Vitea

AI Governance for Healthcare

Where Policy Ends

The State of AI Governance in Healthcare, 2026

What health system leaders report, what they say when the room is candid, and what the network actually shows.

Featuring findings from a CHIME Foundation member survey fielded for Vitea, a CHIME member focus group of senior IT, informatics and clinical leaders, observed data from Vitea Lens assessments, and the year's leading independent research.

Inside this report

Contents

	Foreword	3
	Key findings at a glance	4
	About this research	6
01	AI Arrived Before the Org Chart Did	7
02	Governance on Paper	10
03	Detect: You Can't Govern What You Can't See	12
04	Control: Policy Without a Point of Control	15
05	Monitor: Approval Is a Moment. Risk Is Continuous.	18
06	The Cost of an Ungoverned Moment	22
07	The Rulebook Is Moving In-House	26
08	The Priority Paradox	29
09	From Policy to Practice: What Good Looks Like	31
	The longer horizon	34
	Sources	35
	About Vitea	37

Foreword

Governance has to live where AI is used

Before Vitea, I co-founded Jvion, a clinical AI company that helped hospitals identify patients at risk and act before avoidable harm occurred. That experience showed me what AI could make possible in healthcare. Later, after starting Seed2B Capital and working with AI companies across the fund's portfolio, I began to see the gaps in how these technologies were governed. As AI adoption expanded, so did the questions: Who had approved its use? Who was monitoring its performance? Who would step in when something went wrong?

Those questions have become more urgent as the ways AI enters a health system have multiplied. It arrives in a vendor's quarterly release, inside a browser extension, or through a personal account used at work. In our assessments, we see how wide the gap between expected and actual use can be. Leaders expect dozens of AI tools. They find hundreds.

None of this reflects a lack of effort. Health systems have stood up AI councils, written policies and built intake reviews faster than almost any other industry. What they haven't had is a way to carry those decisions out of the meeting room and into the moment someone actually uses AI: the prompt, the upload, the drafted note, the agent acting on a chart.

That gap is what this report is about. We brought together three sources of our own: a CHIME Foundation member survey of senior technology leaders, a candid focus group with CHIME members, and observed network data from Vitea Lens assessments. Then we set them against the strongest independent research published this year.

The findings are sobering, and they are also useful. Healthcare has already decided to use AI. The open question is whether it will be governed where it is used, every time, or only on paper. I hope this report helps your next governance conversation start in the right place.

Shantanu Nigam
Co-founder and CEO, Vitea

At a glance

Key findings

Health systems have built the structures of AI governance. The AI itself now moves faster than those structures can see, enforce or verify, and the accountability for closing that gap is landing on the health system.

81%

of U.S. physicians now use AI in their practice, more than double the 2023 rate.

American Medical Association, 2026

160–312

AI-enabled applications observed at each health system in Vitea Lens assessments.

Vitea Lens, 2026

6 of 7

health systems that call their AI governance "established" cannot enforce AI policy in real time.

CHIME Foundation member survey fielded for Vitea, 2026

85%

of AI upload volume at one system went to apps its existing security tools did not categorize as AI.

Vitea Lens, 2026

0 of 15

surveyed organizations have automated, continuous monitoring of AI after deployment.

CHIME Foundation member survey fielded for Vitea, 2026

43%

of breached organizations reported a shadow AI incident in 2026, up from 20% a year earlier.

IBM Cost of a Data Breach Report, 2026

#1

AI's rank on both of ECRI's 2026 lists: health technology hazards and patient safety concerns.

ECRI, 2026

Last 3

Detection, real-time control and monitoring rank lowest on 12-month priority lists, despite being the widest gaps.

CHIME Foundation member survey fielded for Vitea, 2026

What the evidence adds up to

Adoption is no longer a decision.
AI is embedded in the EHR, in vendor platforms and in the daily habits of clinicians and staff, across nearly every department.

Visibility is overestimated.
Even organizations with modern security tooling miss AI that isn't labeled as AI, AI inside approved apps and activity that can't be tied to a person.

The costs are now visible.
Class actions over ambient AI, breach costs that lead every industry and documented patient injuries have turned AI governance into a financial and clinical risk issue.

Governance exists on paper.
Committees, charters and intake reviews are common. Control at the point of use is rare.

Approval is treated as the finish line.
Nearly every system tests AI before go-live. Almost none watch it continuously afterward.

Accountability is moving in-house.
Federal oversight is loosening while states, accreditors and plaintiffs' attorneys tighten, leaving the deploying health system to carry the burden of proof.

How we use three key terms

- **Detect** — Continuous discovery of AI in the environment: sanctioned tools, AI embedded in approved software, extensions and plugins, consumer apps and agents, tied to the people and data involved.
- **Control** — Enforcement of policy at the moment of use, applied automatically and consistently rather than interpreted after the fact.
- **Monitor** — Ongoing assurance that approved AI keeps behaving as approved, including accuracy, drift, policy violations and vendor changes.

Reading guide by role

Each chapter stands on its own. If you lead one of these functions, start here.

Role	Your first question	Start with
CIO / CTO	Do we have a trusted, current inventory of every AI capability in our environment?	Chapters 1, 3 and 9
CISO	Can we see and stop PHI moving to AI tools, including apps not labeled as AI and personal accounts?	Chapters 3, 4 and 6
Chief Compliance / Risk Officer	Could we produce audit-ready evidence of AI oversight for a surveyor, regulator or court within 30 days?	Chapters 6 and 7
CMIO / CNIO	How would we know if an approved clinical AI tool started producing worse output?	Chapters 5 and 6

Methodology

About this research

This report combines three proprietary sources with independent research. The survey shows what leaders report. The focus group shows what they say when the conversation is candid. Lens assessments show what is actually happening on the network.

CHIME Foundation member survey, fielded for Vitea

The CHIME Foundation fielded a survey built from Vitea's questions to its member base in 2026. After removing test entries and incomplete responses, 15 complete responses remained from CIO/CTO-level leaders (8), CISOs, CMIOs and CHIOs, and senior IT leaders. Organizations ranged from small hospitals to 1,000-plus-bed enterprise systems. Fifteen responses offer a directional signal from a senior, on-target audience, not a statistically projectable market estimate. We report counts rather than percentages wherever that is the more honest framing.

CHIME member focus group

On September 3, 2026, Vitea convened senior IT, informatics and clinical leaders from a range of healthcare organizations for a candid discussion of AI governance in practice. All comments were anonymized to keep the conversation candid, and they are quoted here without attribution.

Vitea Lens assessments

Observed data comes from three Vitea Lens assessments at multi-hospital U.S. health systems, presented here as Systems A, B and C. Observation windows ranged from roughly 33 to 110 days; System C's figures reflect a sample of about 65% of network logs. Applications flagged as unapproved were flagged only because approval and contract status had not yet been validated, which is a visibility gap, not evidence of misuse. Identifying details have been removed and some figures rounded.

Independent research

External data was published in 2026, with a small number of late-2025 sources used for regulatory context and one hospital survey released in November 2025. Sources are linked where they are cited and listed in full at the end of the report. We excluded research produced by healthcare AI governance vendors.

Chapter 1

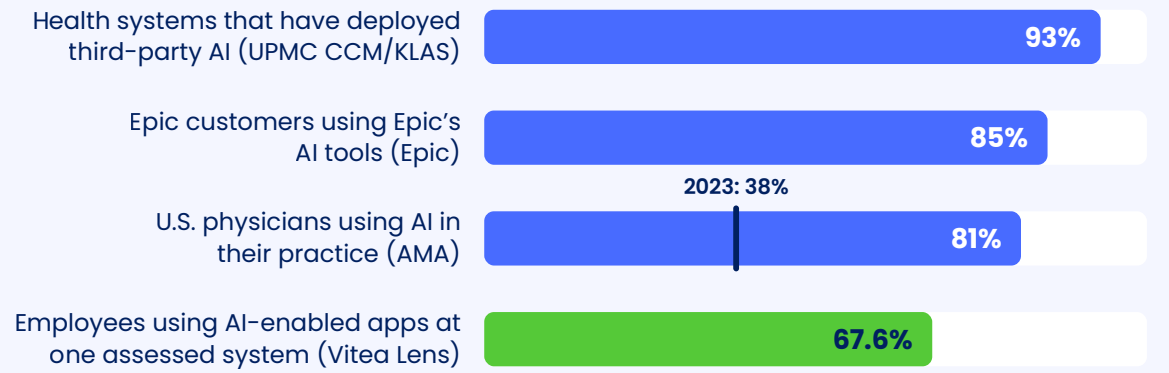
01 AI Arrived Before the Org Chart Did

In most health systems, the question of whether to adopt AI was settled without a vote. It happened one vendor update, one browser extension and one helpful colleague at a time.

Figure 1

AI is already in the building

Share of each group using AI, 2026. Different sources and denominators; shown together for scale.



Sources: UPMC Center for Connected Medicine & KLAS Research (2026); Epic (HIMSS26, 2026); American Medical Association (2026); Vitea Lens assessment, System B.

The American Medical Association's 2026 physician survey found that 81% of physicians now use AI in their practices, up from 38% in 2023, and the average physician uses it for 2.3 tasks instead of 1.1. Physicians also want a seat at the table: 85% want to be consulted or directly involved in decisions about adopting AI.

The platforms clinicians live in are moving just as quickly. Epic reported at HIMSS26 that more than 85% of its customers now use its AI tools, and by its August user group meeting it was offering roughly 120 AI features out of the box. The EHR has become an AI platform in its own right. Third-party AI is just as common: in a 2026

study from the Center for Connected Medicine at UPMC and KLAS Research, 93% of health systems had deployed third-party AI solutions, with clinical documentation the most common use case.

What the network shows

Surveys tell you what people believe is happening. Network data shows what is actually happening. Across Vitea Lens assessments, every system had far more AI in use than its existing inventory reflected.

	System A	System B	System C
Observation window	90 days	110 days	~33 days (~65% of logs)
AI-enabled apps observed	**160+**	**312**	**236+**
Scale of use	3,000+ users; 150,000+ AI-related events	67.6% of nearly 10,000 employees; 2,300+ habitual users	Tens of thousands of employees using AI-enabled apps
Data movement	~30 GB uploaded to AI apps across ~35,000 files	AI activity in 645 of 692 departments	Clinical, productivity and general-purpose workflows

Source: Vitea Lens assessments, 2026. Observed data, not modeled.

Pilots don't reach 645 of 692 departments. At System B, two-thirds of the workforce used an AI-enabled app during the observation window. This is the operating environment, and governance has to start from that fact.

AI is already inside clinical work

Some of the most important activity sits in clinical workflows. At System B, more than 1,700 employees in clinical functions used a general-purpose AI assistant. A clinical evidence tool reached 428 employees and accounted for 58% of active AI users in residency and medical education. Every one of the system's 11 interpreters was a habitual AI user, with a single machine translation tool accounting for about 85% of their activity. At System A, 74 clinical staff generated more than 35,000 events in a writing assistant, a volume consistent with documentation work.

Nothing here describes reckless behavior. Clinicians and staff found tools that help them do their jobs. The gap is that nobody could see these workflows clearly enough to ask the questions that matter. Is there a business associate agreement? Where does the data go? Is this translation tool accurate enough for a conversation about a patient's medication?

The same job, done by many vendors

Every system also showed heavy overlap. At System A, four use cases were each served by three to five vendors, including five conversational AI tools with more than 3,000 users between them. System B had six clinical documentation products in use at the same time. System C had four categories, from clinical scribes to image tools, where multiple products performed the same function.

Each overlapping tool is another contract to review, another data flow to map and another set of guardrails to maintain. It is also, as Chapter 9 explains, one of the fastest places to recover both money and control.

Chapter 2

Governance on Paper

02

Most health systems can produce an AI governance charter on request. Fewer can tell you what that charter did last Tuesday.

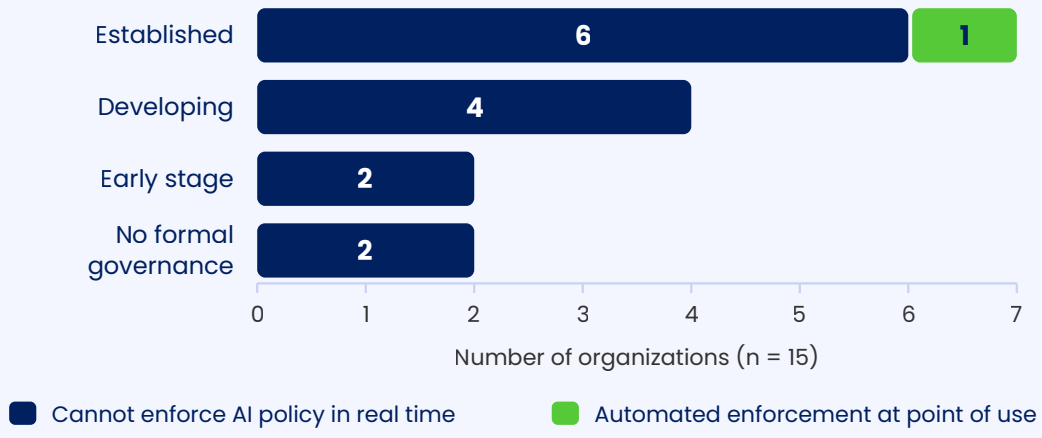
The structures are real. In its February 2026 response to an HHS request for information, the American Hospital Association reported that 74% of hospitals have multiple teams responsible for evaluating predictive AI. The Joint Commission and the Coalition for Health AI recommended formal governance structures in their initial guidance in September 2025, and committees of that kind are now common.

In the CHIME Foundation member survey fielded for Vitea, 7 of 15 organizations described their AI governance as “established,” meaning formal governance with technical controls that is actively monitored. Four more said “developing.” On the surface, that reads like a maturing market.

Figure 2

The maturity illusion

Self-rated AI governance maturity vs. ability to enforce policy in real time. CHIME Foundation member survey fielded for Vitea, 2026.



Source: CHIME Foundation member survey fielded for Vitea, 2026 (n = 15).

The maturity illusion

Then we asked how they enforce policy. Of the seven organizations with established programs, six cannot enforce AI policy in real time. Across all 15 respondents, exactly one has automated enforcement at the point where someone interacts with AI.

The disconnect has a simple cause. As most organizations measure it, maturity means committees, policies and periodic audits. Those things are necessary. They are also slow. A council that meets monthly is governing a portfolio that, in our assessments, changes every week.

Independent research points the same way. In the CCM and KLAS study, 63% of health systems described their AI strategy as developing or ad hoc, and only one leader of the 27 interviewed described their organization's strategy as advanced.

Who owns it?

Ownership is just as uneven. In the CHIME survey, 5 of 15 organizations have no single owner of AI governance; responsibility is informal or spread across a committee. Where there is an owner, it is usually the CIO or CTO, sometimes shared with the CMIO.

There is also a gap between the people who write AI policy and the people who use AI most. Wolters Kluwer Health's 2026 shadow AI survey found administrators were three times more likely than providers to be involved in developing AI policy (30% vs. 9%). Only 29% of providers said they were aware of their organization's main AI policies.

> ***"We talked about feeding our policies into an AI model and having it enforce our policies… but our policies just don't reflect the necessary granularity. They're too generic, and they're written for manual processes."***
>
> CHIME member focus group participant, September 2026

That comment captures the core problem. Policies written for people to interpret later can't be applied by a system in the moment, which is where AI risk actually occurs. The next three chapters look at the three places that gap shows up: what organizations can see, what they can control and what they can verify over time.

Chapter 3 Detect

You Can't Govern What You Can't See

03

Every governance program starts with an inventory. In most health systems, that inventory is wrong, and not by a small margin.

Shadow AI is now routine. In Wolters Kluwer Health's survey of 518 healthcare professionals, 40% had encountered unauthorized AI tools at work, nearly 20% had used one, and one in ten had used one for direct patient care. The most common reason was speed: half pointed to faster workflows. A May 2026 survey of 1,823 clinicians in 25 countries by Heidi found 83% were using AI without guidance from their employer, a formal policy or a recommended tool.

The financial consequences are growing with it. In IBM's 2026 Cost of a Data Breach Report, the share of breached organizations reporting a security incident involving shadow AI more than doubled, from 20% to 43%. Meanwhile, only 38% of breached organizations required IT approval before AI was deployed, down from 45% the year before.

Health system leaders know this. In the CHIME survey, only 1 of 15 organizations has end-to-end, real-time detection and alerting for unsanctioned AI. Five described their visibility as good, based on network or policy-level monitoring. Five have partial visibility with significant gaps, three rely on informal discovery and one has no visibility at all.

Even good tools have blind spots

The five organizations reporting good visibility aren't wrong. They see what their tools are built to see. What Lens assessments show is how much sits outside that view.

System A had a cloud access security broker with a generative AI category. Lens found 64 AI-enabled apps outside that category, used by 1,350 people across more than 50,000 events. Thirty-three of them were rated critical or high risk. More striking

was where the data went: 85% of AI upload volume, about 25.6 GB, flowed to apps the existing tooling didn't classify as AI. A single clinical vendor received 7.4 GB, mostly PDFs, from 70 users.

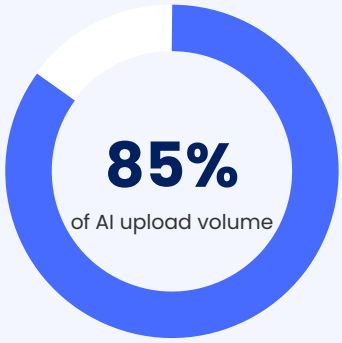
Data-loss rules keyed to a "GenAI" category never fired for that traffic, because the destination wasn't in the category.

Figure 3

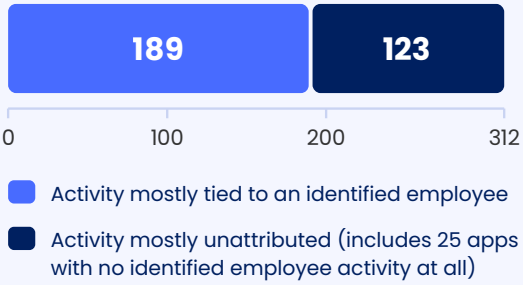
What existing tools didn't see

Observed Vitea Lens data from two assessed health systems

System A: uploads to apps the existing CASB did not categorize as AI



System B: 312 AI-enabled apps by identity attribution



Source: Vitea Lens assessments, Systems A and B, 2026.

Seeing an app isn't the same as seeing who uses it

At System B, most activity in 123 of 312 AI-enabled apps could not be tied to an identified employee, and 25 apps showed no identified employee activity at all. Six service accounts generated half of all unidentified firewall activity. Another 1,740 usernames didn't match the employee roster.

An app you can see but can't attribute is only half visible. You can block it. You can't coach the person using it, check whether their role should have access, or tell whether a clinician, a contractor or an automated process sent the data.

Personal account, enterprise network

Identity gets harder when the same product comes in two versions. At System B, 57.5% of Microsoft 365 Copilot users also used the consumer version of Copilot. On the network, a personal AI account and an enterprise tenant can look nearly identical. One comes with a business associate agreement, retention controls and audit logs. The other comes with a consumer privacy policy. Without identity-aware detection, there is no reliable way to tell which one just received a discharge summary.

AI inside the tools you already approved

The blind spot that worried our focus group most wasn't a new app. It was AI arriving inside software they had already reviewed.

> *"Those vendors are doing upgrades… behind the scenes, they're adding in some of that technology, and there's not an outright notification that we've now got this tool embedded that's AI."*
>
> CHIME member focus group participant, September 2026

Lens data shows the scale. At System A, more than 480 users generated over 13,000 events from AI plugins and extensions running inside apps that were already approved. At System C, embedded AI in enterprise apps, browser extensions and vendor platforms was a major finding in its own right, because these experiences rarely appear in a traditional software inventory. Each new feature lands in a product that already cleared review, so no review is triggered.

Detect: questions to ask your team

1. Can we name every AI capability running today, including features vendors added after we approved their product?
2. Does our discovery depend on a vendor's category label, or on what the traffic actually is?
3. For each AI app, can we tie activity to a named person, role and department, and tell personal accounts from enterprise ones?

Chapter 4 Control

Policy Without a Point of Control

04

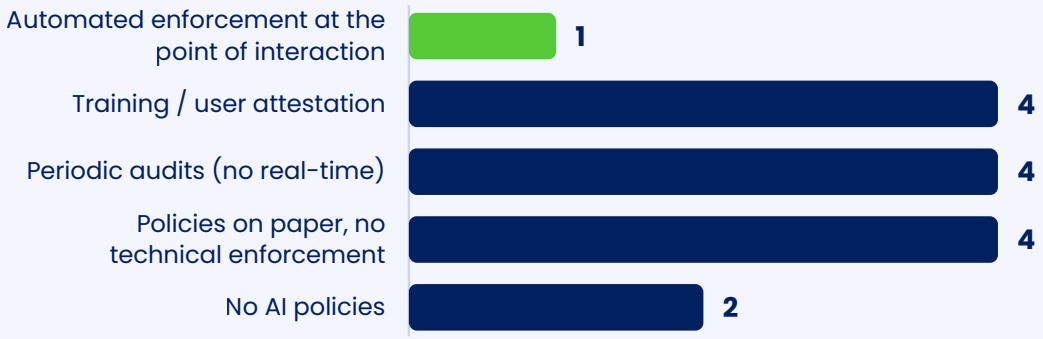
A policy that relies on people remembering it at the busiest moment of their shift isn't really a control. In most health systems, that is still how AI policy works.

The CHIME survey asked how organizations primarily enforce AI policy. Four rely on training and user attestation. Four run periodic audits and cannot enforce in real time. Four have policies on paper with no technical enforcement. Two have no AI policies. One has automated enforcement at the point of interaction.

Figure 4

How health systems enforce AI policy today

Primary enforcement approach, number of organizations (n = 15). CHIME Foundation member survey fielded for Vitea, 2026.



Source: CHIME Foundation member survey fielded for Vitea, 2026 (n = 15).

Put differently, 12 of 15 organizations rely on methods that act before the moment of use, long after it, or not at all. Training happens months earlier. Audits happen months later. The prompt, the upload or the agent's action happens in between, and that is where the risk lives.

Training alone can't close that gap, and the data shows why. Only 29% of providers are aware of their organization's main AI policies. Health systems are also deploying faster than their own approval processes. An Imprivata survey of 250 U.S. healthcare leaders, conducted by Vanson Bourne, found close to three-quarters of organizations deploy AI tools or agents without formal IT approval at least some of the time.

Agents raise the stakes

A chatbot that drafts a reply waits for a person to send it. An agent that can read a chart, message a patient or queue a task may not wait at all.

Twelve of the 15 CHIME respondents already use agentic AI. Of those, 11 have only informal controls, rely on the vendor's guardrails or have no controls specific to agents. One has formal guardrails and monitoring in place. The Imprivata survey found more than a quarter of organizations already run agentic AI in production and another 44% are piloting it, yet only 17% believe their current identity approaches would work for agents without changes.

IBM's breach data shows where that leads. Among organizations that experienced an AI-related breach in 2026, 92% lacked proper AI access controls.

Making policy machine-applicable

Focus group participants kept returning to one idea: using AI to help govern AI. The obstacle is that most policies were written for people. A rule such as "don't share PHI with unapproved tools" is clear to a human reader and nearly useless to a system. Which tools? What counts as PHI inside a prompt? Should the system block, redact, warn or log? Who gets notified? Does the data ever leave the environment while the decision is made?

Machine-applicable policy answers those questions in advance, so the decision can be made in milliseconds and applied the same way across sanctioned tools, embedded features and consumer apps. A useful test for any AI policy is whether a system could apply it consistently to the next 10,000 interactions without a person interpreting it. Most can't yet.

Control: questions to ask your team

1. Which of our AI policies can be enforced automatically today, and which depend on someone remembering them?
2. What happens, technically, when someone tries to send PHI to an unapproved AI tool right now?
3. Does every AI agent have its own identity, scoped permissions and an auditable record of what it did?

Chapter 5 Monitor

05

Approval Is a Moment. Risk Is Continuous.

Every leader in our focus group had a review process for new AI tools. None of them believed that review said much about how a tool would behave six months later, after a model update, a new feature or a change in how staff use it.

The data bears that out. In the CCM and KLAS study, 92% of health systems test third-party AI before deployment, but only 44% have a dedicated platform or environment for that testing. In the CHIME survey, not one organization has automated or continuous monitoring after deployment. Seven rely on end users to report problems, four run occasional manual audits and four have nothing in place.

Figure 5

Approval gets tested. What happens after rarely does.



First two bars: UPMC Center for Connected Medicine & KLAS Research (interviews, May–June 2026). Third bar: CHIME Foundation member survey fielded for Vitea (n = 15). Different samples, shown together to illustrate the pattern.

Sources: UPMC Center for Connected Medicine & KLAS Research (2026); CHIME Foundation member survey fielded for Vitea (2026).

That leaves drift, hallucinations and policy violations to be caught by whoever happens to notice. The patient safety community has noticed. ECRI named misuse of AI chatbots the top health technology hazard for 2026, citing chatbots that suggested incorrect diagnoses, recommended unnecessary tests and even invented body parts. Weeks later, ECRI and the Institute for Safe Medication Practices put the "AI diagnostic dilemma" at the top of their 2026 patient safety concerns, warning about automation bias and the erosion of clinicians' critical thinking.

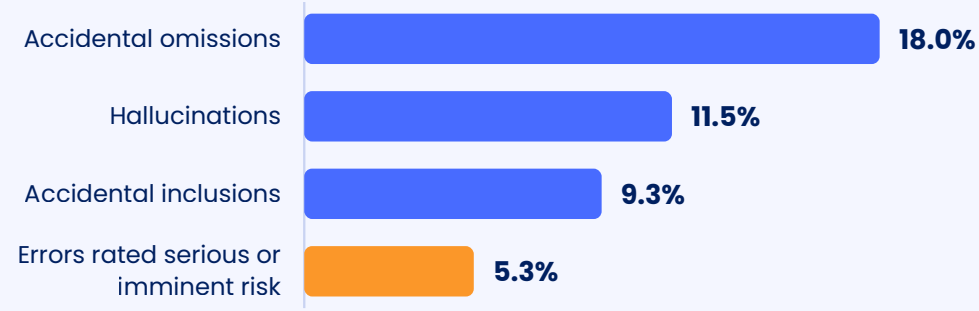
What a monitoring gap looks like in a chart

> *"It said vital signs were stable throughout the shift. It was a pre-op patient, and the only vital documented was weight. That's misleading. I don't want that in my note."*
>
> CHIME member focus group participant, September 2026

Figure 6

What physicians found when they checked AI-drafted notes

Share of 356 physician-reviewed ambient AI notes containing each error type. Prospective pilot study, published April 2026.



Source: Quality of Clinical Notes Created by Ambient Listening Generative AI: Pragmatic Prospective Pilot Study (2026).

That is an AI scribe inventing a reassuring detail, and it is exactly the kind of error that slips through when review depends on a busy clinician catching it. A prospective pilot study published in April 2026 measured how often it happens. Thirty-one physicians used an ambient AI scribe to create 7,545 notes and reviewed 356 of them. Omissions appeared in 18% of reviewed notes, hallucinations in 11.5% and

accidental inclusions in 9.3%. Most errors were mild to moderate, and 94.7% of notes were free of significant errors. But 5.3% contained errors rated as posing serious or imminent risk if left uncorrected.

Most notes were fine. That is precisely what makes this hard. When 95 notes in 100 are clean, reviewers stop looking hard at the other five. Clinicians know the risk: in Wolters Kluwer Health's 2026 Future Ready Healthcare survey, 74% named hallucinations as a major concern. Patients are counting on that vigilance, too; 78% expect their doctors to validate AI-derived information against a trusted source. Focus group participants added a provenance problem: users often can't see where an AI output came from, which makes checking it slower and less likely.

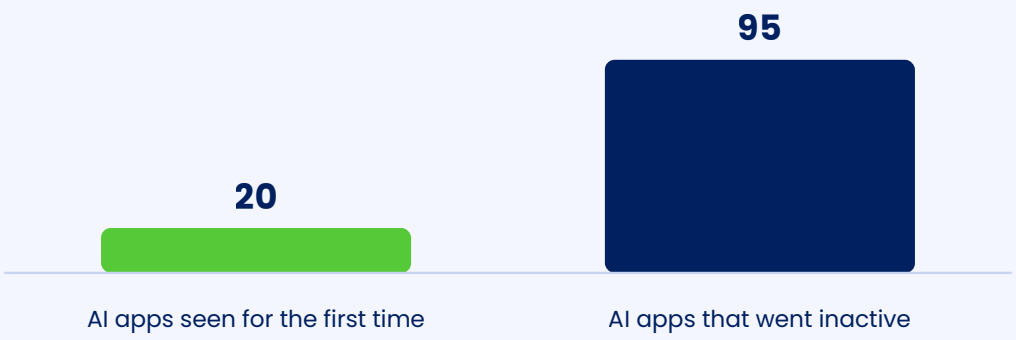
The portfolio never stands still

Even if every tool behaved perfectly, the set of tools keeps changing. In the final 28 days of System B's assessment, 20 AI apps appeared for the first time and 95 went quiet. Usage of one general-purpose assistant fell 43% in that window while another rose 25%. At System A, use of one major assistant was declining while clinical evidence tools grew.

Figure 7

Four weeks of change in one AI portfolio

System B, final 28 days of a 110-day Vitea Lens observation window



Source: Vitea Lens assessment, System B, 2026.

An inventory taken in June is a historical document by August. Timing and trend data don't prove misuse, but they do show where to renew, retire or re-review, and they make a one-time approval process look increasingly out of step.

Vendors change after you approve them

> *"A vendor has added some new feature that becomes available to all of our users without us really knowing that it includes an AI component… it's happening way faster than we're able to document and catch it in a contract renewal."*
>
> CHIME member focus group participant, September 2026

Contract renewals happen every year or three. Vendor releases happen every few weeks. Monitoring has to watch for vendor-side changes as closely as user behavior, because a product that was approved without AI can acquire it overnight.

Who owns testing that never ends?

Our focus group called this the least-settled question in the room. Security and informatics teams test tools today, but no one has a durable mandate for testing that must repeat every time a model changes. Until someone owns it, monitoring defaults to the end user, which is where seven of the 15 CHIME respondents are today.

Monitor: questions to ask your team

1. How would we know if an approved AI tool started performing worse next month?
2. Do our contracts require notice before a vendor adds or changes AI functionality, and would we detect it if they didn't?
3. Who owns repeat testing of AI tools after go-live, and what evidence could we hand an auditor today?

Chapter 6

06

The Cost of an Ungoverned Moment

AI incidents rarely announce themselves. They surface later, in a patient's chart, a plaintiff's complaint or a breach report, often months after the moment that caused them. By then the costs have stacked up in three places: money, legal exposure and people.

The financial cost

Healthcare has had the costliest data breaches of any industry for 13 consecutive years. In IBM's 2026 study, the average healthcare breach cost $6.64 million, and the average U.S. breach across industries reached $11.5 million. Breaches involving shadow AI averaged $5.39 million, and roughly one in five led to a regulatory fine. The governance trend is moving the wrong way: 68% of breached organizations had no policy to govern AI or manage shadow AI, up from 63% the year before.

The legal cost: consent is the new front line

The first wave of AI litigation against health systems isn't about bad algorithms. It is about whether patients knew AI was in the room.

Sharp HealthCare. In November 2025, a patient filed a proposed class action in San Diego Superior Court alleging that Sharp used an ambient AI documentation tool to record his visit without notice or all-party consent, as California law requires. The complaint alleges more than 100,000 patients were affected and that medical records contained boilerplate language saying patients had been advised of and consented to the recording when they had not. California's Invasion of Privacy Act allows $5,000 in statutory damages per violation, and the suit also seeks compensatory and punitive damages, corrections to medical records and an order blocking use of the tool without proper consent.

Sutter Health and MemorialCare. In April 2026, three patients filed a proposed class action in federal court alleging the two systems used the same category of ambient AI scribe without meaningful, informed consent. The complaint goes further than Sharp, applying the federal Electronic Communications Privacy Act to a proposed nationwide class. The vendor is not named as a defendant; the health systems are. According to Law360, at least three other patients have filed similar proposed class actions since the Sharp case.

Neither case has reached a public resolution as of this writing, and the allegations have not been proven. But the arithmetic explains why boards are paying attention. If a court counted each recorded visit as a violation, 100,000 encounters at $5,000 apiece would put theoretical statutory exposure at $500 million. That figure is a ceiling, not a forecast; how violations are counted in ambient AI cases is still unsettled. The legal theories themselves are not new, though. In late 2025, Kaiser Permanente agreed to pay up to $47.5 million to settle class action claims over tracking technologies on its websites and apps, brought under the same family of wiretap and privacy laws now aimed at ambient AI.

Matter	What's alleged or at issue	Financial exposure	Status
Sharp HealthCare Filed Nov. 2025	Ambient AI recorded visits without notice or all-party consent; consent language allegedly inserted into records	$5,000 per violation under CIPA; class alleged at 100,000+ patients	Pending
Sutter Health & MemorialCare Filed Apr. 2026	Ambient AI scribe used without informed consent; federal wiretap claims on behalf of a proposed nationwide class	Statutory damages under CIPA, CMIA and ECPA	Pending
Kaiser Permanente Settled late 2025	Tracking technologies shared member data with third parties (a precedent for the same legal theories)	$46M, up to $47.5M; ~13.4M members	Preliminary settlement
Healthcare breaches involving shadow AI IBM, 2026	Unapproved AI tools involved in security incidents	$5.39M average; ~1 in 5 drew a regulatory fine	43% of breached organizations

Sources: Fisher Phillips; Medical Xpress; TechTarget; EisnerAmper; BankInfoSecurity; IBM Cost of a Data Breach Report 2026. Allegations in pending cases are unproven.

Settlements are only the visible line item. Legal defense, re-consenting thousands of patients, correcting records, retraining clinicians and, potentially, pausing a tool that clinicians have come to depend on don't show up in court filings. Neither does the cost of the governance work that would have to happen afterward anyway.

The question every board should be able to answer

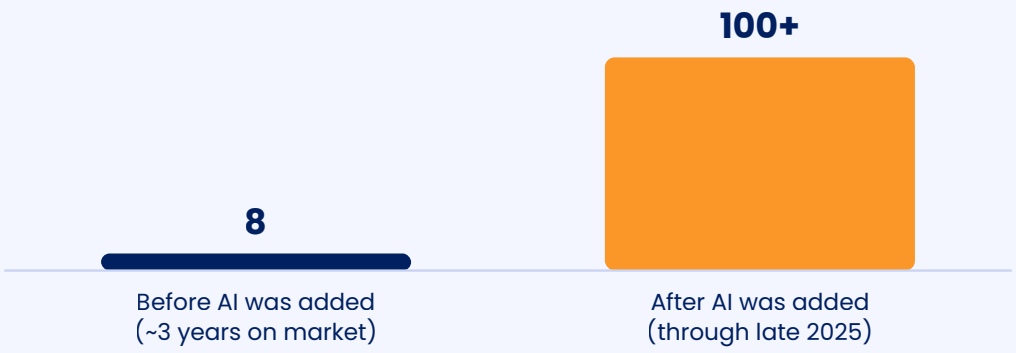
If an AI-generated note or recommendation were challenged tomorrow, by a plaintiff, a surveyor or a patient, could we show which tool produced it, which version was running, what policy applied, who reviewed it and when we last verified it performed as approved? ***For most health systems today, the honest answer is no.***

The human cost

Figure 8

Same device, same vendor, one software change

Unconfirmed malfunction and adverse-event reports to FDA for a sinus surgery navigation system, per Reuters (Feb. 2026)



Source: Reuters, "As AI enters the operating room, reports arise of botched surgeries and misidentified body parts" (Feb. 9, 2026). Reports are unconfirmed and do not establish causation.

The most serious costs can't be expressed in dollars. A February 2026 Reuters investigation examined a sinus surgery navigation system that had been on the market for about three years before its maker added machine-learning features in

2021. Before the change, FDA had received unconfirmed reports of seven malfunctions and one patient injury. After it, the agency received unconfirmed reports of at least 100 malfunctions and adverse events, including at least 10 injuries: a cerebrospinal fluid leak, a puncture at the base of a patient's skull and two strokes after a major artery was damaged. The manufacturer told Reuters there was no evidence linking the AI to any injury, and FDA reports are not designed to establish cause.

Whatever courts decide, the governance lesson stands on its own. It was the same device from the same vendor, used in the same operating rooms. What changed was a software update that added AI, which is exactly the kind of change a one-time approval never sees.

Quieter failures follow the same pattern: a scribe that adds a reassuring detail that never happened, a chatbot that answers a clinical question with confidence and an invented anatomical structure, a translation tool that may never have been validated for clinical language. Each error becomes the clinician's words the moment it is accepted.

Chapter 7

The Rulebook Is Moving In-House

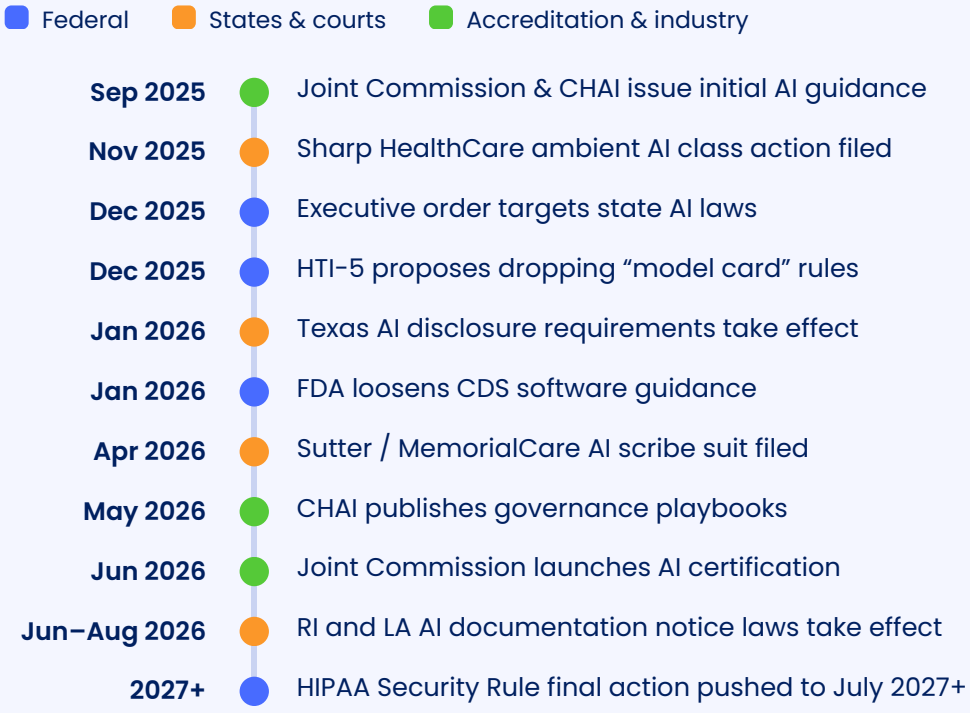
07

Federal oversight of clinical AI got lighter this year. Accountability didn't disappear; it moved to the organizations deploying the tools.

Figure 9

Twelve months that moved accountability to the health system

Selected federal, state, legal and accreditation milestones



Sources: Joint Commission; CHAI; FDA via Covington; ASTP/ONC via Covington; Manatt Health; Transparency Coalition; Interlace Health; court filings as reported.

Washington steps back

- **FDA.** On January 6, 2026, FDA issued revised clinical decision support guidance as a final document without a formal comment period. It widens the range of software FDA won't treat as a regulated device, says nothing specific about AI and sets no new policy for consumer-facing chatbots or symptom checkers.
- **ASTP/ONC.** The HTI-5 proposed rule, published in late December 2025, would eliminate the "model card" requirements that oblige EHR developers to disclose source attributes for predictive decision support, along with related risk management requirements. The comment period closed in February 2026, and the rule had not been finalized at our last check.
- **HIPAA.** OCR has pushed final action on the HIPAA Security Rule overhaul to at least July 2027. The proposal would require a written technology asset inventory and network map for every system that touches ePHI. The existing rule still applies in full.
- **The White House.** A December 2025 executive order directed federal agencies to challenge state AI laws viewed as burdensome. As of August 2026, state AI laws generally remain in effect while that effort plays out in court.

Less federal disclosure from developers means more diligence for deployers. If HTI-5 is finalized as proposed, the documentation health systems rely on to evaluate predictive tools could get thinner just as expectations for oversight rise. That matters because vendor documentation is already the weak link: in a November 2025 Black Book Research survey of 182 hospital leaders, 41% named limited vendor documentation, such as model cards and drift reports, as their top barrier to AI audits.

States fill the gap

In 2025, 47 states introduced more than 250 AI bills affecting health care, and 21 states signed 33 of them into law. The pace continued in 2026, with 27 states introducing 46 bills on AI in care settings and six states enacting seven. Several speak directly to the ambient AI cases in Chapter 6:

- **Rhode Island and Louisiana** now require providers to notify patients when AI is used to record or document a visit, effective June and August 2026 respectively.
- **Iowa** requires patient consent before AI is used to record and transcribe clinical interactions.

- **Texas** began requiring providers to disclose AI use to patients in plain language on January 1, 2026.

Accreditation arrives

The industry's own standards are moving faster than federal rules. On May 27, 2026, the Coalition for Health AI published governance playbooks developed with more than 150 health AI leaders, covering eight elements that include lifecycle management, risk and impact assessment, and third-party management. Days later, on June 1, the Joint Commission launched its voluntary Responsible Use of AI in Healthcare certification, open to organizations whether or not they hold Joint Commission accreditation.

The certification evaluates five areas: governance; data management; risk and bias reduction; monitoring, evaluating and validating safety, performance and responsible use; and transparency, education and training. It certifies how an organization uses AI, not individual products. In other words, it measures exactly the operational capabilities that Chapters 3 through 5 found missing.

What this means for compliance and risk leaders

The regulatory floor is now set by states, the bar by accreditors and the enforcement by plaintiffs' attorneys. All three ask for the same thing: evidence. Only 22% of hospitals in the Black Book survey said they were highly confident they could produce a complete AI audit trail within 30 days for regulators or payers, and just 29% had implemented and enforced policies covering AI model inventory, lineage and sign-offs.

Chapter 8

The Priority Paradox

08

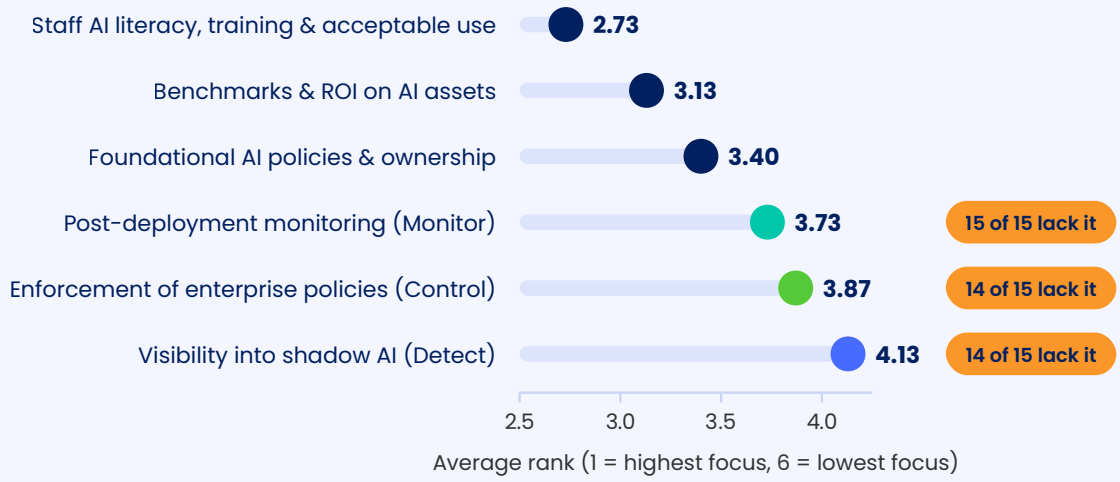
We asked CHIME respondents to rank six governance capabilities by how much focus each will get over the next 12 months. Training came first. Return on investment came second. Foundational policy came third.

The bottom three were visibility into shadow AI, enforcement of enterprise policy and post-deployment monitoring. Those are the same three capabilities with the widest gaps: 14 of 15 organizations lack real-time visibility, 14 of 15 lack real-time enforcement and 15 of 15 lack automated monitoring.

Figure 10

The priority paradox

Planned focus over the next 12 months vs. today's capability gap. CHIME Foundation member survey fielded for Vitea (n = 15).



Source: CHIME Foundation member survey fielded for Vitea, 2026 (n = 15). Lower average rank indicates higher planned focus.

Why it happens

This isn't a judgment failure. The top barrier to governance maturity, named by 7 of 15 respondents, is a lack of budget or dedicated staff. Three more cited competing priorities such as EHR and infrastructure work. Training and policy are familiar, affordable first steps. Real-time detection, control and monitoring sound like large projects that require people a health system doesn't have.

Spending reflects that. The Black Book survey found the median hospital planned to allocate 4.2% of its 2026 IT and quality or safety budgets to AI governance and safety, with small hospitals at 2.3% and large systems at 6.8%.

Why it's a trap

The deferred capabilities are the ones that can't be done by adding people. No team can manually review hundreds of AI apps, thousands of users and millions of interactions, or re-test every tool each time a vendor ships an update. A training program scales with headcount. Detection, control and monitoring only scale with automation, which makes them the best answer to the staffing problem rather than a casualty of it.

Training still matters. Focus group participants were clear that staff need to learn to question AI output, especially clinicians newer to these workflows. But training tells people what the rules are. It can't tell you whether anyone followed them.

The good news is that leaders are not skeptical of AI itself. Nine of 15 CHIME respondents described themselves as cautiously optimistic about AI's impact, two as all in and four as neutral. None called it a distraction. They want to use AI safely. The paradox is about sequencing, and sequencing can change.

Chapter 9

From Policy to Practice: What Good Looks Like

09

None of the leaders we heard from lacked commitment. What separates programs that work from programs on paper is whether governance decisions reach the point of use. Three capabilities do that work.

	Detect	Control	Monitor
The question it answers	What AI is running, who is using it and what data is it touching?	Are our policies applied every time, in real time?	Is approved AI still behaving the way we approved it?
What good looks like	Continuous discovery across sanctioned tools, embedded features, extensions, consumer apps and agents, attributed to people and roles	Machine-applicable policy enforced at the moment of interaction: block, redact, warn or log, with PHI kept inside the environment	Ongoing checks for accuracy, drift, hallucination and policy violations, plus alerts when vendors change AI behavior
Warning signs	Inventory built from surveys and contracts; discovery limited to a vendor's "GenAI" category	Enforcement relies on attestation or after-the-fact audits; agents share human credentials	Problems found by end users; testing ends at go-live; no audit-ready evidence

The Detect, Control, Monitor framework.

A practical maturity model

Level	Detect	Control	Monitor
1. Documented	Intake review for new tools; inventory maintained manually	Written AI policy and acceptable-use training	Pre-deployment testing and pilots
2. Observed	Continuous discovery of AI in use, including embedded AI and shadow tools	Policy violations visible as they happen	Usage and adoption tracked by department and role
3. Enforced	Activity tied to identity; personal and enterprise accounts distinguished	Policies applied automatically at the point of use, including for agents	Vendor changes trigger re-review automatically
4. Assured	Inventory is trusted and audit-ready at any moment	Enforcement evidence available for every interaction	Continuous assurance of accuracy and drift, reported to leadership and the board

Most CHIME respondents sit at Level 1, with some Level 2 visibility. The Joint Commission's certification areas map most directly to Levels 3 and 4.

Where to start: the first 90 days

1. **Get an observed inventory.** Replace a survey-based list with one built from network and identity data. It will be larger than expected, and it is the only reliable foundation for everything that follows.
2. **Start where the intensity is.** AI use is concentrated, which makes it governable in stages. At System B, the top 10% of users produced 75% of attributed activity, and habitual users, about a third of the people using AI, produced 91%. At System A, compliance staff, under 1% of users, used AI at more than three times the enterprise average intensity. That is not evidence of misuse; compliance teams review documents for a living. It is a practical first scope.
3. **Rationalize overlapping tools.** Where three to six products do the same job, consolidation reduces license spend, shrinks the review surface and lets one set of guardrails apply to everyone.
4. **Rewrite your five most important AI policies as rules a system could apply.** Specify the tools, the data, the action and the owner. If a rule still requires interpretation, it isn't ready.

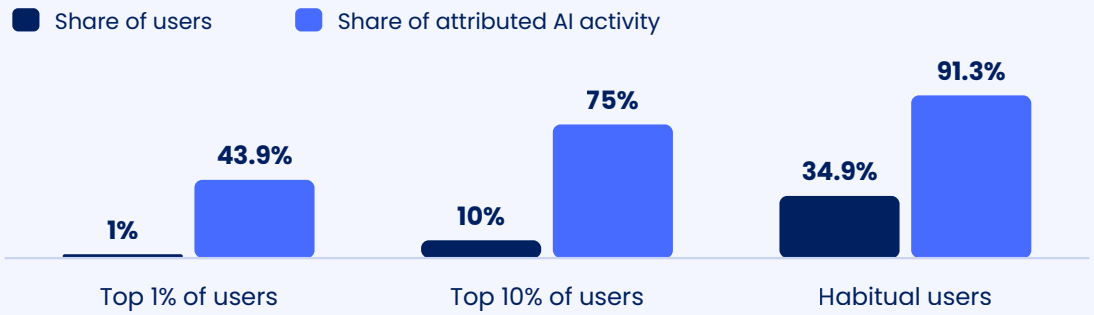
5. **Put approved tools on a monitoring cadence.** Add contract language requiring notice before AI features change, and watch for changes whether or not notice arrives.

6. **Decide who owns a wrong output and who owns repeat testing.** Write both down. These were the least-settled questions in our focus group, and they are the first ones a plaintiff or surveyor will ask.

Figure 11

A small group drives most of the AI activity

System B, attributed activity over a 110-day Vitea Lens observation window



Source: Vitea Lens assessment, System B, 2026.

Three questions for your next governance council

Focus group participants proposed these as the questions every council should be able to answer.

1

What's actually running?

Can we name every AI operating in our environment today, including features a vendor added after we approved their tool and tools staff started using on their own?

2

Can our policies actually be enforced?

Are they specific enough to hold on every AI interaction in real time, or written for a person to interpret after the fact?

3

Who owns a wrong output?

When an AI-generated note, summary or reply is inaccurate, is it clear that it became the user's words the moment they accepted it, and that the organization can show how it was produced?

Closing

The Longer Horizon

When our focus group looked past the next budget cycle, the worry wasn't a single breach or lawsuit. It was a slower erosion of the critical thinking and clinical judgment that governance exists to protect. ECRI raised the same concern in its 2026 patient safety report. The more fluent AI becomes, the easier it is to stop questioning it.

Several leaders compared this moment to healthcare's move from paper to electronic records. The organizations that came through that transition best weren't the ones that moved slowest. They built guardrails early and kept adjusting them as the technology changed.

AI is moving faster than the EHR ever did. That leaves less time to get governance right, and it makes starting now worth more. The policies most health systems have already written are a good beginning. The work ahead is making sure they hold where AI is actually used: on every prompt, every upload, every drafted note and every agent's action.

References

Sources

Proprietary sources: CHIME Foundation member survey fielded for Vitea (2026, n = 15); CHIME member focus group convened by Vitea (September 3, 2026); Vitea Lens assessments at multi-hospital U.S. health systems (2026).

Independent research and reporting

American Hospital Association. AHA Response to HHS RFI on AI in Health Care. February 23, 2026. Link

American Medical Association. AMA: AI usage among doctors doubles as confidence in technology grows (2026 Physician Survey on Augmented Intelligence). March 12, 2026. Link

Becker's Hospital Review. Hospitals face AI governance gaps heading into 2026, report finds (Black Book Research). November 2025. Link

Becker's Hospital Review. Most health systems lack dedicated AI testing platforms: Report (UPMC Center for Connected Medicine and KLAS Research). August 2026. Link

Becker's Hospital Review. 10 top patient safety concerns in 2026: ECRI. March 2026. Link

BankInfoSecurity. Kaiser Permanente to Pay Up to $47.5M in Web Tracker Lawsuit. December 2, 2025. Link

Coalition for Health AI. CHAI Releases Comprehensive Governance Playbooks. May 27, 2026. Link

Coalition for Health AI. Joint Commission and CHAI Release Initial Guidance. September 17, 2025. Link

Covington & Burling. 5 Key Takeaways from FDA's Revised Clinical Decision Support Software Guidance. January 2026. Link

Covington & Burling. HHS Proposes Changes to the Health IT Certification Program in HTI-5 Proposed Rule. January 15, 2026. Link

Cybersecurity Dive. As data breaches grow costlier, ungoverned AI creates new risks. July 29, 2026. Link

Drummond Group. How the HTI-5 Proposed Rule Affects ONC Certification. 2026. Link

EisnerAmper. AI in Healthcare: Liability and Governance Risks. July 30, 2026. Link

Epic (via Digital Health News). Epic Unveils Agent Factory and Expands AI Tools at HIMSS26. March 2026. Link

Fierce Healthcare. ECRI flags misuse of AI chatbots as a top health tech hazard in 2026. January 2026. Link

Fierce Healthcare. Epic expands AI ambitions with agent platform, Cosmos-powered predictions and workflow automation. August 2026. Link

Fierce Healthcare. Imprivata survey finds most health systems deploy some AI tools without formal IT approval. September 15, 2026. Link

Fierce Healthcare. Wolters Kluwer Health survey examines AI use and concerns among clinicians, patients in 2026. 2026. Link

Fisher Phillips. New Class Action Targets Healthcare AI Recordings. December 2025. Link

Health Law Advisor (Epstein Becker Green). Governing Health AI Development and Adoption: Insights from HHS's Strategy. 2026. Link

Health System CIO. Most Health Systems Validate AI Without a Testing Environment, Find KLAS and CCM at UPMC. August 6, 2026. Link

Healthcare Innovation. In Tandem, CHAI Releases AI Governance Playbooks, Joint Commission Launches Certification. June 2026. Link

Heidi. Clinicians Are Adopting AI to Solve the Documentation Crisis, New Global Survey Finds. July 15, 2026. Link

IBM and Ponemon Institute. Cost of a Data Breach Report 2026. July 2026. Link

Joint Commission. Joint Commission Releases Voluntary Responsible Use of AI in Healthcare Certification. June 1, 2026. Link

Law360. Ambient AI notetakers and liability. July 7, 2026. Link

Manatt Health. Health AI Policy Tracker, Q3 2025 Update. October 2025. Link

Medical Xpress. Does AI belong in the exam room? Lawsuit alleges Sharp violated patient privacy. January 2026. Link

PubMed. Quality of Clinical Notes Created by Ambient Listening Generative AI: Pragmatic Prospective Pilot Study. April 17, 2026. Link

Security Management (ASIS). Threat Actors Embrace AI, Helping Push the Average Cost of a Data Breach Higher (IBM 2026). July 2026. Link

TechTarget. AI testing, validation common but inconsistent across healthcare providers. August 2026. Link

Wolters Kluwer Health. Wolters Kluwer survey finds broad presence of unsanctioned AI tools in hospitals and health systems. January 22, 2026. Link

HIPAA Journal. Global Data Breach Cost Rises 12% to Almost $5 Million (IBM Cost of a Data Breach Report 2026). July 29, 2026. Link

Interlace Health. Ambient AI Isn't the Liability. Your Consent Process Is.. September 2026. Link

Kiteworks. The IBM 2026 Cost of a Data Breach Report Proves AI Governance Failure, Not AI Itself, Is Driving Costs. July 2026. Link

LiveCompliance. AI Healthcare Regulations 2026: Federal and State Rules. August 2026. Link

Manatt Health. Health AI Policy Tracker. 2026. Link

Paubox. HHS pushes final HIPAA Security Rule update to 2027. 2026. Link

Reuters. As AI enters the operating room, reports arise of botched surgeries and misidentified body parts. February 9, 2026. Link

TechTarget. Sutter Health, MemorialCare face class action lawsuit over AI scribe use. April 2026. Link

Transparency Coalition. State lawmakers have passed new laws regulating the use of AI in health care. July 27, 2026. Link

Wolters Kluwer Health. 2026 Future Ready Healthcare survey. June 2026. Link

About Vitea

Policy that holds on every AI prompt

Most AI governance stops at the document, the inventory or the alert that arrives after the data is already gone. Vitea is the only platform that enforces your AI policies at the prompt level, in real time, at the moment someone uses AI. That is the gap this report found in 14 of 15 health systems, and it is the gap Vitea was built to close.

Detect

Lens

Discovers the AI in use across your environment, including shadow tools, embedded features, extensions and agents, and ties activity to people, roles and data.

Control

Command

Every prompt, upload and agent action is evaluated against your policies before anything moves, and Command blocks, redacts, warns or logs accordingly. It runs as an inline gateway inside your own cloud or on-premises environment, makes its decision in under 500 milliseconds and never sends PHI outside your walls to do it.

Monitor

Pulse

Continuously monitors approved AI for accuracy, drift and policy violations, and flags when vendors add or change AI in tools you've already approved.

Vitea also provides testing, evaluation and red teaming so health systems can validate AI before and after it reaches clinicians and patients.

Turn your AI policies into controls that hold

Your organization has already written the rules. Our team will show you how Vitea applies them to every prompt, across sanctioned tools, embedded AI features and consumer apps, in real time and without PHI leaving your environment.

Book a call with our team here.

Vitea

AI Governance for Healthcare

vitea.ai

CHAI
COALITION FOR HEALTH AI

CHiME

FIERCE
AI Innovation Awards
WINNER
2026