

DOI: https://doi.org/10.48009/3_iis_2026_144

From concerns to frameworks: Mapping 2021 healthcare professional concerns about AI to post-2021 U.S. governance responses

Barry Humphrey, Marymount University, bah03060@marymount.edu

Xiang Liu, Marymount University, xliu@marymount.edu

Abstract

A 2021 quantitative study of 281 U.S. healthcare professionals and HIPAA officers identified six recurring concerns about artificial intelligence (AI), HIPAA, and the privacy and security of protected health information (PHI): regulatory gaps, HIPAA's perceived obsolescence, vendor accountability outside HIPAA's covered-entity perimeter, re-identification risk, limited workforce AI literacy, and lack of operational guidance. This paper uses a structured document analysis, grounded in regulatory and adaptive-governance theory, to map those concerns to post-2021 U.S. governance responses, including NIST AI RMF, ONC HTI-1, HHS OCR actions, the proposed HIPAA Security Rule update, FDA guidance, FTC enforcement, and selected state laws. Using a four-point rubric (addressed, partially addressed, not yet addressed, contested) with a partial-code taxonomy that distinguishes failures of scope, bindingness, accountability, and enforceability, the analysis finds none of the six concerns is fully resolved by the binding U.S. instruments as of mid-2026. Five are partially addressed through layered, jurisdiction-specific mechanisms; re-identification remains materially unresolved; and several responsive instruments are now contested through pending finalization or non-enforcement. The findings characterize the U.S. response as a regulatory mosaic produced by institutional layering rather than a unified regime, and they highlight continuing gaps in accountability, coordination, and implementation as AI adoption in healthcare accelerates.

Keywords: artificial intelligence, HIPAA, healthcare data privacy, protected health information, AI governance, regulatory gap analysis

Introduction

The integration of artificial intelligence (AI) and machine learning (ML) into healthcare has accelerated rapidly in recent years. Diagnostic imaging algorithms, predictive analytics for clinical deterioration, ambient documentation tools, generative AI assistants, and consumer-facing wearables now process protected health information (PHI) at scales and through architectures that the Health Insurance Portability and Accountability Act of 1996 (HIPAA), as amended by the 2009 HITECH Act and the 2013 Omnibus Rule, did not contemplate. The U.S. Food and Drug Administration (FDA) has authorized a large and growing number of AI/ML-enabled medical devices; its public AI/ML-enabled Medical Device List reported 882 authorizations through March 31, 2024, and 1,016 authorizations through December 2024 (FDA, 2024a; Exponent, 2024). The American Medical Association's Augmented Intelligence Research surveys show physician self-reported use of AI in practice increasing from 38% in 2023 to 66% in 2024 and to 81% in 2026, indicating that AI adoption among physicians has more than doubled since the initial 2023 survey (American Medical Association, 2026).

The 2021 study that provides the empirical foundation for this paper surveyed 281 U.S. healthcare professionals, including 179 clinical and health information technology professionals and 102 HIPAA security, privacy, and compliance officers, and identified six recurring concerns about AI's impact on HIPAA and patient data privacy (Humphrey, 2021). Practitioners reported that AI was outpacing the regulatory architecture, that HIPAA's covered-entity perimeter excluded the technology companies most likely to handle health data, that machine learning could re-identify ostensibly de-identified data, that the workforce lacked AI literacy, and that no clear operational guidance existed for clinical AI use. Rather than drawing conclusions about current law, the 2021 study recommended updates to HIPAA, the expansion of the covered-entity definition to include AI developers, and the development of a federal regulatory framework.

Since 2021, the U.S. governance landscape has changed substantially. The National Institute of Standards and Technology (NIST) released the AI Risk Management Framework 1.0 in 2023 and the Generative AI Profile (NIST AI 600-1) in 2024, both voluntary management-system frameworks rather than binding rules (NIST, 2023; 2024). The Office of the National Coordinator for Health Information Technology (ONC) finalized the Health Data, Technology, and Interoperability (HTI-1) rule in January 2024, establishing binding transparency requirements for Predictive Decision Support Interventions delivered through certified health IT modules (ONC, 2024). The Department of Health and Human Services Office for Civil Rights (HHS OCR) issued the Section 1557 Final Rule extending nondiscrimination obligations to patient care decision support tools used by recipients of HHS financial assistance (HHS OCR, 2024a), and proposed the first major HIPAA Security Rule update in more than a decade (HHS OCR, 2025a). The FDA finalized guidance on Predetermined Change Control Plans (PCCPs) for AI-enabled device software functions (FDA, 2024b). The Federal Trade Commission (FTC) brought a series of enforcement actions against non-HIPAA health-data actors (FTC, 2023; 2024a; 2024b). State legislatures enacted consumer-health-data laws in Washington, Nevada, Connecticut, and Maryland, while California enacted AB 3030 governing disclosure of generative AI in healthcare communications.

In light of this evolving policy and regulatory landscape, this article asks whether and to what extent these post-2021 U.S. instruments have addressed the concerns that healthcare professionals voiced in 2021. This paper bridges the 2021 practitioner concerns and the evolving U.S. governance landscape from 2023 to 2026, using a structured mapping approach to identify which concerns have been addressed, which remain only partially addressed, and where key governance gaps persist. Accordingly, two research questions guide the study:

RQ1: To what extent have the concerns identified in 2021 about AI, HIPAA, and patient data privacy been addressed by post-2021 U.S. governance responses?

RQ2: What governance gaps remain, and what are the implications for policy and practice?

Background and Literature Review

The 2021 Empirical Baseline

The empirical anchor for this paper is a 2021 quantitative study of 281 U.S. healthcare and HIPAA professionals (Humphrey, 2021). The study drew on two surveyed groups, 179 clinical and health information technology professionals and 102 HIPAA security, privacy, and compliance officers, and identified six recurring concerns about AI's impact on HIPAA and protected health information (PHI): (1) the absence of binding regulation for AI in healthcare; (2) HIPAA's perceived obsolescence relative to AI

capabilities; (3) AI developers and vendors falling outside HIPAA's covered-entity and business-associate perimeter; (4) machine learning's capacity to re-identify de-identified data; (5) limited AI literacy among healthcare professionals; and (6) the absence of operational guidance for clinical AI use. These six categories form the analytical spine of the present paper. They are treated as a practitioner-validated demand signal and used to assess, concern by concern and instrument by instrument, how the post-2021 U.S. governance architecture has responded.

A Governance-Accountability Lens Grounded in Regulatory Theory

The 2021 study used Protection Motivation Theory (PMT) because it was designed to examine how healthcare professionals perceived AI-related threats and their ability to respond to them. The present paper analyzes regulatory and governance documents rather than practitioner perceptions, and a behavior-motivation theory is ill-suited to that object. It therefore adopts a governance-accountability lens, situated within three established strands of regulatory scholarship that together explain why the post-2021 response took the shape it did.

First, the pacing problem—the structural lag between fast-moving technology and slower legal-ethical oversight (Marchant, Allenby, & Herkert, 2011; Bennett Moses, 2011)—predicts that AI would outrun a statute drafted in 1996 and would be governed, if at all, through whatever authorities already existed. Second, theories of polycentric and decentered regulation (Black, 2001) predict that, absent a single integrating statute, governance will be distributed across multiple agencies acting through their own jurisdictional mandates. Third, institutional layering (Mahoney & Thelen, 2010) names the precise mechanism observed here: rather than replacing HIPAA, regulators have layered new AI obligations onto pre-existing civil rights, certification, device, and consumer protection authorities. Against these descriptive theories, adaptive and anticipatory governance (Sarewitz, 2011) supplies the normative yardstick—whether the architecture has the feedback, coordination, and revision capacity to keep pace with the technology.

Within this frame, the governance-accountability lens is the operational instrument. For each post-2021 instrument, the analysis asks three questions: does it create a binding obligation, does it assign accountability to a defined actor, and is it enforceable through a defined mechanism? These tests operationalize the layering-and-pacing diagnosis into a codable rubric. The question is therefore not whether practitioners now feel less concerned, but whether the legal and standards architecture has changed in ways materially responsive to their 2021 concerns.

The Post-2021 U.S. Governance Landscape

The corpus distinguishes binding rules, proposed rules, and voluntary frameworks.

Voluntary frameworks. The NIST AI Risk Management Framework 1.0 (NIST, 2023) and its Generative AI Profile, NIST AI 600-1 (NIST, 2024), provide a voluntary, outcome-oriented governance scaffold organized around the Govern, Map, Measure, and Manage functions. NIST AI RMF is not a regulation; it imposes no binding obligations and has no enforcement mechanism. Its influence is normative: it functions as a reference architecture that regulators, procurement officials, plaintiffs, and certifying bodies can increasingly point to as a benchmark for reasonable AI governance.

FDA guidance documents. FDA guidance, including the December 2024 final guidance on PCCPs for AI-enabled device software functions (FDA, 2024b) and the January 2025 lifecycle draft guidance (FDA, 2025), is non-binding by the agency's own definition. Guidance reflects the FDA's current thinking, and manufacturers may use alternative approaches that satisfy the underlying statutory and regulatory requirements. In practice, however, such guidance is highly influential because it shapes the expectations against which premarket submissions are evaluated.

Binding rules in force. The ONC HTI-1 Final Rule (89 Fed. Reg. 1192, January 9, 2024) created binding requirements for developers of certified health IT through the Decision Support Interventions (DSI) certification criterion at 45 CFR § 170.315(b)(11), with maintenance-of-certification obligations applicable as of January 1, 2025 (ONC, 2024). The DSI criterion includes source-attribute disclosure for Predictive DSIs, intervention risk-management practices, and feedback mechanisms, and represents the first U.S. regulatory action to attach binding transparency-related obligations to predictive AI embedded in certified health IT. Separately, the HHS OCR Section 1557 Final Rule (89 Fed. Reg. 37,522, May 6, 2024) imposed binding nondiscrimination obligations on recipients of HHS funding through 45 CFR § 92.210, which prohibits discrimination through the use of patient care decision support tools, defined to include both automated tools such as AI/ML algorithms and non-automated tools such as clinical flowcharts. The general nondiscrimination obligations took effect July 5, 2024, and the affirmative duty under § 92.210 had a delayed applicability date approximately May 1, 2025 (HHS OCR, 2024a; Mintz, 2024). Importantly, in May 2025 HHS announced that it would not enforce certain provisions, including § 92.210(b)–(c) governing patient care decision support tools, pending ongoing litigation; the status of the affirmative AI duties is therefore now contingent rather than settled, a point developed in the Findings.

Portions of the rule unrelated to AI, specifically certain sex-discrimination provisions, remain subject to stays and injunctions in *Tennessee v. Becerra*, *Texas v. HHS*, and related litigation (HHS OCR, 2025c).

Proposed but not yet final. The HIPAA Security Rule Notice of Proposed Rulemaking (90 Fed. Reg. 898, January 6, 2025) would eliminate the distinction between required and addressable safeguards, mandate technology asset inventories and network mapping that would include AI assets processing ePHI, impose semiannual vulnerability scanning and annual penetration testing, and require rapid business-associate breach notification (HHS OCR, 2025a). The comment period closed March 7, 2025, with approximately 4,700 comments received. Although finalization remained on OCR's regulatory agenda for May 2026, that target has since passed with no final rule issued, and OCR has not confirmed whether or when one will issue; a deregulatory administration and industry calls for rescission leave its survival uncertain (Alston & Bird, 2025; HIPAA Journal, 2026). Until finalization, the NPRM creates no binding obligations.

FTC enforcement and state law. The FTC has used the Health Breach Notification Rule (16 CFR Part 318, amended April 2024) and Section 5 of the FTC Act to discipline non-HIPAA health-data actors, including GoodRx, BetterHelp, Easy Healthcare/Premom, Cerebral, Monument, and Rite Aid (FTC, 2023; 2024a; 2024b), building a case-by-case perimeter around health data outside HIPAA's reach. The April 2024 Health Breach Notification Rule amendments codified the rule's application to non-HIPAA health apps and trackers. State activity supplements the federal landscape: Washington's My Health My Data Act and Nevada SB 370 took effect March 31, 2024, with Washington providing a private right of action; Maryland's Online Data Privacy Act (MODPA) took effect October 1, 2025; and California AB 3030 (2024) requires generative-AI disclosure in covered healthcare communications. In this paper, state activity is treated as supplementary evidence within the U.S. matrix rather than as separate mapping rows. The European Union AI Act (Regulation (EU) 2024/1689) and ISO/IEC 42001:2023 are reserved for the Discussion as international comparators and are not part of the primary mapping matrix.

Presentation of the Main Contribution of the Paper

Recent research on healthcare AI governance has cataloged frameworks, surveyed national strategies, and proposed maturity models. Hussein et al. (2026), for example, reviewed 35 frameworks for AI implementation in healthcare published between 2019 and 2024 and proposed a five-level Healthcare AI

Governance Readiness Assessment. Legal and policy scholarship on HIPAA and AI has documented regulatory shortcomings (Greene, 2020; Hoffman, 2019; Rezaeikhonakdar, 2023; Tovino, 2024). What remains underdeveloped is a governance gap analysis anchored in empirically elicited practitioner concerns. This paper addresses that gap by using the six concerns identified in the 2021 study as a fixed reference frame for evaluating post-2021 U.S. governance responses, and in doing so makes four contributions. First, it provides a longitudinal bridge between practitioner concerns identified in 2021 and the evolving U.S. healthcare AI governance landscape from 2023 to 2026. Second, it offers a structured and replicable mapping method—including a partial-code taxonomy that exposes why instruments fall short—for assessing regulatory responsiveness as rules, guidance, and enforcement actions continue to evolve. Third, it identifies the structural gaps that remain, particularly those involving accountability, coordination, and the continuing limits of the covered-entity framework. Fourth, it offers a conceptual synthesis of healthcare AI governance that distills the findings into four interrelated dimensions: regulatory architecture, HIPAA coverage boundaries, operational guidance, and AI workforce literacy.

Methodology

Research Design

This study uses a structured document analysis (Bowen, 2009) operationalized as a regulatory gap analysis. Document analysis is appropriate when the research question concerns the textual and structural content of authoritative documents and their fit with an external evaluative criterion. It is also appropriate here because new primary data collection would not improve construct validity beyond the existing dataset (Humphrey, 2021) that established the baseline concerns. The unit of analysis is the concern–instrument pair. For each of the six 2021 concern categories, each post-2021 U.S. governance instrument is coded against a four-point rubric.

Document Corpus

The corpus comprises seven federal U.S. instrument categories and four state laws included as supplementary evidence within the matrix. Documents were retrieved from primary sources, including the Federal Register, agency websites, and state legislative websites, between January 2023 and mid-2026.

The federal instruments are: NIST AI Risk Management Framework 1.0 (January 2023) and NIST AI 600-1 (July 2024); the ONC HTI-1 Final Rule (89 Fed. Reg. 1192, January 9, 2024) and accompanying Federal Health IT Strategic Plan; the HHS OCR Section 1557 Final Rule (89 Fed. Reg. 37,522, May 6, 2024) and the OCR Dear Colleague Letter on Section 1557 and AI of January 10, 2025; the HHS OCR HIPAA Security Rule NPRM (90 Fed. Reg. 898, January 6, 2025); the FDA AI-enabled Device Software Functions guidance suite, including the Marketing Submission Recommendations for a Predetermined Change Control Plan final guidance (December 2024) and the Lifecycle Management and Marketing Submission Recommendations draft guidance (January 2025); and FTC enforcement actions in *In re GoodRx*, *In re BetterHelp*, *In re Easy Healthcare/Premom*, *In re Cerebral*, *In re Monument*, and *In re Rite Aid Corporation*, together with the 2024 Health Breach Notification Rule final amendments (16 CFR Part 318).

State activity is treated as supplementary contextual evidence rather than as a comprehensive fifty-state survey; the matrix includes selected illustrative state instruments as a single supplementary column: the Washington My Health My Data Act (RCW 19.373), Nevada SB 370 (NRS 603A.400 et seq.), the Maryland Online Data Privacy Act, and California Assembly Bill 3030 (2024). The European Union AI Act (Regulation (EU) 2024/1689) and ISO/IEC 42001:2023 appear in the Discussion as international comparators only.

Coding Rubric

Each concern–instrument pair receives one of four codes.

A — Addressed. The instrument creates a binding obligation that materially closes the 2021 gap, with a defined accountable actor and an enforcement mechanism. The operational tests are: (i) mandatory language such as “must” or “shall”; (ii) a named accountable actor; (iii) an identifiable enforcement or conformity-assessment mechanism; and (iv) substantive scope covering the concern as stated. Voluntary frameworks and non-binding guidance cannot receive an “A” code on their own.

P — Partially addressed. The instrument speaks to the concern but leaves at least one of scope, bindingness, accountability, or enforceability unresolved. Each “P” carries a subcode identifying the failing test: P-bind (no binding obligation; voluntary or guidance), P-scope (substantive coverage incomplete relative to the concern as stated), P-enf (an accountable actor exists but enforcement is absent, ex post, or not yet in force), and P-acct (no clearly defined accountable actor).

N — Not yet addressed. No instrument in the corpus contains a substantive provision targeting the concern, or all relevant provisions are voluntary aspirations or pending finalization without an enforceable obligation as of mid-2026.

C — Contested or contingent. The instrument once addressed the concern, but its status has become uncertain because of litigation, withdrawal, non-enforcement, leadership shift, pending finalization, or rescission. For each pairing, the matrix cites the relevant provision and provides a one- to two-sentence justification, with representative justifications given in the Findings. At the cell level, the code reports the most protective status available within that instrument’s own scope; at the aggregate level, the column reports whether the concern is resolved across the corpus as a whole.

Reliability

To address coder subjectivity, each concern–instrument pairing was independently coded by the primary coder and a second reviewer using the rubric, including the partial-code taxonomy. The coders compared results after the initial round, and disagreements were resolved through discussion by returning to the four operational tests: bindingness, accountable actor, enforceability, and substantive fit with the concern. Persistent disagreements were retained as range codes, such as “P/N,” with both interpretations explained where applicable. Cohen’s κ for the 42 pairings was 0.81, indicating substantial agreement.

Results/Findings

Table 1 presents the 6×7 mapping of the 2021 practitioner concerns against the post-2021 U.S. governance instruments. Each entry shows the code with its partial-code subtype and the most relevant provision. The aggregate column reflects the overall status of each concern across the full U.S. corpus.

Table 1. Mapping of 2021 Practitioner Concerns to Post-2021 U.S. Governance Instruments (mid-2026)

2021 Concern	NIST AI RMF	ONC HTI-1	HHS OCR §1557	HIPAA Sec. Rule NPRM	FDA PCCP	FTC + HBN R	State laws	Aggregate
C1. Lack of binding AI regulation	P-bind	A	C	C	P-bind	P-enf	P-scope	P

C2. HIPAA outdated for AI	—	A	C	C	—	A	P-scope	P
C3. Vendors as non-covered entities	P-bind	A	P-scope †	P-scope	P-bind	A	P-scope	P
C4. Re-identification of de-identified data	P-bind	P-scope	—	P-scope	—	P-scope	P-scope	N → P
C5. Workforce AI literacy	P-bind	P-scope	P-scope †	P-enf ‡	—	—	—	P
C6. Lack of clear operational guidance	P-bind	A	C	P-enf ‡	P-bind	P-enf	P-scope	P

Note. A = Addressed (binding, accountable, enforceable). P = Partially addressed, failing test shown: bind = non-binding; scope = coverage incomplete; enf = enforcement absent, ex post, or not yet live; acct = no defined accountable actor. N = Not yet addressed. C = Contested or contingent. — = does not bear materially on the concern. † Contingent on the May 2025 non-enforcement posture for § 92.210(b)–(c). ‡ NPRM provision responsive but not yet in force as of mid-2026. “A” cannot appear in the NIST or FDA columns; both are non-binding, so their ceiling is P-bind

The aggregate pattern across the U.S. corpus is striking. Of the six 2021 concerns, none is fully resolved by the binding U.S. instruments reviewed. Five are partially resolved (Concerns 1, 2, 3, 5, and 6). One remains materially unresolved (Concern 4, re-identification), coded N → P because FTC enforcement provides partial protection within the consumer-protection perimeter even though HIPAA’s de-identification methodology remains unchanged. Five cells now carry “C,” reflecting the unfinalized HIPAA Security Rule NPRM (Concerns 1 and 2) and the contingent enforcement status of the Section 1557 § 92.210 affirmative duties (Concerns 1, 2, and 6).

Finding 1: Broad Surface Area, Narrow Scope at Each Node

The five partial aggregate outcomes do not reflect regulatory inactivity. Rather, they reflect an instrument-by-instrument response in which AI has been addressed through preexisting statutory authority: Section 1557 operates through civil-rights authority, HTI-1 through certification authority, FDA PCCPs through device authority and non-binding guidance, and the FTC through Section 5 and the Health Breach Notification Rule. As institutional-layering theory predicts, AI in healthcare is now governed along the boundaries of each agency’s existing jurisdiction. A predictive model embedded in certified electronic health record software is comparatively well governed, whereas a similar model deployed through a non-HIPAA wellness application may be governed only if the FTC takes enforcement action. The 2021 respondents framed Concern 1 as a call for federal regulation of AI in healthcare; the 2026 answer is more fragmented, with the federal government responding in multiple ways but not through a unified regulatory framework.

Finding 2: The Covered-Entity Perimeter Remains Unresolved

Concern 3 that AI developers and vendors operate outside HIPAA’s covered-entity perimeter was clearly identified in the 2021 study and remains among the least resolved issues in the post-2021 architecture. The HIPAA Security Rule NPRM would tighten business-associate obligations and require covered entities to incorporate AI-vendor risk into their security risk analysis, but it does not redraw the definition of covered entities. FTC actions against GoodRx, BetterHelp, Cerebral, Monument, Premom, and Rite Aid create a de facto perimeter for non-HIPAA health-data actors, but only through ex post enforcement. State laws such as Washington’s My Health My Data Act and Maryland’s MODPA extend protection in some jurisdictions while also creating the kind of patchwork that 2021 respondents identified as problematic. As a matter of

legal structure, no federal U.S. instrument has redrawn the covered-entity perimeter identified by the 2021 study as a central concern.

Finding 3: Re-identification Remains the Weakest Area

The HIPAA Privacy Rule's two de-identification methods, Safe Harbor and Expert Determination, remain unchanged, and none of the reviewed instruments updates them. AI-era re-identification is best understood not as a single risk but as a family of attack pathways, each with a distinct relationship to those methods: (1) record-linkage attacks match de-identified records to external datasets on quasi-identifiers (Rocher et al., 2019); (2) membership inference determines whether a known individual's data was in a model's training set (Shokri et al., 2017); (3) model inversion or reconstruction recovers training features or approximate records from model outputs (Fredrikson et al., 2015); (4) attribute inference predicts undisclosed sensitive attributes; and (5) generative re-identification, in which large generative and multimodal models lower the cost and skill threshold for several of the above. HIPAA's methods were designed largely against pathway (1) and address (2)–(5) weakly or not at all. NIST AI 600-1 identifies model-inversion and data-integrity risks but is voluntary. The FTC's BetterHelp action offers the closest practical development, treating hashed identifiers as personally identifiable where the recipient could reasonably relink them—a contextual-risk standard that has not migrated into HIPAA. This remains the most materially unresolved technical concern in the current U.S. healthcare-specific regulatory framework.

Finding 4: Workforce AI Literacy Remains Weak

Across all seven U.S. instrument categories, every applicable entry for Concern 5 is coded “P,” and no U.S. instrument imposes a binding obligation for workforce AI literacy. NIST AI RMF recommends training but is voluntary (P-bind); the HHS OCR Dear Colleague letter of January 10, 2025, recommends training but is not binding; the HIPAA Security Rule NPRM proposes training requirements that remain pending (P-enf); and HTI-1's real-world testing and intended-use risk-management requirements apply to certified health IT developers, not to clinical end users (P-scope). This gap is best understood as a coping-capacity deficit: the formal architecture increasingly recognizes AI-related threats, but it imposes no binding obligation to build the workforce competence required to act on them, leaving response capacity dependent on voluntary uptake.

Finding 5: A Layered Architecture Rather Than a Unified Regime

The matrix shows that AI in U.S. healthcare is now governed through four overlapping layers: a device layer (FDA PCCPs and lifecycle guidance) for AI as a medical device; a certification layer (ONC HTI-1) for AI embedded in certified health IT; a civil-rights layer (Section 1557) for AI used by federally funded providers; and a consumer-protection perimeter (FTC Health Breach Notification Rule and state consumer-health-data laws) for AI outside the HIPAA universe. NIST AI RMF operates beneath these layers as voluntary governance scaffolding, and a finalized HIPAA Security Rule would add a stronger security layer for electronic PHI. What remains missing is a horizontal AI-in-healthcare regulation that would apply across actor types, sectors, and deployment contexts. In retrospect, this is the structural gap the 2021 respondents identified when they called for coordinated leadership.

Finding 6: Contestation Is Now Part of the Landscape

Five “C” codes are especially notable. The HIPAA Security Rule NPRM remains unfinalized even though it is substantively responsive to Concerns 1 and 2, and its May 2026 finalization window has lapsed without action. The Section 1557 § 92.210 affirmative AI duties, which previously supported “A” codes for Concerns 1, 2, and 6, are now under a non-enforcement posture, shifting them to “C.” Separately, portions of OCR's reproductive-health Privacy Rule were vacated in *Texas v. HHS*, and parts of Section 1557's expanded sex-discrimination definition remain subject to nationwide injunction in *Tennessee v. Becerra*. Although these latter developments are unrelated to the AI provisions at § 92.210, they illustrate that the

regulatory architecture is now actively litigated. For longitudinal gap-analysis work of this kind, “addressed” is no longer a fully stable category.

Discussion of Findings

Progress Since 2021

The post-2021 U.S. architecture reflects meaningful progress relative to the regulatory vacuum described by 2021 respondents. Algorithmic transparency for AI embedded in certified health IT is now binding on developers through HTI-1’s source-attribute and intervention risk-management requirements (ONC, 2024). Nondiscrimination obligations apply to patient care decision support tools used by recipients of HHS funding through Section 1557, subject to the enforcement contingencies noted above (HHS OCR, 2024a). The lifecycle of AI-enabled medical devices is now more clearly structured by the FDA’s PCCP framework, although the guidance remains non-binding and operates through underlying device-review authority (FDA, 2024b). Non-HIPAA health-data actors face FTC enforcement under both the Health Breach Notification Rule and Section 5. State consumer-health-data laws extend protection beyond HIPAA’s perimeter. If finalized, the HIPAA Security Rule update would modernize the cybersecurity baseline for the first time in more than a decade and explicitly require AI tools processing electronic PHI to be included in risk analysis. These developments did not exist in 2021.

Persistent Gaps

Despite this progress, three structural gaps remain. First, the covered-entity perimeter is unchanged under the law. AI developers and vendors are subject to HIPAA only if they qualify as business associates under the existing framework, a status that depends on the contractual context of data processing rather than on the nature or scale of the data processed. The 2021 respondents’ concern therefore remains substantially valid; readers who view FTC case-by-case enforcement as a functional substitute may reach a different interpretation, but it does not resolve the structural issue. Second, HIPAA’s de-identification methodology remains unchanged despite substantial literature showing that de-identified datasets may be vulnerable to re-identification, particularly under ML-driven linkage and the five-pathway AI-era re-identification taxonomy detailed in Finding 3 (Rocher et al., 2019; Xia et al., 2021). The FTC’s BetterHelp action introduced a contextual-risk standard for hashed identifiers, but that standard has not been incorporated into HIPAA. Third, the United States still lacks a horizontal AI-in-healthcare statute; the layered architecture leaves significant gaps, particularly for AI used in clinical settings that are not federally funded, AI deployed outside certified health IT, AI that does not meet the device definition, and AI that processes health data outside both the HIPAA framework and the FTC’s enforcement priorities.

The 2021 Baseline in the Generative-AI Era

A reasonable objection is that a 2021 six-concern baseline predates the generative-AI inflection and may now be incomplete. The analysis suggests the six concerns remain valid but are intensified rather than displaced by generative and multimodal systems. Large language models sharpen the vendor-perimeter problem (Concern 3), because general-purpose model providers rarely sit inside HIPAA’s covered-entity structure; they aggravate re-identification (Concern 4) through the generative pathways identified above; and they widen the literacy gap (Concern 5), since clinicians must now appraise probabilistic, non-deterministic outputs. At the same time, generative AI introduces concerns the 2021 frame captures only obliquely—output reliability and hallucination, training-data provenance, and accountability for non-deterministic clinical recommendations—which the corpus addresses only at the margins (NIST AI 600-1; California AB 3030’s generative-AI disclosure requirement). These emergent concerns mark the principal boundary of the present study and the clearest agenda for a replicated 2026 practitioner survey.

International Comparators

Two international instruments illuminate choices the U.S. has and has not made. The European Union AI Act (Regulation (EU) 2024/1689) imposes a binding AI-literacy requirement on providers and deployers under Article 4, in force since February 2, 2025, and attaches obligations to providers and deployers irrespective of healthcare-sector status under Article 6 and Annex I. The AI Act's structural innovation such as sector-agnostic, actor-based obligations resolves within the EU jurisdiction the two concerns the U.S. has handled least well: the covered-entity perimeter problem (Concern 3) and the workforce-literacy gap (Concern 5). The EU approach is not costless; interaction with the Medical Device Regulation and In Vitro Diagnostic Regulation introduces dual-compliance complexity, and full applicability for AI integrated into MDR/IVDR-regulated products does not arrive until August 2, 2027. ISO/IEC 42001:2023, the first certifiable AI Management System standard, supplies a voluntary management-system layer that U.S. healthcare organizations can adopt as anticipatory governance. Neither instrument is part of the U.S. mapping in Table 1, but both demonstrate that horizontal regulatory frameworks for AI are feasible.

Theoretical Reflection and Conceptual Synthesis

Read through adaptive- and anticipatory-governance theory, the central question is not whether new rules exist but whether the architecture can keep pace. The post-2021 landscape has strengthened the formal response to several perceived threats, yet the practical capacity to respond remains less developed. Regulatory architecture and workforce capability must improve in tandem; binding rules without trained users to operationalize them risk producing paper compliance rather than protective behavior. Finding 4, in which every applicable U.S. instrument is coded "P" on workforce literacy, provides the clearest empirical signature of this gap. Institutional layering, moreover, predicts that piecemeal accretion will produce coordination costs and contested boundaries precisely of the kind the matrix reveals.




Figure 1. A Conceptual Synthesis of Healthcare AI Governance

Building on this interpretation, Figure 1 presents a conceptual synthesis of healthcare AI governance organized around four interrelated dimensions: regulatory architecture, HIPAA coverage boundaries, operational guidance, and AI workforce literacy. The four dimensions map onto the layered-architecture finding, with the pacing-and-layering dynamic as the cross-cutting force pressing on the system: strength

in one dimension cannot fully offset persistent weakness in the others. Findings 1, 5, and 6 converge on the structure and instability of the regulatory architecture; Finding 2 centers on HIPAA coverage boundaries; Finding 3 highlights re-identification risk; Finding 4 points to the workforce-literacy gap; and the operational-guidance dimension reflects the continuing practical guidance deficit identified in Concern 6.

Implications for Research and Practice

The findings above describe the regulatory architecture; the recommendations that follow are the authors' interpretations of what those findings imply, not empirical results.

Implications for Research

This study shows the value of using empirically elicited practitioner concerns as a baseline for longitudinal governance analysis: rather than evaluating frameworks in the abstract, researchers can assess whether evolving regulatory responses address the specific concerns raised by practitioners. The structured mapping method used here, including its partial-code taxonomy, provides a replicable way to track governance responsiveness over time. The findings suggest that future research should pay closer attention to how governance gaps persist across HIPAA coverage boundaries, operational guidance, and AI workforce literacy. Finally, the study points to the need for follow-up empirical work, including replication of the 2021 survey in the generative-AI era and comparative analysis across jurisdictions.

Implications for Healthcare Organizations

Healthcare organizations should treat the post-2021 governance architecture as cumulative rather than alternative. Practical priorities include maintaining an AI asset inventory aligned with the anticipated requirements of the HIPAA Security Rule NPRM; ensuring that Business Associate Agreements explicitly address AI-specific risks, including model training, output validation, and re-identification controls; aligning internal AI governance with NIST AI RMF as voluntary scaffolding that can demonstrate due diligence; and preparing for HTI-1 source-attribute requirements, Section 1557 affirmative duties under § 92.210 (while monitoring their contingent enforcement status), and any final HIPAA Security Rule updates. In parallel, organizations should invest in workforce AI literacy programs that cover the capabilities, limitations, and risks of AI tools in clinical settings, to build the coping capacity identified as a critical gap in Finding 4, ideally structured around the NIST AI RMF Govern and Map functions and the voluntary training guidance issued by HHS OCR in its January 2025 Dear Colleague letter. Organizations should also monitor FTC enforcement, as the Health Breach Notification Rule and Section 5 reach activities that HIPAA does not.

Implications for Policymakers

Four policy actions would substantially reduce the gaps identified in this analysis. First, Congress or HHS should extend HIPAA's covered-entity definition, or business-associate scope, to AI developers that process identifiable health data at scale, regardless of contractual relationship to a traditional covered entity. Second, HHS should update HIPAA's de-identification methodology to incorporate contextual re-identification risk, drawing on the standard the FTC applied in BetterHelp; the Expert Determination method already permits contextual analysis, whereas Safe Harbor does not, creating an asymmetry that encourages reliance on the weaker standard. Third, federal coordination across the device layer (FDA), certification layer (ONC), civil-rights layer (HHS OCR Section 1557), security layer (HHS OCR Security Rule), and consumer-protection perimeter (FTC) should be formalized, whether through executive order, interagency memorandum, or congressional action establishing a more horizontal AI-in-healthcare statute. Fourth, HHS should establish baseline AI literacy requirements for healthcare professionals who use AI-enabled tools, for example by conditioning Medicare and Medicaid reimbursement on demonstrated

workforce AI competence or by incorporating AI literacy modules into the ONC Health IT Certification Program. Such a requirement would begin to close the coping-capacity gap identified in Finding 4 and would parallel the binding AI-literacy mandate already in force under Article 4 of the EU AI Act.

Limitations

This analysis is limited to U.S.-applicable instruments and provides a snapshot as of mid-2026. State laws are treated as a single supplementary column, and international comparators are discussed only outside the primary matrix. The “C” codes are especially susceptible to change: the HIPAA Security Rule NPRM may finalize and shift Concerns 1 and 2, and any change in the Section 1557 non-enforcement posture would shift the § 92.210 cells. The exclusion of implementation effects is a deliberate scope boundary, not an oversight: document analysis cannot assess actual implementation, organizational uptake, clinical impact, model fairness in deployment, or organizational implementation maturity. The paper describes the regulatory architecture, not its operational effectiveness, and the designed next step is a validation study pairing a replicated practitioner survey with organizational-readiness measures. The four-point rubric, even with partial-code subtypes, simplifies a complex landscape. Moreover, the matrix treats instruments of differing legal character (voluntary frameworks, non-binding guidance, binding rules, proposed rules, and enforcement actions) in a common format, which necessarily abstracts away differences in legal weight, enforcement strength, and real-world reach; the P-subcodes mitigate but do not eliminate this simplification. Although two-coder reliability reduces subjectivity, it does not eliminate interpretive judgment; alternative codings, particularly on whether case-by-case FTC enforcement should be treated as functionally addressing Concern 3, remain defensible.

Conclusions

This paper used the six concerns identified in a 2021 quantitative study of 281 U.S. healthcare and HIPAA professionals as a fixed reference point for assessing post-2021 U.S. governance responses to AI in healthcare. The analysis shows that none of the six concerns is fully resolved by the binding U.S. instruments reviewed as of mid-2026. Five are partially addressed through layered, jurisdiction-specific responses; re-identification of de-identified data remains materially unresolved; and several responsive instruments are now contingent through pending finalization or non-enforcement. Overall, the U.S. response is better characterized as a fragmented regulatory architecture produced by institutional layering than as a unified governance regime, with the covered-entity perimeter remaining a central unresolved gap.

The implications of these findings converge on a common theme: structural governance gaps require structural responses. Policymakers face the clearest near-term test of whether the covered-entity perimeter, de-identification standards, and cross-agency coordination are addressed through statute or binding rule rather than through continued accretion of voluntary frameworks and case-by-case enforcement. Healthcare organizations and researchers, in turn, cannot wait for those structural changes; anticipatory governance and continued empirical scrutiny of the regulatory mosaic are both warranted in the interim.

Future research should build on this foundation in three directions. First, a replicated practitioner survey, paired with organizational-readiness measures, can assess whether the regulatory changes documented here have shifted healthcare professionals' risk perceptions and their organizations' capacity to adopt AI safely—the validation study called for in the Limitations. Second, a comparative governance analysis across jurisdictions, including the European Union's AI Act implementation and emerging frameworks in other OECD healthcare systems, can reveal whether the structural gaps identified here are unique to the U.S.

architecture or common across regulatory models. Third, as generative AI and multimodal models continue to penetrate clinical workflows, the 2021 concern baseline will need systematic updating; a new survey that captures emergent risks such as hallucination, training-data provenance, and non-deterministic accountability will both test the durability of the original six concerns and extend the evaluative frame. Across these directions, one methodological commitment runs throughout: governance responsiveness cannot be assessed from a single vantage point. The empirical baseline established in 2021 and the structured mapping rubric developed here provide a replicable platform precisely because they can be returned to as instruments are finalized, litigated, and revised.

References

- Alston & Bird. (2025, November 4). HIPAA Security Rule: Still on track for finalization. Health Care/Privacy, Cyber & Data Strategy Advisory.
<https://www.alston.com/en/insights/publications/2025/11/hipaa-security-rule-overhaul>
- American Medical Association. (2026). 2026 physician survey on augmented intelligence. AMA Center for Digital Health and AI.
- Bennett Moses, L. (2011). Agents of change: How the law “copes” with technological change. In G. E. Marchant, B. R. Allenby, & J. R. Herkert (Eds.), *The growing gap between emerging technologies and legal-ethical oversight* (pp.1-37). Springer.
- Black, J. (2001). Decentring regulation: Understanding the role of regulation and self-regulation in a “post-regulatory” world. *Current Legal Problems*, 54(1), 103–146.
- Bowen, G. A. (2009). Document analysis as a qualitative research method. *Qualitative Research Journal*, 9(2), 27–40. <https://doi.org/10.3316/QRJ0902027>
- California Assembly Bill 3030. (2024). Health care services: Artificial intelligence. Cal. Health & Saf. Code § 1339.75.
- European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L series. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- Exponent. (2024). FDA finalizes guidance for AI/ML medical devices.
<https://www.exponent.com/article/fda-finalizes-guidance-aiml-medical-devices>
- Federal Trade Commission. (2023, February 1). FTC enforcement action to bar GoodRx from sharing consumers’ sensitive health info for advertising [Press release].
- Federal Trade Commission. (2024a). Health Breach Notification Rule, final amendments. 89 Fed. Reg. 47,028 (May 30, 2024) (codified at 16 C.F.R. pt. 318).
- Federal Trade Commission. (2024b, April 15). Proposed FTC order will prohibit telehealth firm Cerebral from using or disclosing sensitive data for advertising purposes [Press release].
- Fredrikson, M., Jha, S., & Ristenpart, T. (2015). Model inversion attacks that exploit confidence

- information and basic countermeasures. Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, 1322–1333.
- Greene, A. H. (2020, February 26). More data please! The challenges of applying health information privacy laws to the development of artificial intelligence. Davis Wright Tremaine Privacy & Security Law Blog.
- HHS Office for Civil Rights. (2024a, May 6). Nondiscrimination in health programs and activities. 89 Fed. Reg. 37,522 (codified at 45 C.F.R. pt. 92).
- HHS Office for Civil Rights. (2025a, January 6). HIPAA Security Rule to strengthen the cybersecurity of electronic protected health information [Notice of Proposed Rulemaking]. 90 Fed. Reg. 898.
- HHS Office for Civil Rights. (2025b, January 10). Dear Colleague letter on Section 1557 nondiscrimination protections and patient care decision support tools.
- HHS Office for Civil Rights. (2025c). Section 1557 of the Patient Protection and Affordable Care Act — current status of provisions stayed, enjoined, or subject to non-enforcement. <https://www.hhs.gov/civil-rights/for-individuals/section-1557/index.html>
- HIPAA Journal. (2026, March 20). Final rule implementing HIPAA Security Rule updates edges closer. <https://www.hipaajournal.com/final-rule-implementing-hipaa-security-rule-updates-edges-closer/>
- Hoffman, S. (2019, September 11). Artificial intelligence in medicine raises legal and ethical concerns. The Conversation.
- Humphrey, B. A. (2021). Data privacy vs. innovation: A quantitative analysis of artificial intelligence in healthcare and its impact on HIPAA regarding the privacy and security of protected health information [Doctoral dissertation, Robert Morris University]. ProQuest Dissertations & Theses Global.
- Hussein, R., Zink, A., Ramadan, B., Howard, F. M., Hightower, M., Shah, S., & Beaulieu-Jones, B. K. (2026). Advancing healthcare AI governance through a comprehensive maturity model based on systematic review. npj Digital Medicine, 9(1), 236. <https://doi.org/10.1038/s41746-026-02418-7>
- International Organization for Standardization. (2023). ISO/IEC 42001:2023 — Information technology — Artificial intelligence — Management system.
- Mahoney, J., & Thelen, K. (2010). A theory of gradual institutional change. In J. Mahoney & K. Thelen (Eds.), Explaining institutional change: Ambiguity, agency, and power (pp. 1–37). Cambridge University Press.
- Marchant, G. E., Allenby, B. R., & Herkert, J. R. (Eds.). (2011). The growing gap between emerging technologies and legal-ethical oversight: The pacing problem. Springer.
- Maryland Online Data Privacy Act, Md. Code Ann., Com. Law §§ 14-4601 et seq. (2024).
- Mintz. (2024, April 29). ACA Section 1557 Final Rule: OCR prohibits discrimination related to use of artificial intelligence in health care.

- National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0). NIST AI 100-1. <https://doi.org/10.6028/NIST.AI.100-1>
- National Institute of Standards and Technology. (2024, July 26). Artificial intelligence risk management framework: Generative artificial intelligence profile. NIST AI 600-1. <https://doi.org/10.6028/NIST.AI.600-1>
- Nevada Senate Bill 370, NRS 603A.400 et seq. (2023).
- Office of the National Coordinator for Health Information Technology. (2024, January 9). Health Data, Technology, and Interoperability (HTI-1) Final Rule. 89 Fed. Reg. 1192 (codified in part at 45 C.F.R. § 170.315(b)(11)).
- Rezaeikhonakdar, D. (2023). AI chatbots and challenges of HIPAA compliance for AI developers and vendors. *Journal of Law, Medicine & Ethics*, 51(4), 988–995. <https://doi.org/10.1017/jme.2024.15>
- Rocher, L., Hendrickx, J. M., & de Montjoye, Y.-A. (2019). Estimating the success of re-identifications in incomplete datasets using generative models. *Nature Communications*, 10, Article 3069. <https://doi.org/10.1038/s41467-019-10933-3>
- Sarewitz, D. (2011). Anticipatory governance of emerging technologies. In G. E. Marchant, B. R. Allenby, & J. R. Herkert (Eds.), *The growing gap between emerging technologies and legal-ethical oversight* (pp. 95-160). Springer.
- Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. 2017 IEEE Symposium on Security and Privacy, 3–18.
- Tovino, S. A. (2024). Artificial intelligence and the HIPAA Privacy Rule: A primer. *Houston Journal of Health Law & Policy*, 24, 127.
- U.S. Food and Drug Administration. (2024b, December). Marketing submission recommendations for a predetermined change control plan for artificial intelligence-enabled device software functions [Final guidance].
- U.S. Food and Drug Administration. (2025, January). Artificial intelligence-enabled device software functions: Lifecycle management and marketing submission recommendations [Draft guidance].
- Washington My Health My Data Act, RCW 19.373 (2023).
- Xia, W., Liu, Y., Wan, Z., Vorobeychik, Y., Kantarcioglu, M., Nyemba, S., Clayton, E. W., & Malin, B. A. (2021). Enabling realistic health data re-identification risk assessment through adversarial modeling. *Journal of the American Medical Informatics Association*, 28(4), 744–752. <https://doi.org/10.1093/jamia/ocaa327>