• A Credo AI Playbook · Healthcare Edition

AI Governance as a Risk and Compliance Foundation

A practical guide for healthcare AI and governance leaders preparing for the next audit cycle, written in the language of the leaders we've sat across from.

Companion to The State of AI Governance in Healthcare Report
The Joint Commission 2027 · ISO 42001 · HITRUST · CMS 0057-F · EU AI Act · HIPAA

Table of contents

A NOTE ON THIS EDITION

We wrote this playbook after working with healthcare AI leaders who were struggling to build a business case for AI Governance focused on ROI. They were running against a deadlines from The Joint Commission 2027, ISO 42001, HITRUST recertification, the EU AI Act. This healthcare edition leads with that lens, not with governance theory.

SECTION 01

Healthcare AI leaders aren't debating governance. They're reacting to dates on a calendar.

AI Healthcare Leaders aren't debating whether to build a governance program or not. They're trying to understand how their program can produce an evidence trail before six external deadlines that are looming.

2027 · APPROACHING

The Joint Commission

New standards for AI-assisted clinical decisions: documented oversight, monitoring, evidence on file not a one-time approval.

ACTIVE · CERT CYCLES

ISO 42001

The AI Management System standard. Increasingly required of vendors and orgs pursuing broader quality or security certifications.

NOW · R2 / I1 CYCLES

HITRUST

AI-specific control requirements landing in HITRUST r2 and i1 audits. Governance documentation is becoming a required audit artifact.

ENFORCEMENT UNDERWAY

EU AI Act

Clinical decision support classified high-risk. Conformity assessments, transparency, post-market monitoring. Obligations land on the deployer.

IN FORCE · API JAN 2027

CMS 0057-F

72-hour prior-auth decisions and denial documentation already required for MA, Medicaid, and marketplace plans. FHIR API mandate lands Jan 2027 & CMS holds deployers accountable for vendor AI.

ALREADY IN FORCE

HIPAA, today

Access controls, audit trails, and minimum-necessary use apply to every AI system that touches PHI. Most of yours do.

> *"We've got AI features going into products. There's some in production, some in early access, a lot being discussed. It's a bit of a free for all at this moment in time."*

AI Governance Lead, large healthcare payer

THE STRUCTURAL PROBLEM

Most healthcare AI programs were designed to govern 1–2 use cases. They were not designed for 20, 50, or 130+ which is where most large health systems and payers are already rapidly approaching.

SECTION 02 · CURRENT STATE

Your program may look complete, but will it survive external scrutiny?

The root cause isn't lack of intent. It's structural: AI governance in healthcare typically lives across three or four functions — security, privacy, compliance, the AI/data team — with no single owner and no shared system of record.

Today, only 12% of organizations use a dedicated AI governance platform. The rest govern with spreadsheets, SharePoints, and documents circulated via email. As deployment scales, the burden compounds: 45% of governance teams spend 11–20 hours a week on manual intake, climbing to 21–40 hours for the largest programs. And 60%+ of enterprises now deploy AI across multiple departments. In healthcare, that means clinical tools, revenue cycle AI, ambient documentation, and vendor-embedded AI all running simultaneously, often governed (or not) by different teams with no central view.

From the outside, the posture has four predictable cracks:

- Risk reviews exist, but they're point-in-time not ongoing.
- **Documentation lives in multiple places** nobody can locate under audit conditions.
- **Vendor AI tools are approved once** and never monitored again.
- **The AI registry is a spreadsheet** someone started maintaining six months ago and stopped updating.

> *"The process has grown from a vendor management privacy review that's had AI considerations bolted on. We don't have anyone formally running this, and I'm at least three months behind where I want to be."*
> *Director of AI, regional health system*

THE AUDIT GAP, PLAINLY

What compliance programs ask for isn't a policy document. It's an **evidence trail**: documented risk assessments per AI system, ongoing monitoring records, access logs, vendor evaluations, change control documentation.

Spreadsheets don't produce evidence trails. Governance platforms do.

SECTION 03 · DIAGNOSTIC

Six common gaps that map to existing regulations.

A structured way to see your own exposure. Every gap below is a control an auditor, certification body, or regulator can ask you to evidence.

01

No centralized AI inventory

You can't govern what you can't see. Most healthcare programs don't have a complete list of every AI system in production, including vendor-embedded tools inside critical systems.

WHERE IT'S REQUIRED

Joint Commission 2027 · ISO 42001 · EU AI Act · CMS 0057-F all require a documented inventory of every AI system in scope, including vendor tools.

02

No risk tiering

Every use case moves through the same review, regardless of clinical impact or PHI exposure. Low-risk tools slow down unnecessarily; high-risk tools don't get proportional scrutiny.

WHERE IT'S REQUIRED

HIPAA minimum-necessary principle · EU AI Act high-risk classification framework.

03

No ongoing monitoring

Pre-deployment approval is not governance. Clinical AI drifts. Models trained on pre-pandemic data behave differently post-pandemic. Regulatory expectations evolve. None of it is captured by a risk assessment from 18 months ago.

WHERE IT'S REQUIRED Joint Commission & ISO 42001 expect evidence of continuous oversight, not initial approval. CMS 0057-F holds payers accountable for what AI produces in production and deployment.

04

No vendor AI oversight

Your EHR vendor has embedded AI. Your payer portal has AI. Your ambient documentation tool is an AI system. For most healthcare orgs, evaluating a single AI vendor takes hours without a scalable process or shared scorecard.

WHERE IT'S REQUIRED

HIPAA BAAs · Joint Commission · ISO 42001 supplier management · EU AI Act deployer obligations · CMS 0057-F (prior-auth workflow).

05

Fragmented evidence collection

When a regulator asks for documentation, you should be able to produce it in 24 hours — not 24 days. Most programs can't, because evidence lives across SharePoints, email threads, team drives, and vendor portals.

WHERE IT'S REQUIRED

Every framework above expects a **single, on-demand evidence trail**.

06

No defined ownership

AI governance in healthcare is typically handed to whoever has relevant expertise or bandwidth. CISOs, compliance managers, AI directors all end up owning pieces. Without a designated owner and a shared platform, there's no accountability and no continuity.

WHERE IT'S REQUIRED

ISO 42001 & HITRUST require documented roles and a defined accountability model.

SECTION 03 · DIAGNOSTIC

Five questions you should be able to answer without a fire drill.

If any of these don't have an immediate answer, you have a compliance gap. Use this as the working checklist before your next audit cycle.

01

What AI systems are currently in production, including every vendor-embedded tool?

If the answer involves a spreadsheet someone stopped updating six months ago, you don't have an inventory.

02

When was each system last assessed for bias, accuracy, and regulatory alignment?

Joint Commission and ISO 42001 expect ongoing evidence, not a point-in-time approval from 18 months ago.

03

If The Joint Commission requested documentation tomorrow, where would you find it?

An evidence trail should resolve to a single link, not a 24-day scramble across SharePoints, email threads, and vendor portals.

04

Who is responsible for the governance of each AI system post-deployment?

"Whoever raised their hand" isn't an answer ISO 42001 or HITRUST will accept. Both regulations require documented roles and a defined accountability model.

05

How long does it take your team to evaluate a new AI vendor?

If the honest answer is "hours or days", you can't scale to the 20+ AI vendor relationships a typical mid-size health system already manages.

SECTION 04 · FRAMEWORK

Compliance-ready AI Governance isn't a policy document. It's an operational system.

BEFORE · FRAGMENTED GOVERNANCE

An auditor asks one question. Three weeks pass.

THE ASK

"Show oversight for AI in clinical decision support."

DAY 01	Email AI/Data lead	⊗
DAY 04	Pull HIPAA review.docx — v2, stale	⚠
DAY 09	SOC2 SharePoint — access denied	🔒
DAY 15	Vendor portal login — BAA expired	⚠
DAY 24	**Partial bundle delivered —** drift evidence missing	✕

24 days. Incomplete answer. Risk reviews are point-in-time. Vendor approvals never re-checked.

AFTER · CENTRALIZED GOVERNANCE

Same question. One link. Same afternoon.

THE ASK

"Show oversight for AI in clinical decision support."

Evidence bundle · CDS systems — 7 systems

Scope	All AI influencing clinical decisions
Risk tier	All classified · 5 high · 2 medium
Owners	Tagged · current as of 2026-05-08
Monitoring	Live · 0 drift alerts open
Last assessed	Q1 2026 · bias, accuracy, vendor

Export for Joint Commission — PDF + appendix · 1 click

Hours, not weeks. Same question, same export, every time the regulator asks.

COMPONENT 01

Centralized visibility

A single, continuously updated inventory of every AI system — internal, co-developed, vendor, and AI features embedded in existing platforms. Each entry: risk class, regulatory obligations, owner, current status.

COMPONENT 02

Automated, risk-tiered workflows

Low-risk use cases move through streamlined approval lanes. High-risk ones get the scrutiny they require, with documented evidence at each stage. No more queue stalling for three months.

COMPONENT 03

Continuous monitoring & audit-ready docs

Post-deployment monitoring tracks drift, flags risk, and triggers re-review when systems change. When Joint Commission, ISO auditors, or your board ask for documentation, the answer is a link — not a scramble.

SECTION 05 · THE VENDOR PAIN POINT

Healthcare orgs buy software. AI Governance hasn't caught up.

Unlike organizations that tend to build AI internally, healthcare organizations are comfortable purchasing new systems. The downside to this approach is that you don't control the training data, the bias methodology, or the update cadence.

THE VENDOR MATH, MID-SIZE HEALTH SYSTEM

20+
Vendor relationships at a typical mid-size health system that employ AI in their solutions

×

3 hrs
Is the average amount of time it takes a team to manually evaluate a single vendor

=

60+ hrs / yr
on vendor triage as tools update or onboard — with no standardized scorecard, no reusable artifacts, no systematic re-evaluation cadence.

FRAMEWORK	WHAT IT ASKS OF YOU FOR VENDOR AI
HIPAA	BAAs and documentation of PHI handling for every AI system processing protected health information.
Joint Commission	Evidence of oversight for AI-assisted clinical decisions, including third-party tools embedded in the EHR.
CMS 0057-F	Payers are accountable for what vendor AI produces in the prior-authorization workflow, including 72-hour decision timelines and denial documentation, even if the AI was procured, not built.
ISO 42001	Supplier management requirements for AI systems within the management system scope.
EU AI Act	Deployer obligations for high-risk third-party AI systems. You bear responsibility even if you didn't build it.

SECTION 06 · THE BENEFIT

The orgs building governance for compliance are getting an unexpected benefit. Their AI programs move faster.

Once the compliance floor is in place, the shift from “why you must” to “what you gain” happens almost on its own. Four things start to compound.

Risk tiering clears the queue

When low-risk use cases have an accelerated approval lane, the backlog clears. Programs stuck in 90-day review cycles complete in weeks. The teams who were losing appetite because of long review cycles are now staying involved and using the systems they requested.

The board conversation changes

Instead of fielding questions about AI risk, AI leaders proactively demonstrate oversight. The governance dashboard becomes evidence of control, compliance, and program maturity instead of a hand-wave at process.

Early movers get a window

Only **19%** of organizations have fully implemented AI governance frameworks. The orgs that get there first earn a window of competitive advantage in deployment speed, in vendor negotiations, and in stakeholder trust before the rest of the market catches up.

30–40%

THE DOWNSTREAM OUTCOME

Organizations with structured AI governance report 30–40% faster AI deployment timelines not in spite of compliance, but *because* a tiered, automated process replaces the queue.

SECTION 07 · THE SOLUTION

Not a GRC tool, a spreadsheet, or a framework document.

Each capability below maps to one of the six gaps on page 5. The order matters — visibility before workflow, workflow before monitoring, monitoring before audit.

CLOSES GAP 01 **AI Inventory & Registry**	Centralized, continuously updated inventory across every AI system including internally built, co-developed, and third-party vendor tools, especially AI features embedded in platforms you already own.
CLOSES GAP 02 **Automated risk tiering**	Low / medium / high classification with intelligent routing. The right level of scrutiny that is automatically applied so low-risk use cases stop blocking the queue, and high-risk ones get full review.
CLOSES GAPS 01, 03, 05 **Policy Intelligence & Compliance Packs**	Pre-built, configurable frameworks for Joint Commission, HIPAA, ISO 42001, HITRUST, CMS 0057-F, and the EU AI Act. Map your AI systems to controls without starting every program from scratch.
CLOSES GAP 04 **Vendor Assessment Module**	Out-of-the-box vendor scorecards and ongoing monitoring. Scale third-party AI oversight without scaling your headcount, with re-evaluation triggers whenever a vendor's risk profile changes.
CLOSES GAP 05 **Audit-ready documentation**	Continuous evidence collection across every AI system and workflow. Exportable on demand for when The Joint Commission, ISO, or your board asks. The response is a link, not a scramble.
CLOSES GAP 06 **Advisory Services**	A 6–12 week engagement that builds the program foundation — risk classification, operating model, governance workflows — before or alongside platform implementation. So the platform lands on an operationalized foundation, not aspiration.

Deployment flexibility

SaaS and self-hosted / on-premise options for organizations with strict PHI data residency requirements.

Built with healthcare leaders

Frameworks, workflows, and language tested with AI Leaders and CISOs across health systems, payers, and healthcare technology companies.

• Three steps before your next audit cycle

Stop looking for a roadmap. Get started today.

These three steps are the minimum prep for any audit or certification engagement. None requires a platform decision. All three sharpen the case for one.

01

Complete your AI inventory

List every AI system currently in use like internal models, third-party tools, and AI features embedded in existing platforms. Include who owns each one and whether it handles PHI. This is the prerequisite for everything else.

02

Apply risk tiers

Not all AI carries the same compliance burden. Classify each system as low, medium, or high risk based on clinical impact, PHI exposure, and applicable regulations. High-risk systems get prioritized for full governance buildout. Low-risk systems get a streamlined lane.

03

Identify your evidence gaps

For each high-risk system, assess: Is there a documented risk assessment? Ongoing monitoring? Vendor evaluation artifacts? A designated owner? These gaps are your compliance exposure and your immediate action list.

NEXT STEP

See how Credo AI can close these gaps for your healthcare AI program.

Book a 30-minute slot with a healthcare AI governance specialist that is tailored to your current frameworks and deadlines.

Book a working session →