99TH NHA ANNUAL Convention

SCAN HERE

Operationalizing AI Governance for Hospitals

Brenton Hill, JD, MHA
Head of Operations and General Counsel
Coalition for Health AI

NHA | NEBRASKA HOSPITALS

Operationalizing AI Governance



Brenton W. Hill, JD, MHA
Head of Operations and General Counsel,
Coalition for Health AI

2026

Nebraska Hospital Association | October 22, 2026

Why CHAI's lawyer is talking to Nebraska hospitals

Brenton W. Hill, JD, MHA

Head of Operations and General Counsel,
Coalition for Health AI (CHAI)

- Attorney and health administrator; previously at Mayo Clinic
- Leads CHAI's working group of 40+ health system attorneys
- Facilitates CHAI's AI Monitoring Subgroup of Chief AI Officers
- Convenes CHAI's partnership with the Joint Commission on responsible AI guidance and certification
- Oklahoma City based, so the rural and critical access picture is familiar

The Coalition for Health AI, in three numbers

3,000+ member organizations: health systems, payers, vendors, patient advocates and academics

8 governance playbooks that the Joint Commission crosswalked line by line into its Responsible Use of AI certification

57 standardized vendor intake controls in the draft AI Intake Playbook, distilled from 800+ questions in use at health systems

CHAI is an industry-led, public-private partnership. It does not certify, evaluate, or receive PHI; its frameworks set the bar that others assess against.

Agenda

Where we are going in the next 50 minutes

01	**The governance baseline** Joint Commission expectations and the CHAI playbooks that sit underneath them	**7 min**
02	**Contracting for AI** The specific terms that structure the vendor agreement, and what is realistically winnable	**12 min**
03	**The operating model** How leading systems are staffing AI: product managers and designers embedded with end users	**9 min**
04	**Where monitoring is going** Risk-tiered metrics per use case, and runtime evidence for traceable governance	**7 min**

Discussion and questions: bring your hardest vendor clause and your smallest hospital — **15 min**

01

The governance baseline

What accreditors expect, and the playbooks that get you there

CHAI

The Joint Commission has defined responsible AI

RUAIH guidance, issued with CHAI in September 2025; voluntary certification opened July 2026

1 AI policies and governance structures

Named accountable leaders, a multidisciplinary group, an inventory of AI in use

2 Patient privacy and transparency

Tell patients when AI touches their care; keep disclosures accurate

3 Data security and use protections

Minimum necessary, de-identification, limits on secondary use by vendors

4 Ongoing quality monitoring

Risk-based, context-appropriate monitoring; responsibility settled in the contract

5 Reporting of AI safety events

A process to report AI-related events internally and to a PSO

6 Risk and bias assessment

Assess risk and bias before deployment and across the lifecycle

7 Education and training

Upskill clinicians, staff and leaders on safe use

Line-level crosswalk

The certification standards map directly onto CHAI's eight-playbook governance series. Build to the playbooks and you are building to the survey.

"Note that monitoring responsibility should be discussed as part of third-party procurement and contracting."

Joint Commission and CHAI, RUAIH Guidance, Element 4

CHAI playbooks make the seven elements operational

CHAI Governance Framework: four domains, five process subdomains, one playbook each

Domain 1

AI Policy

One organization-wide policy for buying, building and using AI. The document surveyors ask for first.

Domain 2

Organizational Structure

A cross-functional governance group with named roles from intake through monitoring.

Domain 3

Organizational Resources

The data, tooling, compute and people AI needs, and how they are protected.

Domain 4

Organizational Processes: where governance becomes daily work

4.1 Responsible AI lifecycle management

Intended use, inventory, lifecycle monitoring

4.2 Risk and impact assessments

Tier the use case before deployment

4.3 Responsible data management

Acquisition, de-identification, bias and quality checks

4.4 Third party management

Liability, disclosure and responsibilities with vendors

4.5 Feedback and training

User feedback, incident reporting, clinician training

Free at chai.org/workgroup/cross-cutting/ai-governance. Companion drafts: AI Intake Playbook (57 vendor controls, public comment) and the Ambient AI Implementation Playbook.

Three questions set the risk tier of any AI use case

Adapted from CHAI Subdomain 4.2. The tier decides the review, the contract terms, and the monitoring.

1 How close is it to patient care?

Scheduling and back office sit far from the bedside. Diagnosis, triage and treatment recommendations sit on top of it.

2 Does PHI flow in or out?

Names, notes, recordings, claims data. If patient information feeds the tool, privacy duties attach: BAA, minimum necessary, consent.

3 Who checks the output, and when?

A human reviewing every output is your strongest brake. Autonomous action or auto-filed text removes it.

LOW risk

Admin tasks, no PHI, human always in control. Light review, standard terms.

MEDIUM risk

Assists clinicians or touches PHI, human still signs off. Standard review plus the AI addendum.

HIGH risk

Influences care decisions, acts autonomously, or handles PHI without clear oversight. Full review, full terms, defined monitoring before go-live.

Worked example: an ambient scribe is Medium risk (PHI-heavy, but a physician signs every note). A sepsis agent that writes to the chart is High. A free public chatbot with PHI pasted in is High, always: no BAA, no contract, PHI leaving the building.

02

Contracting for AI

The terms that structure the agreement, drawn from CHAI Subdomain 4.4

CHAI

Your BAA predates AI. The AI addendum closes the gap

Three documents, three jobs. Most hospitals have the first two and are missing the third.

BAA

Business Associate Agreement

PHI access, transmission and permitted uses under HIPAA. Most BAAs are silent on model training and secondary fine-tuning.

DPA

Data Processing Agreement

Purposes, retention schedules, deletion obligations, security standards, subprocessor list. Covers data HIPAA does not reach.

AI Addendum

The AI-specific terms

Training prohibition, performance warranty, change notice, monitoring duties, incident notice, suspension right, indemnity carve-out, deletion certification, precedence over click-wrap.

Two public models worth borrowing

Mercy's AI Terms and Kaiser Permanente's AI Attachment for Vendors are both reproduced in CHAI Subdomain 4.4. Kaiser's is posted publicly on its supplier site. Borrow the structure, adapt the numbers with counsel.

Contract language in use today: two patterns

Reproduced in CHAI Subdomain 4.4 Third Party Management. Quoted terms are from the public documents.

Mercy AI Terms

Performance to spec (§ 4)

Must "perform in accordance with the accuracy, precision, and performance specifications set forth in the applicable SOW."

Ongoing compliance (§ 3)

Vendor continuously monitors for drift; notifies Customer "within two (2) business days" of degradation or an adverse incident, with a remediation plan.

No-fee suspension (§ 3)

Customer may suspend use "in its sole discretion" pending review, at "no additional fee."

Material AI Modification

Triggers pre-deployment notice and consent "regardless of whether it is accompanied by a version number" change.

Kaiser Permanente AI Attachment

Training prohibition (§ 2.2)

Vendor is "expressly prohibited from using… KP Data, Outputs, or Personal Information (including anonymized, de-identified or aggregated)… to train, fine tune, improve…"

Output ownership (§ 2.3)

"Customer will be the owner of all Outputs and all Intellectual Property Rights therein," with irrevocable assignment.

Termination (§ 2.5)

Return Outputs, then destroy within "thirty (30) calendar days" plus "written certification of the destruction."

Precedence and click-wrap (§ 2.1, 2.7)

The AI Attachment "takes precedence and governs"; any click-wrap terms "will be null and void."

The pattern: name the data, name the trigger, name the deadline, name the remedy.

Structuring the agreement: terms 1 to 5

Data and performance. Use this as the checklist for every Medium- and High-risk tool; tier decides how hard you push.

Term	What to ask for	Why it matters	Anchor
1. Training and secondary use	Flat prohibition on using your data, outputs or PHI (including de-identified or aggregated) to train or improve products; consent before any secondary use	"De-identified" data can feed their models forever; AI raises re-identification risk	KP § 2.2
2. Performance warranty	Perform to documented accuracy, precision and performance specs in the SOW; not "substantially as documented"	An AI tool with no accuracy promise is a liability transfer	Mercy § 4
3. Change notice and consent	30 to 90 day notice of material AI changes, with the right to test, delay or opt out; applies even without a version number change	The tool you validated is not the tool you will be running in six months	Mercy
4. Monitoring and incident notice	Vendor monitors for drift; notifies within 2 business days of degradation or adverse incident with a remediation plan; shares monitoring data you can act on	Satisfies RUAIH Element 4; puts the monitoring burden where the telemetry lives	Mercy § 3; RUAIH
5. Suspension and rollback	Suspend use at your sole discretion, no fee, pending review; independent ability to disable the tool without vendor cooperation	Your kill switch should not depend on a support ticket	Mercy § 3

Sources: CHAI AI Governance Playbook, Subdomain 4.4 Third Party Management; Mercy AI Terms; Kaiser Permanente Artificial Intelligence Requirements for Vendors (public). Adapt with counsel.

Structuring the agreement: terms 6 to 10

Risk, data, exit and paper. The second half protects you when something goes wrong and when the relationship ends.

Term	What to ask for	Why it matters	Anchor
6. Liability and indemnity	Indemnification for AI errors, regulatory non-compliance and UPL claims; AI-error carve-out from the liability cap	You own every clinical decision while their exposure is three months of fees	CHAI 4.4
7. Output ownership	Customer owns outputs and IP in them; vendor keeps pre-existing IP; configurations and derived improvements addressed explicitly	Avoids a dispute over your own chart notes and models	KP § 2.3
8. Deletion and return	Return, then destroy on a stated timeline (30 days) with written certification	"Standard practices" is not a deadline	KP § 2.5
9. Subprocessors and audit	Named subprocessor list, flow-down of all terms, annual audit rights on data handling, AI-specific breach notice including model exposure	Your contract is with the primary vendor; the chain behind it is opaque	CHAI 4.4
10. Precedence and exit	AI addendum governs over click-wrap and portal terms; termination for cause with no-cost transition services and data migration	Click-through terms quietly undo negotiated ones	KP § 2.1, 2.7

Sources: CHAI AI Governance Playbook, Subdomain 4.4 Third Party Management; Mercy AI Terms; Kaiser Permanente Artificial Intelligence Requirements for Vendors (public). Adapt with counsel.

Negotiating reality: what you can actually win

CHAI Subdomain 4.4, Common Challenges in AI Vendor Contracting

The power imbalance

- Dominant EHRs, hyperscalers and foundation-model providers rarely give bespoke AI terms to a single hospital
- Embedded AI auto-updates inside health IT whether you negotiated or not
- Vendors routinely push back on liability and broad indemnification

Realistically winnable

- Advance notice of new AI features and the ability to opt out of activation
- Disclosure of known model limitations and what the vendor monitors; treat as non-negotiable even when dollars are
- 30 to 90 day change notice, version lock or opt-out for clinical uses
- A notice-and-cure pathway: written notice, remediation plan, feature suspension, termination for cause

When terms stay thin, control the risk yourself

- Human-in-the-loop: outputs surface as drafts; sign-off requires affirmative human action
- Independent rollback: disable the tool without vendor cooperation
- Default-off in sensitive contexts; track override and edit rates
- No indemnity clause prevents harm; it argues about the bill afterward. Operational controls stop the injury

Leverage you have that a 25-bed hospital alone does not: the Nebraska Hospital Association, your GPO, and a shared addendum. Vendors negotiate with a state once; they do not negotiate with 90 hospitals separately.

CHAI

Five vendor questions, and what a dispute could look like

1. **Is your product regulated by FDA as a medical device?**
   Classification? Clearance current? Cleared is not the same as safe, compliant or permanent.

2. **How do you handle PHI?**
   BAA and DPA in place? Training-data restrictions? Deletion obligations? Storage outside the US?

3. **What are your performance warranties?**
   Specific benchmarks? What happens when the model changes? Change-notification window?

4. **Who is liable if AI output contributes to patient harm?**
   Indemnification scope? AI-error carve-out from the cap? Adverse event reporting?

5. **What are your security practices and subprocessor oversight?**
   Terms flow down? No-cost data migration on vendor breach?

Get the answers in writing and attach them to the contract. A vendor who will not answer these has answered them.

CommonSpirit v. Innovaccer

Del. Super. Ct., filed April 2026, $32M+ claimed. As alleged in the complaint. Since settled; the parties opted to continue partnering.

- Vendor allegedly misrepresented its ability to cure deficiencies and migrate data onto a new platform
- Allegedly altered its own project plan to conceal delays; refused contractually required service credits (about $3.4M)
- Allegedly refused to cooperate on post-termination transition and data migration

Contract lessons

- Specific, measurable performance warranties
- Milestones with automatic financial penalties for vendor delay
- Capability reps tied to your data environment and workflows
- Written approval for every project-plan change
- Termination for cause plus no-cost transition services

03

The operating model

How leading systems are staffing AI: product management, embedded with the people who use it

Product management is the trending operating model

Leading systems run AI like a product portfolio: vendor management, configuration, governance, redefined roles

DISCOVER

Embedded insights

Product managers own the relationship with the vendor and the internal customer. They sit with end users, map journeys, and define the problem before the tool.

GOVERN

Support from above

A central AI governance function provides intake support, risk assessment, recommended controls and contracting support. It does not own the product.

DELIVER

Decision ownership

Business and clinical leaders partner on ongoing management of the solution (performance, bias, education) with the product owner and central governance supporting.

The Chief AI Officer seat changes what gets built

Where systems have a Chief Health AI Officer (UC San Diego, Mass General Brigham, Inova, AdventHealth), AI sits at the executive table and is staffed with product, engineering, ops and clinical roles not slotted under IT. Across CHAI's membership, organizations with a named AI executive consistently move first.

Why this matters for a community hospital

This is the structure Chief AI Officers are choosing to move away from data science and engineering-heavy teams. The scarce skill is not building models; it is owning a vendor relationship and an end-user workflow. That is a role a 100-bed hospital can staff, often from people it already has.

UC San Diego's forward-deployed automation team

Chief Health AI Officer Karandeep Singh, MD. Core builder team plus staffing for evaluation.

Tech. Product Owner

Defines product priorities, translates needs into technical requirements, guides delivery through deployment

Ops. Project Manager

Leads requirements gathering, process optimization, value realization and change management

Automation Designers

Explore new opportunities, prototype quickly, surface promising agents for development

Config. Analysts

Optimize live agents through parameter and prompt tuning; adapt existing agents

Full Stack Engineer

Designs and builds new agents, deploys tooling, configures APIs across enterprise systems

Data Scientist

Builds predictive and generative capabilities; provides statistical performance evidence

EHR Analysts

Deep workflow expertise across EHR workflows, interfaces and enterprise platforms

Advisory Teams

Regulatory, patient experience, risk and safety, ELSI, provider wellness and clinical guidance

4,000+

AI preparatory calls for colonoscopy and endoscopy since launch

84%

net promoter score on those calls (average 9.0 of 10)

10-15% to 6%

no-show and same-day cancellation rate, baseline to date

500+

test calls before go-live, including 6+ physicians and 20 nurses red-teaming

Source: Karandeep Singh, MD, UC San Diego Health, presentation to CHAI, April 2026. Figures as reported at that time.

Inova's AI product governance: four stages, one team

AI Product Management sits at the intersection of technology strategy, governance, and the voice of the customer

1 Idea assessment

Formal pre-assessment: R&D, user journey mapping, initial risk and market assessment

Key questions

What clinical or operational gap does it address? How is it regulated? Who are the stakeholders? Do we already own something similar?

2 AI product profile

Intake and due diligence: vendor profile across user journey, feature set, data privacy and security; RFI if needed

Key questions

Workflow fit, HIPAA check, security architecture review, clinical evidence and bias assessment

3 Governance approval

Multidisciplinary review across clinical, operational and technical stakeholders

Key questions

Clinical leadership sign-off, IT and security validation, AI ethics subcommittee review, executive sponsor approval

4 Deploy and performance oversight

Implementation support, post-launch monitoring, safety checks, re-evaluation as needed

Ongoing activities

Business and operational metrics decomposed into technical KPIs with data governance, IT and the vendor

"There is no wrong door."

Have an idea: engage us with your IT leader. Evaluating a vendor: we keep a portfolio of vetted vendors. Worried about risk: that is our job. Already in progress: we embed at any stage to add governance, vendor management or post-launch monitoring.

Source: Inova Health, AI Product Management overview, 2026. Reproduced with the team's presentation to CHAI.

Forward deployed: product teams sit with end users

Governance in a committee room cannot see the workflow. Governance in the clinic can.

How Inova listens

- Regular recurring interviews with functional end-user groups and customer sponsors
- Co-navigating user experiences and journeys with IT teams, SMEs and service lines
- User journey research and contextual inquiry (watching the work, not surveying it)
- Quantitative outcomes agreed with the sponsor before build

Problem definition framework

[Person] wants to [achieve something] and needs to feel [an emotion or state] but cannot because of [something in the way] which happens when [timing] because [root cause].

Define the problem in the user's words before you define the tool.

The same model, sized for a Nebraska hospital

One AI product owner

A nurse informaticist, a quality leader or an operations director. Owns the vendor relationship and the workflow. Not a new hire in most cases.

One clinical champion per use case

The person whose notes, calls or orders the tool touches. They do the pre-go-live testing and the post-go-live review.

Existing structures, not new ones

Quality, patient safety and compliance committees already exist. RUAIH says to use them. Add AI as a standing agenda item.

A shared bench through the association

Vendor vetting, model cards, contract addenda and monitoring templates do not need to be rebuilt in every hospital.

04

Where monitoring is going

Defined metrics for every use case, scaled to risk, and runtime evidence for traceable governance

CHAI

Defined metrics for each use case, scaled to its risk tier

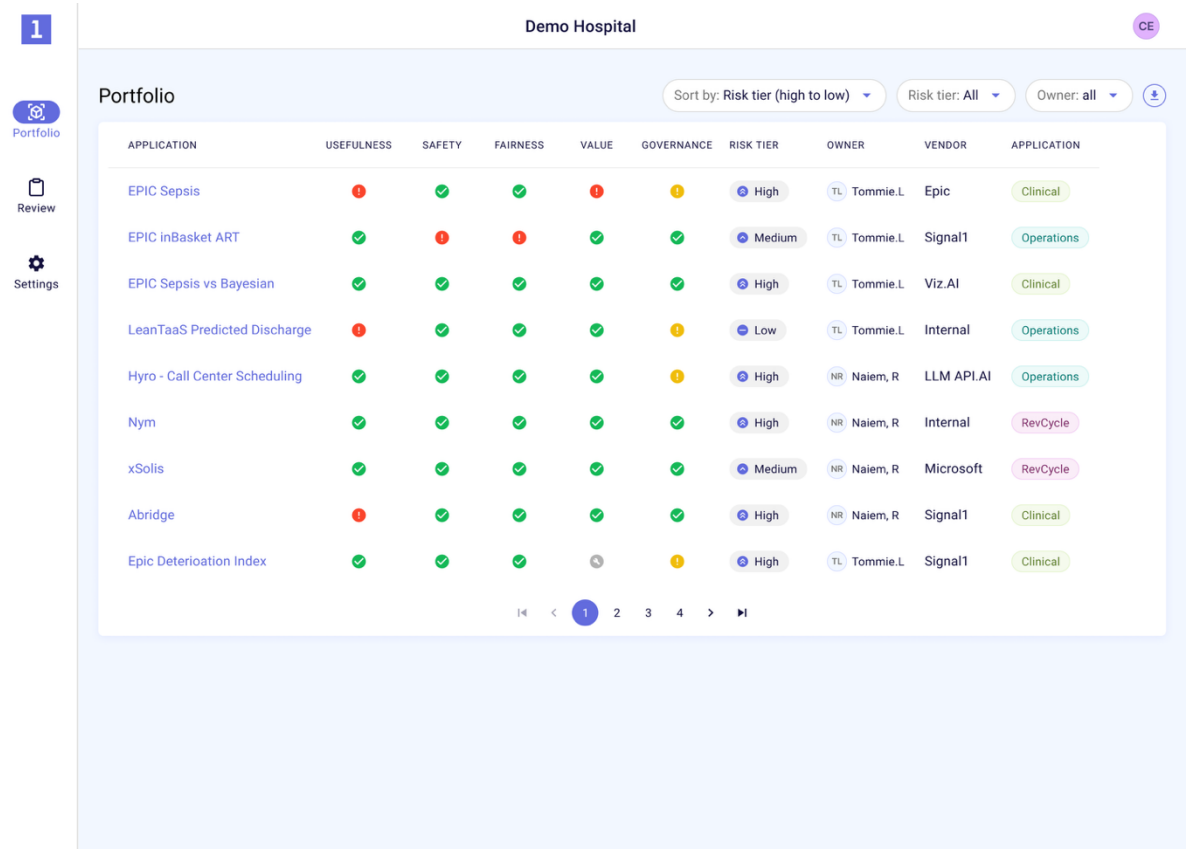
RUAIH: risk-based, scaled to your setting. Inova: decompose business metrics into technical KPIs.

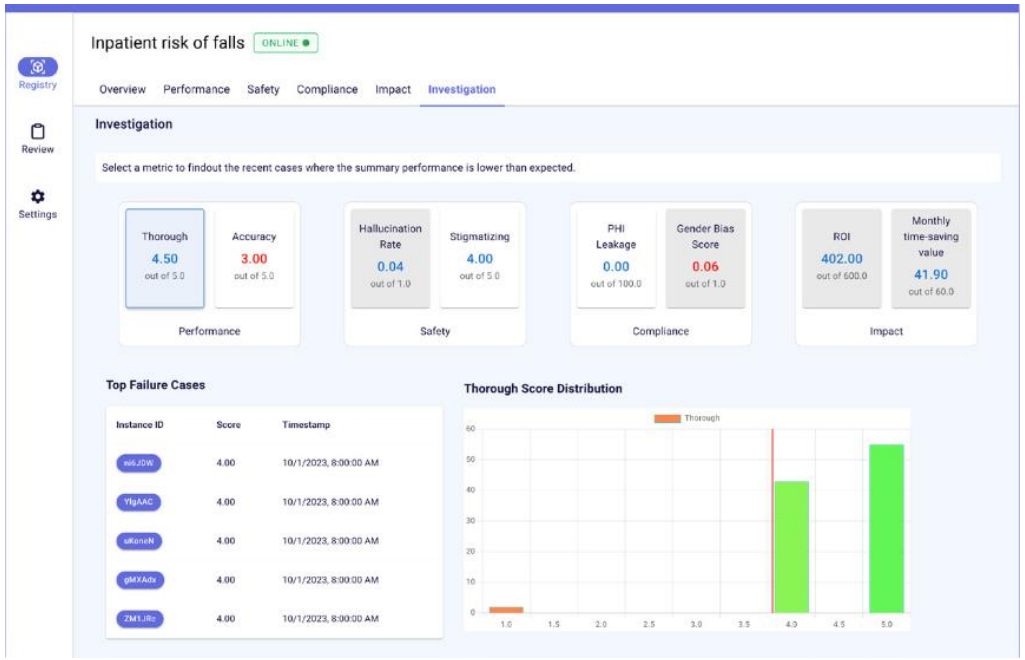
Use case	Tier	Business or clinical metric	Technical KPIs to watch	Cadence and owner
Ambient scribe	MEDIUM	Note turnaround time, clinician after-hours EHR time, patient complaints	Edit rate per note, hallucination and omission rate on sampled notes, consent-capture rate per encounter, model version	Monthly sample review by clinical champion; quarterly to quality committee
Sepsis or deterioration model	HIGH	Time to antibiotics, mortality, ICU transfers, alert fatigue	PPV and sensitivity by unit and population, alert volume per 100 patients, override rate, drift vs. validation baseline	Weekly dashboard; monthly clinical review; immediate review on drift threshold
Patient-facing scheduling agent	MEDIUM	No-show rate, connected rate, net promoter score, appointments booked	Escalation rate to humans, transcript-flag rate, wrong-answer rate on red-team scripts, call completion	Daily transcript review of escalations by nursing; monthly to ops leader
Coding or claims assistant	LOW	Denial rate, days in A/R, coder productivity	Suggestion acceptance rate, audit-sample accuracy, code-mix shift vs. baseline	Quarterly audit sample by revenue cycle; compliance review annually
Public website chatbot	HIGH	Call deflection, complaints, UPL and privacy exposure	Share of responses outside the approved scope, medical-advice leakage on test prompts, PHI capture events	Weekly test-prompt run by marketing plus compliance; kill switch owned by IT

Illustrative starting points, not CHAI standards. The CHAI AI Monitoring Subgroup's evidence paper on monitoring practice is due this fall.

What a monitoring layer looks like in practice

Inova and Signal 1, AI operations UI (beta wireframe). Centralized monitoring, federated management.





What to notice

Every tool has an owner, a vendor and a risk tier. Each is scored on usefulness, safety, fairness, value and governance. The drilldown turns a metric into the specific cases where performance fell short.

Source: Inova Health and Signal 1, AI Management Solution UI wireframe (beta), 2026. Demo data.

CHAI

Runtime evidence: policy binder to patient encounter

Accreditation asks whether governance exists. Runtime evidence shows it ran when care was delivered.

Policy and governance
What should exist

Certification review
Is the program defined?

Covered event
A patient-facing event occurs

Control executes
Consent, guardrail, monitoring check

Evidence receipt
Timestamped, signed record of execution

Independent verification
Surveyor, insurer or board checks the receipt

Documented governance

Runtime assurance

The standard

OVERT (Observable Verification Evidence for Runtime Trust): an open, royalty-free specification for verifiable operational evidence from AI systems. CHAI became co-steward in September 2026. Privacy-preserving by design: no protected content leaves the operator's environment.

Already in use

Utah's Office of AI Policy requires a valid receipt per covered event as a condition of operating in its AI sandbox. Ambient clinical AI vendors emit a receipt per encounter showing consent was captured before recording began. Lloyd's Lab is exploring it as the assurance layer for underwriting AI risk.

Start narrow

First covered event: consent for ambient recording. One encounter, one receipt. Expand to guardrails, model version and clinician review as the standards work matures.

One evidence layer for certification, insurers and courts

1 Certification

Certification today relies on documented policies and a survey-day review. Runtime evidence lets a surveyor sample encounters and verify that authorization, data boundaries and clinician review actually executed. Members told us a certification without it could feel performative.

2 The insurance market

Underwriters want a real-time baseline of how a system performs. CHAI's attorney working group has been asked to help co-design metrics with an actuary and broker. Likely effect: liability shared between deployer and developer, priced on evidence.

3 Defensibility

The six-months-later question ("why did the system generate this code?") gets answered with a record instead of a reconstruction. Audit logs of what the AI decided, when, and who reviewed it become the first document requested.

What a hospital can do now

Ask every vendor with a Medium or High-Risk tool three questions: What do you log per event? Can I retrieve it for a specific patient encounter? Can a third party verify it without your help? Write the answers into the monitoring clause. That is runtime evidence readiness without buying anything.

Close

Monday morning: three commitments, no new headcount

THIS WEEK

Inventory and tier

List every AI tool in use, including the ones embedded in the EHR and the free ones staff use anyway. Rate each one Low, Medium or High with the three questions.

THIS MONTH

Attach the addendum

Pull the Mercy and Kaiser patterns from CHAI Subdomain 4.4. Draft a one-page AI addendum. Attach it to the next Medium- or High-risk renewal. Ask the five vendor questions in writing.

THIS QUARTER

Owners and metrics

One product owner and one clinical champion per Medium- or High-risk tool. Two metrics each: one business, one technical. Put AI on the quality committee agenda as a standing item.

Free resources

CHAI AI Governance Playbooks
chai.org/workgroup/cross-cutting/ai-governance

Subdomain 4.4 contract anchors
governance.chai.org/subdomain-4-4-third-party-management

Joint Commission and CHAI RUAIH guidance
jointcommission.org (Responsible Use of AI in Healthcare)

Kaiser Permanente AI vendor requirements
supplier.kp.org (Artificial Intelligence Requirements for Vendors)

Questions

Bring your hardest vendor clause and your smallest hospital.

Brenton W. Hill, JD, MHA
Head of Operations and General Counsel, Coalition for Health AI
brenton@chai.org

99TH NHA ANNUAL
Convention

Tell us your thoughts on this session!

This quick survey takes less than a minute, and your feedback helps us improve future sessions.

SCAN HERE

NHA | NEBRASKA HOSPITALS