERACT WITH AI



A further 9.1% were unsure.

TAKEAWAY

Organizations still governing AI as an emerging technology are applying pilot-stage controls to production-stage risk.

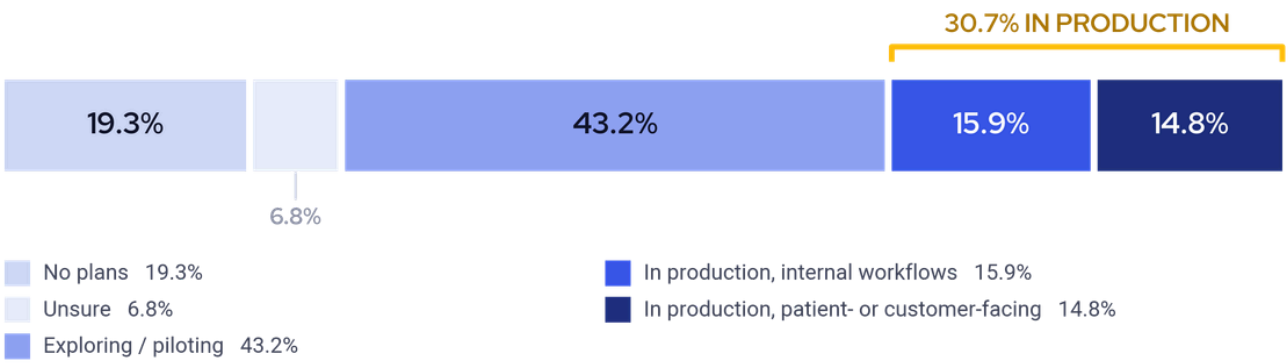
THE LANDSCAPE

Agentic AI is already operating.

About one in seven has it in patient-facing or customer-facing environments.

Agentic systems change what governance has to cover. A model that generates text creates content risk. An agent that calls an application programming interface, operates a tool, or writes to a record creates action risk, and action risk is governed through identity, privilege, logging, and containment rather than through model accuracy. Seventy-four percent of respondents are exploring agentic AI or already running it in some form. Only 19.3% report no plans. The industry is deploying capability faster than it is extending the controls that capability requires.

AGENTIC AI DEPLOYMENT



FIVE QUESTIONS TO ASK ABOUT ANY AGENT

1. What identities and privileges does the agent hold?
2. What systems and data can it reach?
3. Which actions require human authorization?
4. Can its activity be reconstructed after the fact?
5. Can it be suspended in minutes?

CLEARWATER'S INSIGHT

Most organizations can describe what an agent is supposed to do. Far fewer can describe what it is actually allowed to do: which systems it can reach, which credentials it holds, and whether those permissions were ever revisited after launch.

Too often, agent identity is provisioned once, at the broadest scope needed to get a pilot working, and never scaled back. It is the same pattern that created years of over-permissioned service accounts in traditional IT, except now those permissions can be exercised at the speed of an API call.

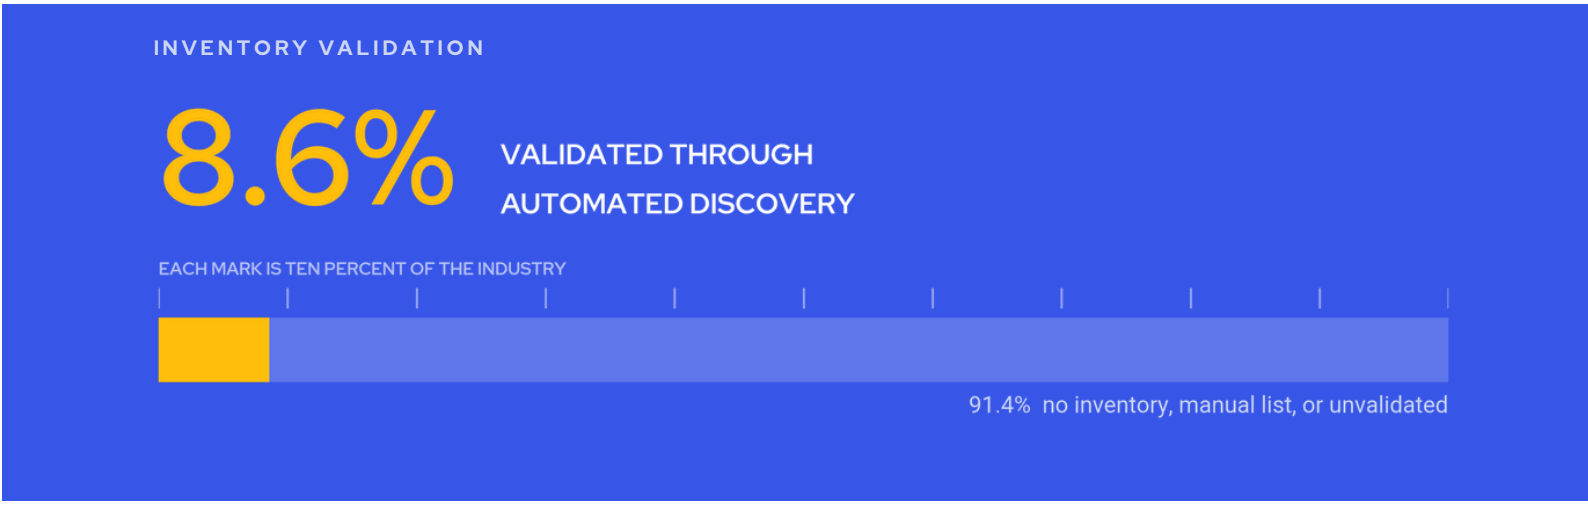
GAP ONE — THE INVENTORY

Only 8.6% know what AI is actually running.

Most organizations can tell you how many AI tools they believe they have. Only 8.6% validate that count through automated discovery.

The remaining 91.4% rely on a manual list, an unvalidated inventory, or no inventory at all. That leaves a basic governance question unanswered: does the AI recorded in policy and oversight processes match the AI actually operating in the environment?

The gap matters because AI no longer enters the organization only through a deliberate purchase or approved deployment. It can arrive embedded in software already licensed, enabled through a routine vendor update, added through a browser extension, or built into an existing workflow by an internal team. None of those necessarily look like 'AI adoption' when they happen.



WHY IT MATTERS

30.8% name unapproved employee AI use as their greatest AI-related data risk.

42.3% name sensitive data exposure through AI tools as a leading concern.

TAKEAWAY An inventory of approved AI is not an inventory of actual AI. Without automated discovery, most organizations cannot verify whether the AI they govern is the AI they actually have.

CLEARWATER'S INSIGHT

When we compare automated discovery with a client's approved AI list, the two rarely match. The issue is usually not that the approved list was wrong. It is that the environment changed faster than the list did.

Discovery often surfaces AI capabilities embedded in platforms the organization has used for years, along with individual accounts and tools that never entered a formal approval process. That makes inventory less of a documentation exercise and more of a continuous monitoring function.

The goal is not simply to produce a more accurate list once. It is to maintain enough visibility to see when the environment changes.

GAP TWO – THE DATA

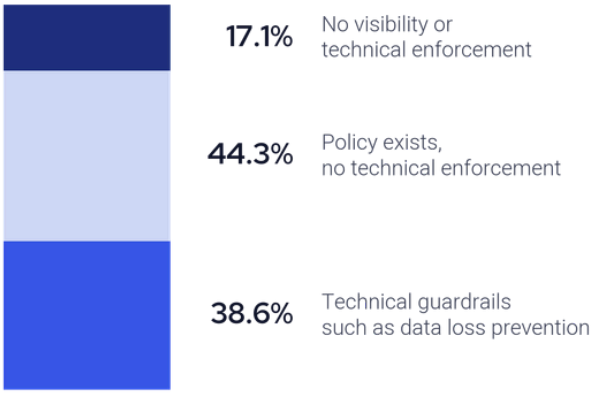
The most common control is a document.

Nearly half rely on policy alone. Only 38.6% enforce data protection technically.

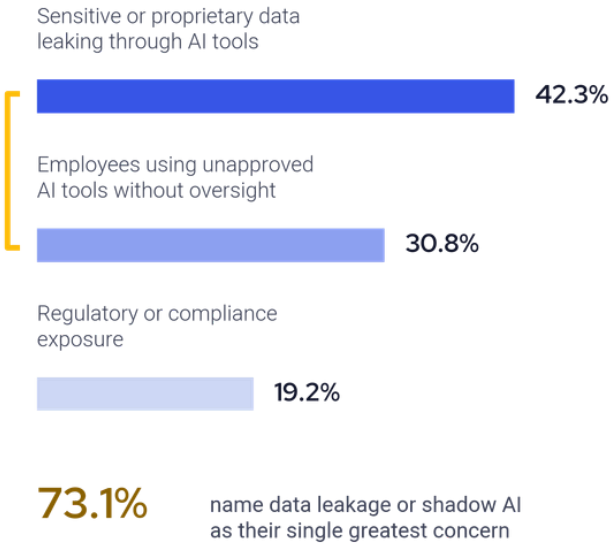
Sensitive data leaking through AI tools is respondents' single largest internal concern at 42.3%, followed by unapproved employee AI use at 30.8%. The most common control is a document: 44.3% report a policy with no technical enforcement, while only 38.6% use technical guardrails such as data loss prevention.

A policy can establish intent and accountability, but it cannot detect an upload, block a paste, inspect a prompt, or show what's left after the fact.

CONTROL POSTURE



GREATEST INTERNAL CONCERN



73.1% name data leakage or shadow AI as their single greatest concern

> A, B, and C, all three, keep me up at night.
>
> *Anonymous survey respondent*

TAKEAWAY

Healthcare organizations recognize the risk of sensitive data entering AI, but most still cannot technically prevent, detect, or evidence that movement.

GAP THREE – THE VENDOR

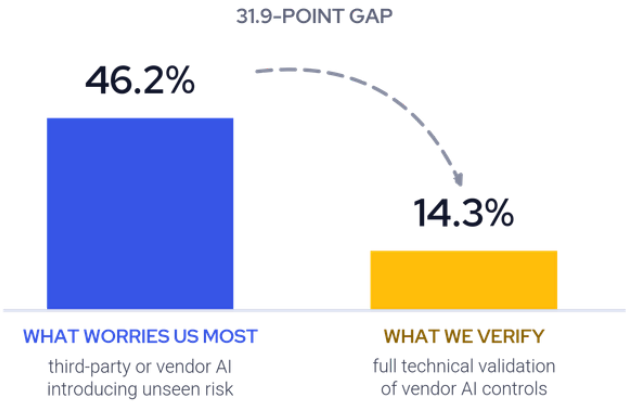
A questionnaire is not evidence.

Most organizations inherit AI from third parties, but only 14.3% fully validate those controls

Most healthcare organizations inherit a meaningful amount of AI from third parties, whether they intend to or not. AI introduced through third parties is the leading external concern at 46.2%, ahead of AI-powered attacks at 34.6% and reputational damage from a public AI failure at 11.5%.

Yet only 14.3% perform full technical validation of third-party AI controls. The gap is material: organizations remain accountable for risk introduced through software, services, and updates they may not recognize as new AI adoption.

CONCERN VERSUS VERIFICATION



ALSO NAMED AS THE GREATEST EXTERNAL CONCERN

AI-powered attacks (deepfakes, AI-enhanced phishing)	**34.6%**
Reputational damage from a public AI failure	**11.5%**
None significant	**5.1%**

WHAT A QUESTIONNAIRE CANNOT CATCH

- AI functionality added after contract signature
- Subprocessors and external foundation models
- Prompt and output retention
- Agents granted standing access to clinical systems
- Model changes shipped as routine updates

TAKEAWAY

Most healthcare organizations inherit meaningful AI risk from third parties regardless of intent, but remain accountable for controls they have not technically validated.

CLEARWATER'S INSIGHT

A security questionnaire is a point-in-time attestation. The AI behind a vendor relationship is not a point-in-time fact.

Subprocessors change. Foundation models are swapped, updated, or fine-tuned. A routine product release can introduce a new AI feature that sends data to a different inference endpoint. None of that requires a new contract, so none of it reliably triggers a new review.

Closing this gap takes more than a better questionnaire. It requires a standing right to verify how vendor AI is actually handling data, supported by technical evidence and refreshed on a regular cadence rather than filed away at signature.

ECOSYSTEM VIEW

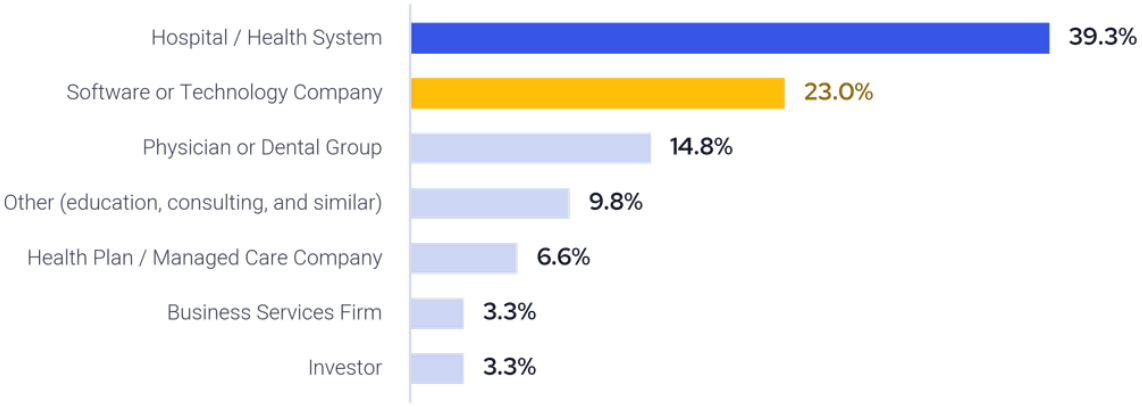
A shared ecosystem challenge.

Organizations building and delivering healthcare technology face similar AI supply-chain visibility and assurance challenges as their clients.

Healthcare's greatest external AI concern is third-party risk it cannot see. This study's respondent base reflects that ecosystem directly: software and technology companies, organizations that build or deliver healthcare software and technology services, made up 23.0% of respondents who identified their organization type, the second-largest group after hospitals and health systems at 39.3%.

Many inherit AI from foundation-model providers, cloud platforms, subprocessors, and features introduced through routine product updates, just as their healthcare clients inherit AI through the solutions they buy. Existing healthcare AI governance guidance has focused more on care delivery organizations than on the companies that enable their AI capabilities, leaving both sides without sufficiently specific validation practices.

WHO ANSWERED THIS STUDY, BY ORGANIZATION TYPE



TAKEAWAY

Healthcare organizations and their technology ecosystem partners are both inheriting third-party AI, but neither side has consistently moved from contractual assurance to technical validation.

CLEARWATER'S INSIGHT

When we assess AI-enabled healthcare technology companies directly, we often find that the underlying security program is more mature than the AI-specific controls layered on top of it. It is the same maturity gap we see on the client side, just one level up the supply chain.

Healthcare organizations often assume a vendor's general security certification also covers its AI-specific practices: how prompts and outputs are retained, which subprocessors touch model inputs, whether a foundation-model change gets reviewed before it ships.

Too often that assumption is untested on both sides of the contract.

INVESTMENT

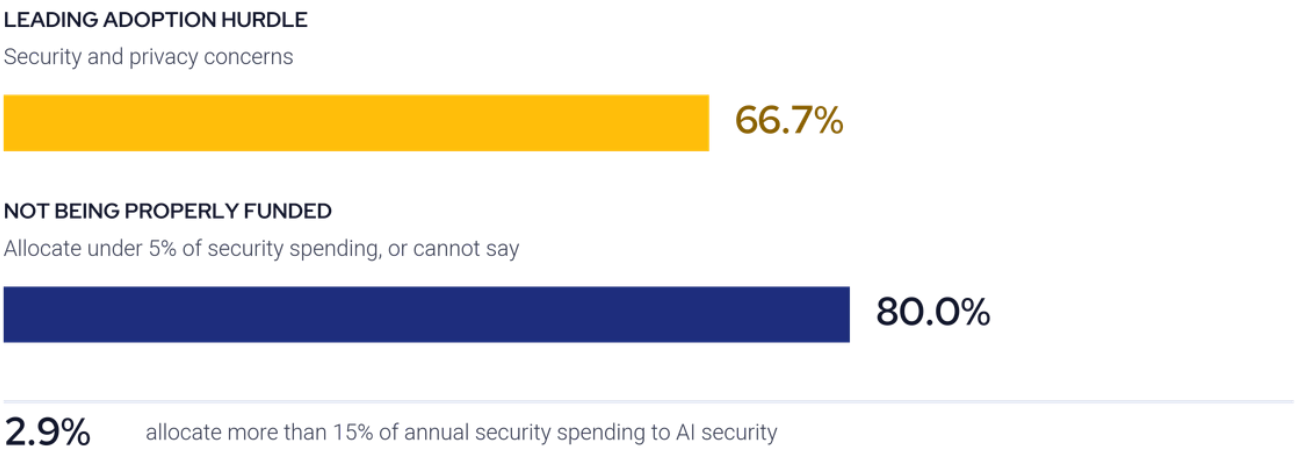
Security is the barrier, and the thing least funded.

Security and privacy are the leading barrier to AI adoption, cited by 66.7%.

Security and privacy concerns are the leading hurdle to AI adoption at 66.7%, well ahead of other factors. The result links adoption directly to assurance: organizations are not only asking whether AI can create value, but whether they can see its footprint, control data movement, and validate the vendors behind it.

AI security is not funded like the leading adoption barrier. A majority, 51.4%, allocate less than 5% of annual security spending to AI security, and 28.6% do not know what they allocate. The gap appears to have two causes: organizations lack guidance and a clear execution plan, while AI investment remains focused on adoption and new capability rather than protection.

CONCERN AND SPENDING, SIDE BY SIDE



2.9% allocate more than 15% of annual security spending to AI security

TAKEAWAY

Security is the leading barrier to AI adoption, but most organizations fund it below 5% of security spending or cannot state the allocation.

CLEARWATER'S INSIGHT

The gap between naming security as the top barrier and funding it accordingly comes from two places.

First, most organizations lack a specific plan of execution for AI security. The budget conversation stalls because no one has translated 'we're concerned about this' into a scoped set of controls with a cost attached.

Second, AI investment to date has been aimed almost entirely at adoption: licensing, integration, and unlocking new capability. Protection has been treated as a line item to revisit later rather than a cost of deployment.

Closing that gap does not require rebuilding the security budget. It requires a defined AI security workstream with its own scope, ownership, and funding rather than leaving it buried inside the broader security program.

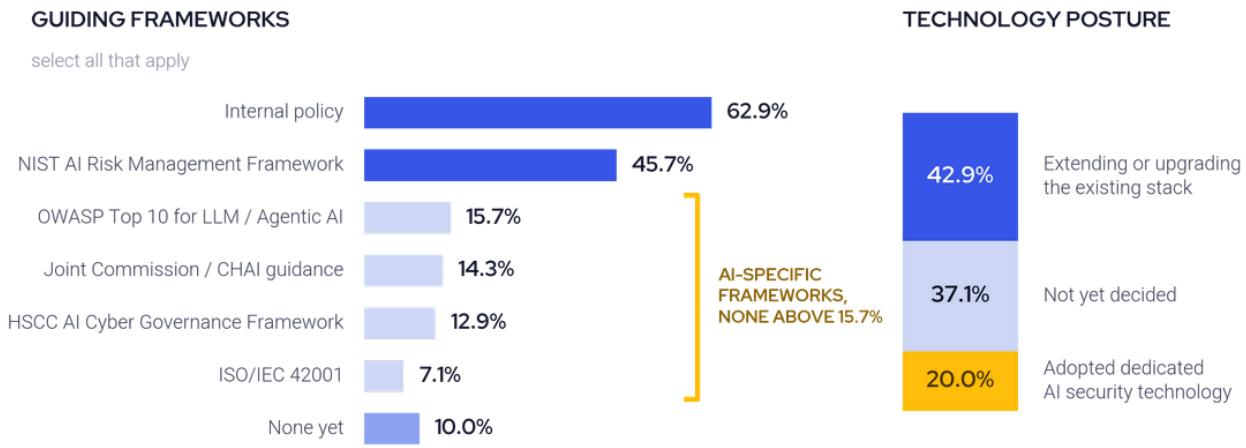
FRAMEWORKS & TOOLS

The frameworks nobody is using, and the tools that match.

Organizations rely on internal policy and broad risk guidance more than AI security-specific frameworks and technology.

Two thirds of respondents, 62.9%, of respondents point to internal policy as a guiding framework. The NIST AI Risk Management Framework follows at 45.7%, but it addresses AI risk and trustworthiness more broadly than security. The pattern leaves organizations with governance language but limited guidance for discovery, technical enforcement, vendor validation, monitoring, and evidence production.

The tooling picture matches. While 42.9% are extending or upgrading their existing security stack, only 20.0% have adopted dedicated AI security technology, and 37.1% have not decided. Most organizations are still determining whether existing tools can produce the assurance their AI programs require.



JUNE 2026

On June 1, 2026, the Joint Commission launched Responsible Use of Artificial Intelligence in Healthcare certification, recognizing organizations that can demonstrate governance, safeguards, monitoring processes, and education for AI use. It is built on the guidance released with the Coalition for Health AI in September 2025. The certification asks organizations to demonstrate. That is the same word this study has been measuring.

TAKEAWAY

The frameworks and the tools tell the same story. Neither is specific to AI security, and voluntary guidance in healthcare has a way of becoming the expectation.

CLEARWATER'S INSIGHT

Healthcare organizations should not wait for one framework to emerge as the definitive standard. The priority is to choose a credible, security-specific AI framework now and use it to build a control structure the organization can actually operate.

That means translating guidance into defined controls, assigning ownership, testing whether those controls work, and revisiting them as models, vendors, data flows, and use cases change.

Framework adoption alone does not create assurance. The value comes from putting the controls into practice and operating them continuously.

The organizations that move first on this will be better positioned as today's guidance becomes tomorrow's expectation.

CONCLUSION

The standard is changing.

1 **Treat AI as an operating environment rather than a technology project.**

Seven in ten respondents in this study are past the pilot stage. The controls have not caught up.

2 **Extend governance past policy and committee structure into enforceable technical control.**

Nearly half rely on policy with no technical enforcement behind it.

3 **Measure maturity by evidence.**

Discovery, enforcement, validation, and containment either produce artifacts or they do not exist.

TAKEAWAY

Healthcare leaders have spent three years answering whether they have an AI governance program. The question arriving now, from accreditation bodies and from clients, is whether they can prove it works.

91.4%

of healthcare organizations still cannot verify the AI running in their environment

THE 2026 BASELINE – WE WILL MEASURE IT AGAIN NEXT YEAR

CLEARWATER'S INSIGHT

For boards and executive leadership, the near-term priority is: fund AI security as its own workstream, not as a subset of the broader security budget or an afterthought to the AI initiatives it is meant to protect.

Ask your teams the three questions in this report's self-assessment (next page) and expect an evidence-based answer, not a policy citation.

For security, compliance, and technology teams already doing the work, the priority is to move from describing controls to proving they work: continuous discovery instead of an annual inventory, evidence instead of attestation from AI vendors, and agent governance before the next agent goes into production, not after.

None of this requires waiting for a mandate. The organizations furthest ahead in this study are the ones treating 2026 as the year to build the muscle, rather than waiting for regulation to force it.

SELF-ASSESSMENT

Score your own gap.

Three questions. The same three we asked the industry.

Answer the three main questions this study asked. Score two points where you can produce evidence, one where you have a policy or process, zero where you have neither. Six points means you can demonstrate what you report. Three or below puts you with the majority of the industry, and the industry rate beside each row shows how few organizations can score the full two.

DETACH AND SCORE

	EVIDENCE (2)	POLICY OR PROCESS (1)	NEITHER (0)	INDUSTRY RATE
Our AI inventory is verified against what is actually running	☐ Discovery output, dated	☐ A maintained list	☐ No inventory	8.6%
Data movement into AI services is controlled	☐ Enforcement logs	☐ A written policy	☐ No visibility	38.6%
Vendor AI controls are verified	☐ Test or validation results	☐ A completed questionnaire	☐ Contract language only	14.3%

YOUR TOTAL

_____

out of 6

WHERE THAT PUTS YOU

0	1	2	3	4	5	6

0 to 3 with the majority of the industry

4 to 6 you can demonstrate what you report

ORGANIZATION _____

ASSESSED BY _____

DATE _____

TAKEAWAY

Every row scores two points only where an artifact exists. That is the same standard accreditation bodies and clients are beginning to apply.

HOW CLEARWATER CAN HELP

Close your own gap.

FIVE STEPS TO GET STARTED

1. **Charter the committee.** Name an executive sponsor and adopt a charter that sets the mandate, decision rights, and meeting cadence.
2. **Set your guardrails.** Approve an AI policy and define risk appetite and tolerance, acceptable versus unacceptable uses.
3. **See what you have.** Stand up AI intake and inventory: discover, catalog, and risk-tier every system, use case, model, and vendor.
4. **Assign ownership.** Map the ten governance functions to named owners so nothing falls through the cracks.
5. **Enable your people.** Launch responsible-AI training and a simple intake path so employees know how to use AI and where their responsibilities lie.

AT THE CENTER OF THE CONVERSATION

Clearwater can help you at every step in your AI journey. We have been at the center of the AI security and governance conversation for the past few years.

- Contributing to Information Security Media Group's reporting on the advancement of generative AI in healthcare in early 2024: *Healthcare Edition: Generative AI Survey Report*
- Authoring the toolbook *AI Governance Strategy & Alignment: Empowering Decision Making* for The Governance Institute in April 2024
- Bringing together our experts and other industry leaders for virtual summits on Responsible AI in Healthcare in the summer of 2025 and 2026

WE HELP CLIENTS DISCOVER, GOVERN, VALIDATE, AND SECURE THE AI ALREADY RUNNING IN THEIR ENVIRONMENT

AI Discovery and Management

"What AI is already in my environment, and what is my team putting into it?"

Technology-enabled discovery of AI services, data access, and system integration.

Integrated AI Assessment

"I want to understand my AI security posture and risk areas."

Integrated assessment to help you understand your posture, align with AI frameworks, and prepare for upcoming regulations.

AI Governance Program Development

"I want to build foundational AI governance but don't know where to start."

Flexible delivery: a toolkit to DIY your own program, white-glove service to build it for you, or ongoing program support where Clearwater runs the cadence.

Readiness Diagnostics Against AI-Powered Threats

"I want to be prepared for a cyberattack powered by a frontier AI model."

Self-paced or Clearwater-led diagnostics against your readiness to defend against AI-powered threats.

AI Advisory

"I want general guidance on AI security, privacy, and compliance as I mature."

An AI advisor you can trust, who understands your environment, whenever you need one.

TALK TO CLEARWATER

Schedule a 30-minute call with our team to discuss where your organization stands and how to move forward.

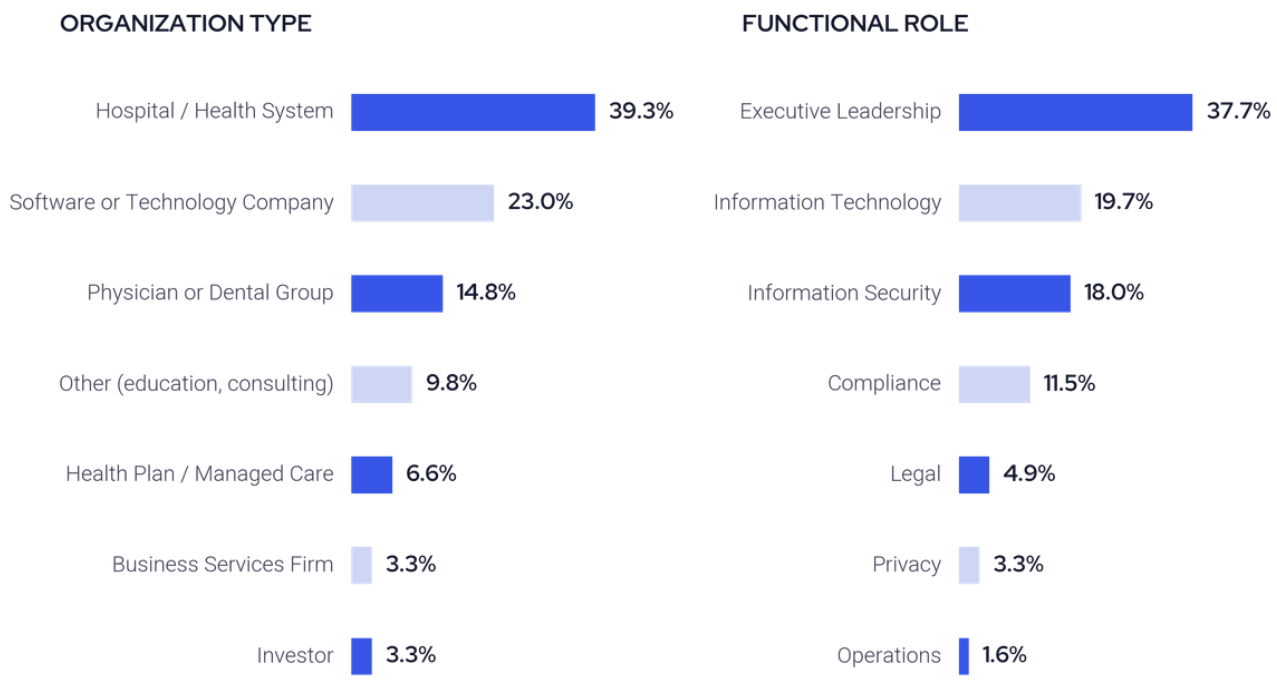
www.ClearwaterSecurity.com | 800.704.3394 | info@clearwatersecurity.com

METHODOLOGY

Respondent breakdown

Responses were collected from late July through early September 2026.

Respondents included hospitals and health systems, software and technology companies, physician and dental groups, health plans and managed care organizations, academic medical and dental education programs, business services firms, and investors. Individual responses are confidential and no respondent ororganization is identified in this report.



Results represent a directional industry benchmark rather than a statistically representative measure of the healthcare sector.