curate partners

2026 Edition

Navigating AI Compliance in Healthcare & Financial Services: Patterns From the Field

AI Regulatory Landscape Brief for Healthcare & Financial Services

Vijay Kukreja, Chief Strategy Officer, Curate Partners
Key Contributions from Curate Data & AI Experts:

- Jonathan Pulliza

Executive Summary

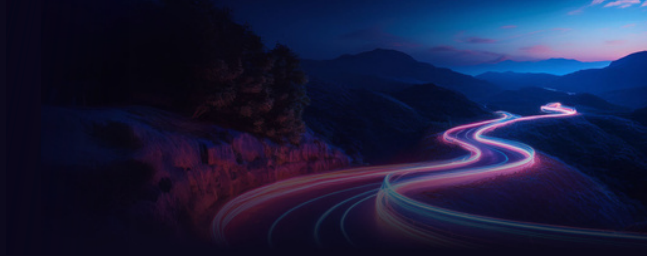

This brief captures what Curate Partners has seen firsthand across AI programs in healthcare and financial services during Q4 2025 and Q1 2026 — practitioner-level insights shaped by architecture reviews, governance workshops, and deployment readiness assessments where compliance obligations meet operational reality.

Most organizations are not waiting for new AI-specific legislation before they deploy. But the ones moving fastest are the ones who realized the rules already exist — they just haven't been mapped to AI yet.

Key findings:

Your existing regulations already govern AI. HIPAA, GLBA, SR 11-7, and FINRA rules were not written for large language models but they apply fully. Organizations that treat AI as a regulatory greenfield are accumulating compliance debt from day one.

The "black box" defense is dead in both sectors. FINRA, the CFPB, and the FDA all require explainability. If your AI cannot articulate why it made a decision, you have a compliance violation not a technology limitation.

Audit trails are the universal non-negotiable. Across every regulatory body we examined (including HHS, SEC, OCC, FINRA, the Joint Commission) the single most consistent requirement is traceability. If you cannot reconstruct what the AI did, why, and with what data, you are not ready for production.

The LLMOps lifecycle creates new compliance surface area. Prompt management, RAG pipelines, fine-tuning, and human feedback loops all introduce regulatory touchpoints that traditional MLOps frameworks do not address.

The bottom line: The regulatory frameworks governing AI in healthcare and financial services are not emerging — they are already here, embedded in existing law. The organizations moving fastest are the ones who mapped these obligations before they built, not after.

The State of AI Regulation: What We're Seeing

Over the past two quarters, Curate Partners has assessed AI deployment readiness across healthcare payers, financial institutions, and organizations operating in both verticals. This brief distills the regulatory patterns we observed firsthand — not from legal abstracts or policy summaries, but from the rooms where architecture teams ask "Can we do this?" and compliance teams answer "Not like that."

The central finding is this: the organizations most paralyzed by AI compliance are the ones waiting for new regulation and explicit guidance. The organizations making progress recognized that existing frameworks — HIPAA, GLBA, SR 11-7, FDA guidance, FINRA rules — already define the boundaries. The work is not inventing new governance. It is mapping existing obligations to new technology.

Core Insight

Regulated industries do not need new AI laws to know what is required. They need disciplined mapping of existing compliance frameworks to the AI development lifecycle — from problem definition through production monitoring.

Pattern 1:
Your Existing Regulations Already Govern AI

The most common misconception we encounter: organizations believe AI operates in a regulatory gray area, and that compliance requirements will become clear "once the government catches up." This is incorrect — and dangerously so.

What We Observed

At a healthcare payer, an innovation team had been piloting an LLM for member communication — but progress had stalled for months. The team assumed that existing regulations like HIPAA hadn't yet been interpreted for AI, and they were convinced that patient safety and compliance organizations would shut them down the moment the pilot surfaced. So they waited — for new guidance, for clearer rules, for someone to tell them it was safe to proceed. When we mapped the data flows and walked through HIPAA's Security Rule requirements alongside them — risk analysis, Business Associate Agreements, audit logging, minimum necessary data use — the team realized the compliance path was already defined. The rules that governed their EHR systems and claims databases applied equally to the LLM. Once that mapping was in place, the pilot moved forward in weeks — with compliance as an enabler, not a blocker.

At a financial institution, a trading desk deployed an AI assistant for research summarization. No one flagged that SEC Rule 17a-4 requires retention of records related to trading decisions — including the AI's outputs and the data it relied on. The tool had been running for weeks with no logging.

In both cases, the issue was the same: teams treated AI as a new category that existing rules hadn't addressed. In reality, the rules applied from the moment the system touched regulated data.

The Regulatory Baseline

Sector	Core Regulations	What They Require for AI
Healthcare	HIPAA / HITECH, GINA, state laws (CMIA, CCPA, TX Medical Privacy Act)	Risk analysis for all systems with ePHI, access controls, encryption, breach notification, minimum necessary data use
Financial Services	GLBA, SEC Reg S-P, NYDFS Cybersecurity Regulation	Privacy notices, consumer opt-out rights, Administrative/technical/physical safeguards, Data security obligations

FOR CIOs AND GENERAL COUNSEL

AI does not get a grace period — but it does not require one either. The moment an LLM processes protected health information or customer financial data, existing regulations like HIPAA, GLBA, and SR 11-7 apply fully. That same regulatory clarity means you already have the framework to move forward confidently. Map your obligations, engage the right expertise, and treat AI deployment as a compliance event — not something to defer until new rules arrive.

Pattern 2:
The "Black Box" Defense Is Dead

Across both healthcare and financial services, we observed a consistent regulatory expectation: if your AI cannot explain its decisions in terms humans can evaluate, it is not ready for production.

What We Observed

Across financial services and healthcare engagements, we consistently see teams experimenting with LLM-based tools without clear frameworks in place to validate the reasoning behind the model's outputs — or even the reliability of those outputs. The enthusiasm to deploy is there, but the rigor to explain and defend the AI's decisions has not caught up.

Consider what this looks like in practice. A financial institution deploys an LLM-based credit model. During an internal compliance review, the team is asked to demonstrate how the model produces adverse credit decisions — a requirement under ECOA Regulation B, which mandates specific, articulable reasons for any denial. The model cannot provide them. The team assumed that the AI's statistical sophistication compensated for its lack of interpretability. Under Regulation B, it does not. Without explainability built into the model's architecture, the institution faces a fair lending exposure before the tool ever reaches a regulator's desk.

In healthcare, the pattern mirrors. The FDA's CDS guidance expects that clinicians can independently review the basis of AI recommendations — meaning the AI must provide information in medically accepted terms, not opaque confidence scores. The Joint Commission requires demonstration that AI tools were evaluated for usability and accuracy before clinical deployment.

The Explainability Requirement Across Sectors

Regulator	Requirement	Consequence of Non-Compliance
CFPB / ECOA	Specific reasons for adverse credit decisions	Fair lending violations; enforcement action
FDA CDS Guidance	Clinician must independently review AI rationale	Reclassification as regulated medical device
FINRA Notice 24-09	"Black box" is not a defense for non-compliance	Supervisory violations; enforcement
OCC / Fed SR 11-7	"Effective challenge" — staff must be able to question AI assumptions	MRM deficiency findings
NIST AI RMF	Transparency and documented performance metrics	Governance gaps in audit findings

FOR ALL C-SUITE LEADERS

Explainability is not a nice-to-have. It is a regulatory requirement in both healthcare and financial services. If your AI vendor says "the model is too complex to explain," that is not a feature — it is a deployment blocker. Build interpretability into your architecture from day one, or plan to retrofit it at significant cost. And interpretability alone is not enough — you also need the operational processes to validate and deliver explanations consistently, which may require dedicated product features, tooling, and workflow development to make explainability repeatable at scale.

Pattern 3:
Audit Trails Are the Universal Non-Negotiable

If there is one requirement that unifies every regulatory body across both sectors, it is this: traceability. Every regulator expects that AI-driven decisions can be reconstructed, reviewed, and attributed.

What We Observed

Across our engagements, audit trail readiness was the single most common gap. Organizations had deployed AI tools — chatbots, summarization engines, decision support systems — without clear visibility into what the AI processed, what it produced, or what data it relied on.

In healthcare, the audit trail requirement extends well beyond the Electronic Health Record (EHR). Consider a payer deploying an AI tool to help plan members estimate out-of-pocket costs before scheduling a care event. To generate those estimates, the AI pulls from claims history, benefit design, provider contract rates, and accumulated deductible data — all of which constitute protected health information under HIPAA. Every query the member makes, every data source the AI accesses, and every cost estimate it produces must be logged and attributable. If the AI surfaces an inaccurate estimate and a member makes a care decision based on it, the payer needs to be able to reconstruct exactly what data the model used, what logic it applied, and what it returned. Without that audit trail, the organization has a compliance gap and a member trust problem simultaneously.

In financial services, the requirements are equally explicit. SEC Rule 17a-4 requires retention of trading decision logs matched with the model used. FINRA expects that all AI-driven decisions or communications can be traced and reviewed. If an AI chatbot talks to customers, transcripts must be retained and supervised like any human representative's conversation.

Audit Requirements by Regulatory Body

- HIPAA / ONC — Audit logs for all ePHI access by AI systems; auditable within EHR logging for legal discovery
- FDA Post-Market Surveillance — Manufacturers must report malfunctions or adverse events from regulated AI
- HL7 / FHIR — AI decisions must be traceable, with outputs labeled and tracked in health records
- SEC 17a-4 / 204-2 — AI inference logs, training data lineage, and model parameters are regulated records
- FINRA — All AI-driven communications and decisions reconstructable for supervisory review
- OCC — Incident response, fallback procedures, and system resilience documentation required

FOR CIOs AND CISOs

Before any AI system reaches production, verify four things: (1) every AI interaction with regulated data is logged, (2) logs include the model version, input data, and output produced, (3) logs are retained for the period required by your sector's recordkeeping rules, and (4) your compliance and oversight teams have the operating model, tooling, and access they need to analyze this data and surface patterns or issues quickly. If you cannot reconstruct what the AI did six months from now, you are not compliant today.

Pattern 4:
The LLMOps Lifecycle Creates New Compliance Surface Area

Traditional MLOps governance was built for predictive models with deterministic outputs. Large language models introduce fundamentally different operational patterns — prompt engineering, retrieval-augmented generation, context window management, human feedback loops — each of which creates new regulatory touchpoints.

What We Observed

Across our engagements, we found that even organizations with mature ML governance programs were encountering new areas of regulatory exposure when applying those same frameworks to LLM deployments. The gaps were not a reflection of weak governance — they were a reflection of how fundamentally different LLM operations are. Three stages stood out as the highest-risk areas:

- Data Engineering & Management — RAG pipelines ingest unstructured documents that may contain PHI or regulated financial data. Prompt management requires versioning and lineage tracking. Organizations were treating prompts as informal artifacts rather than governed data inputs.
- Model Development & Adaptation — Fine-tuning and prompt refinement change model behavior in ways that require re-validation. Organizations were iterating on prompts without documenting changes or assessing whether adaptations introduced new risks.
- Continuous Monitoring & Feedback — LLMs require hallucination detection, semantic drift monitoring, and human feedback integration that traditional drift detection tools don't address. Organizations lacked specialized monitoring for fabricated outputs.

Mapping the LLMOps Lifecycle to Regulatory Requirements

Lifecycle Phase	Healthcare Regulations	Financial Services Regulations	Cross-Industry Recommendation
Problem Definition & Governance	HIPAA Security Rule (upfront risk analysis); FDA GMLP & CDS (intended use / risk classification); NIST AI RMF (governance structures)	SR 11-7 (model purpose, scope, assumptions); OCC (inventorying / tiering); FINRA (governance like any production system)	Conduct a formal risk assessment and define intended use, scope, and governance ownership before development begins. Inventory and tier every AI use case by regulatory exposure.
Data Engineering & Management	HIPAA (minimum necessary; no uncontrolled PHI retention); ONC (standard terminologies, auditable access); HL7 FHIR (output labeling)	SR 11-7 / OCC (data lineage, quality controls); SEC 17a-4 / 204-2 (data as regulated records); CFPB / ECOA (bias in training data)	Enforce minimum-necessary data access, maintain end-to-end data lineage for every input (including RAG sources and prompts), and treat all AI-processed data as regulated records subject to retention and bias review.
Model Development & Adaptation	FDA GMLP (controlled development, validation of changes); ISO 14971 (hazard analysis); Joint Commission (validation before deployment)	SR 11-7 (conceptual soundness reviews); OCC (independent review of GenAI adaptations); FINRA (no "black box" defense)	Require independent validation and documented change control for every model adaptation — including prompt changes and fine-tuning — with hazard and risk analysis before any update reaches production.
Evaluation & Transparency	FDA CDS (post-market surveillance); ONC (explainable to clinicians); NIST AI RMF (transparency, bias evaluation)	SR 11-7 (validation, continuous monitoring); OCC ("effective challenge"); CFPB / ECOA (explainability for adverse decisions)	Build explainability into the model architecture from day one — ensure every AI output can be reviewed, challenged, and explained in domain-appropriate terms by qualified human reviewers.
Deployment & Serving	HIPAA Security Rule (access controls, encryption, audit logging); FDA (labeling, intended use); Joint Commission (workflow impact)	FINRA Rule 2210 (supervision of AI communications); SEC (inference logs as records); OCC (incident response, resilience)	Deploy with production-grade audit logging (model version, inputs, outputs), role-based access controls, encryption, and a documented incident response plan — treating every AI interaction as a supervised, reconstructable event.
Continuous Monitoring & Feedback	FDA (post-market surveillance for malfunctions); HIPAA (continuous access monitoring); NIST AI RMF (feedback loops, drift detection)	SR 11-7 (ongoing monitoring, re-validation); FINRA (continuous supervision); SEC / CFPB (ability to halt harmful AI behavior)	Implement continuous monitoring for hallucination, semantic drift, and adverse outcomes — with automated alerting, human-in-the-loop feedback integration, and a documented kill switch to halt harmful AI behavior immediately.

FOR CTOs AND AI PLATFORM LEADS

Your LLMOps pipeline is a compliance pipeline. Every stage — from RAG ingestion to prompt versioning to human feedback — introduces regulatory obligations that traditional MLOps playbooks don't cover. Build compliance checkpoints into each lifecycle phase, not as a post-deployment audit.

What the Field Taught Us: The Beliefs That Hold AI Programs Back

	The Misconception	The Reality	What to Do About It
1	"There are no AI regulations yet — we'll figure out compliance later."	HIPAA, GLBA, SR 11-7, FINRA, and FDA guidance already govern AI the moment it touches regulated data or production workflows. There is no grace period.	Map every AI use case to the specific regulations that already apply — HIPAA, GLBA, SR 11-7, FINRA — and document your compliance posture before the first model reaches production.
2	"Our model is too complex to explain— regulators will understand."	FINRA, the CFPB, and the FDA all require explainability. "Black box" is explicitly cited as an unacceptable defense. Interpretability must be designed in, not apologized for.	Require interpretability as an architecture criterion from day one. Mandate that every model produce human-reviewable rationale for its outputs, and build the tooling and processes to deliver explanations consistently at scale.
3	"We'll add audit trails once we're in production."	Every major regulator — HHS, SEC, FINRA, OCC, the Joint Commission — requires traceability from day one. Retrofitting audit trails is orders of magnitude harder than building them in.	Instrument your AI pipeline to log every interaction — model version, input data, and output — from the first deployment. Ensure logs meet your sector's retention requirements and that compliance teams have tooling to query and review them.
4	"Our existing ML governance covers LLMs."	LLMs introduce prompt management, RAG pipelines, hallucination risk, and non-deterministic outputs that traditional MLOps governance does not address. Extend your framework or accept gaps.	Extend your model risk management framework with LLM-specific controls: prompt versioning and change management, RAG data lineage, hallucination detection, and semantic drift monitoring. Treat each as a governed compliance checkpoint

About Curate Partners

Curate Partners is a strategic consulting and talent firm specializing in AI adoption, data governance, and digital transformation for regulated industries. We operate as fractional C-suite leaders embedded in our clients' organizations, providing cross-industry pattern recognition that single-company executives cannot access.

Ready to Map Your AI Compliance Landscape?

The regulatory patterns in this brief are not theoretical — they are the obligations we help organizations navigate every week. Whether you are deploying your first AI use case in a regulated environment or scaling an existing program, Curate Partners can help you:

- Map your regulatory obligations to every phase of the AI development lifecycle — before you build, not after
- Classify your AI use cases against FDA, FINRA, and banking regulatory frameworks to determine your compliance path
- Design audit trail architectures that satisfy cross-regulatory requirements from day one
- Extend your model risk management program to cover LLM-specific risks: prompt sensitivity, hallucination, and non-deterministic behavior

Start the conversation.

Visit curatepartners.ai to explore our AI services, or reach out directly to schedule a complimentary AI Readiness Assessment with our team.