AffixIO Technical Paper · WP-025
June 2026
affix-io.com

AFFIXIO WHITE PAPER · WP-025

HIPAA-Compliant Clinical AI: Protecting Patient Data with Zero-Knowledge Attestation

Healthcare AI needs patient data to work. Here is how to use it without violating HIPAA or exposing medical records.

AffixIO | United Kingdom | affix-io.com | June 2026

ABSTRACT

Clinical AI is moving fast. Diagnostic support tools, treatment recommendation engines, clinical trial matching platforms, and drug discovery systems are becoming genuine parts of the healthcare workflow, not just research curiosities. Every one of those systems touches protected health information (PHI), and that puts them squarely inside one of the most demanding compliance frameworks in the world. HIPAA does not distinguish between AI systems and any other form of healthcare data processing: the obligations are the same, and the penalties for breaches are severe. This paper describes how zero-knowledge attestation provides a structural solution to the PHI problem in clinical AI. The approach allows eligibility determinations to be made, clinician credentials to be verified, audit trails to be

generated, and patient data to cross jurisdictional borders, all without the underlying medical records ever leaving the clinical data environment. AffixIO's attestation and audit API operates as one component in a broader HIPAA-compliant architecture, and this paper is intended to help healthcare technology teams understand where it fits and how to deploy it correctly.

CONTENTS

SECTION 1

Healthcare AI Meets HIPAA

Healthcare has become one of the most active sectors for artificial intelligence deployment. Oncology AI systems can flag potential malignancies in radiology scans before a radiologist has opened the case. Drug discovery platforms compress years of molecular screening into weeks. Clinical trial matching tools promise to connect patients with trials they would never otherwise have found. Treatment recommendation engines are beginning to move from research contexts into live clinical workflows, suggesting

interventions based on patterns drawn from millions of prior patient records. The potential benefits for patient outcomes and for the efficiency of healthcare systems are genuinely significant.

But every one of these applications shares a fundamental characteristic: they depend on patient data. Not aggregated, anonymised, or synthetic data, but real protected health information (PHI), the category of personal data that HIPAA defines and protects more rigorously than almost any other class of information in US law. Diagnostic AI needs imaging studies and clinical notes. Trial matching needs diagnosis codes, medication histories, and lab results. Treatment recommendation needs the full clinical picture. The more sophisticated the AI application, the more granular the patient data it tends to require.

This creates a structural compliance tension that is not going to go away on its own. HIPAA was designed for a world where patient data moved between defined covered entities, disclosed for specific treatment, payment, or healthcare operations purposes, governed by carefully negotiated Business Associate Agreements. It was not designed for a world where machine learning models process millions of patient records to identify patterns, where eligibility determination algorithms run against live EHR data in near real time, or where a single AI system may be accessed by clinicians across dozens of hospitals with different HIPAA compliance postures. Healthcare AI teams are navigating this tension every day, often with imperfect tools.

Zero-knowledge attestation offers a different approach to the problem. Rather than asking how to share PHI compliantly, ZK attestation asks a different question: can we achieve the clinical or compliance objective without sharing PHI at all? In many cases, the answer is yes. A clinical trial does not need to know a patient's diagnosis in order to determine that the patient meets the eligibility criteria; it needs to know the binary outcome of that determination. A hospital credentialling committee does not need a physician's licence number; it needs confirmation that the physician holds a valid licence for the relevant jurisdiction and speciality. An audit process for AI treatment recommendations does not need access to the underlying patient records; it needs a tamper-resistant record of what the AI system

recommended, for a patient who met the relevant eligibility criteria, using a specific model version. ZK proofs make it possible to provide all of these things without the underlying data.

AffixIO provides a ZK proof API that healthcare technology teams can use to build this kind of attestation layer. The API produces cryptographic eligibility proofs, anchors them in Merkle trees, and signs governance records with ML-DSA-65 post-quantum signatures. This paper explains how those components fit into a HIPAA-native clinical AI architecture, covering four concrete use cases, the relevant regulatory framework, practical implementation patterns, and an honest account of the limitations of the approach.

SECTION 2

The PHI Problem in Clinical AI

Understanding the compliance challenge requires a clear picture of what HIPAA actually protects and why existing approaches to PHI management often fall short in AI contexts.

What counts as PHI

HIPAA defines protected health information as individually identifiable health information held or transmitted by a covered entity or its business associates in any form. The regulation specifies 18 categories of identifier that, if present alongside health information, render that information PHI. These include names, dates (other than year) directly related to an individual, geographic data smaller than state level, telephone numbers, email addresses, social security numbers, medical record numbers, health plan beneficiary numbers, account numbers, certificate and licence numbers, vehicle identifiers, device identifiers and serial numbers, URLs and IP addresses, biometric identifiers, full-face photographs, and any other unique identifying number or code. The breadth of this definition is not accidental. It reflects Congress's recognition that health information can be re-identified through combination with other data, and that the protections need to be broad enough to address that risk.

De-identification and its limits in AI contexts

HIPAA provides two methods by which information can be de-identified and thereby removed from HIPAA's scope. The Safe Harbour method requires the removal or suppression of all 18 categories of identifier, along with any other information the covered entity has actual knowledge could be used to identify an individual. The Expert Determination method allows a statistician to certify that the risk of re-identification is "very small," conventionally interpreted as below a 0.04 threshold. Either method, properly applied, produces data that is no longer PHI and can therefore be used, shared, or licensed without HIPAA's restrictions.

In practice, neither method is as straightforward as it sounds when applied to clinical AI. Research has repeatedly demonstrated that de-identified medical datasets can be re-identified with surprisingly high accuracy using external information that was not present in the dataset. A 2019 study in Nature Communications demonstrated re-identification rates above 99% for anonymised patient records using only 15 demographic attributes. Medical imaging data is particularly difficult to de-identify robustly: facial reconstruction from 3D imaging has been demonstrated, and the patterns of incidental findings in imaging studies can serve as biometric identifiers in combination with other data. As AI models become more powerful at extracting patterns from data, the re-identification risk from de-identified training datasets increases rather than decreases.

Beyond the training data question, clinical AI raises a PHI issue that de-identification cannot address: the inference output problem. A clinical AI system trained on de-identified data will, when deployed, receive real patient data as inputs and produce outputs that reflect what it has learned from training. Those outputs, even if they are expressed as recommendations rather than facts, may implicitly reveal PHI. A treatment recommendation of "consider prophylactic cranial irradiation" in response to a patient record, presented in a context where that recommendation is linked to a patient, may effectively disclose a diagnosis of small-cell lung cancer. De-identification of the training data does not prevent the inference outputs from creating PHI exposure at deployment time.

The Business Associate Agreement: necessary but not sufficient

When a healthcare organisation provides patient data to an AI vendor, it is required by HIPAA to have a Business Associate Agreement (BAA) in place with that vendor. A BAA establishes the permitted uses and disclosures of PHI, requires the business associate to implement appropriate safeguards, and sets out the obligations of both parties in the event of a breach. BAAs are a necessary component of HIPAA-compliant data sharing, and no healthcare AI deployment should proceed without one.

What a BAA does not do, however, is prevent breaches. A BAA transfers contractual liability and establishes legal obligations; it does not reduce the technical risk that the data shared under it will be exposed. When a healthcare AI vendor suffers a data breach, the BAA provides a legal framework for the aftermath, but the patient data has still been exposed. The 2024 Change Healthcare breach, which affected an estimated 190 million individuals, occurred despite the existence of BAAs throughout the supply chain. The contractual apparatus of HIPAA compliance does not substitute for a technical architecture that limits what data is shared in the first place.

The Minimum Necessary standard

HIPAA's Minimum Necessary standard, codified at 45 CFR 164.502(b), requires that covered entities and their business associates make reasonable efforts to limit PHI use and disclosure to the minimum necessary to accomplish the intended purpose. In a clinical AI context, this standard is often more honoured in principle than in practice. An AI vendor that requests access to a full patient record in order to determine trial eligibility is not applying the Minimum Necessary standard if the eligibility determination could have been made from a much smaller set of data points. ZK attestation provides a technical implementation of the Minimum Necessary standard: if only binary eligibility signals are extracted and passed to the attestation layer, and those binary signals are processed in memory without being retained, then by construction only the minimum necessary information has been used.

SECTION 3

Use Case 1: Clinical Trial Eligibility Without Diagnosis Disclosure

Clinical trials are one of the highest-value applications of patient data matching in medicine, and one of the most PHI-intensive. A trial's inclusion and exclusion criteria typically specify a primary diagnosis, a disease stage or severity range, an age bracket, an absence of specified contraindications, prior treatment history, and laboratory value thresholds. Matching patients to trials that fit these criteria is a genuinely difficult problem: estimates suggest that 80% of clinical trials fail to meet enrolment targets on time, and that delayed enrolment is one of the leading causes of trial failure and consequent drug development cost inflation. AI trial matching tools that can scan patient records and identify eligible candidates represent a significant potential improvement over manual screening, but they require access to exactly the kind of PHI that is most sensitive: diagnosis, treatment history, and lab values.

The conventional approach to trial eligibility screening involves a trial coordinator, a research nurse, or an AI system with access to medical records reviewing patient files and assessing them against the protocol's inclusion and exclusion criteria. At each hospital site where screening occurs, this creates PHI exposure: the patient's diagnosis, treatment history, and relevant clinical parameters are shared with or accessed by the trial coordination infrastructure. For a multi-site trial operating across dozens of hospitals, the total PHI exposure surface is large. The patient may not have consented to their records being reviewed by multiple external parties as part of a trial screening process they are ultimately not enrolled in. The situation is particularly sensitive in oncology AI and other contexts where the underlying diagnosis carries significant personal implications.

The ZK approach

A ZK-based clinical trial eligibility approach restructures where the eligibility determination happens. Rather than sharing patient records with the trial's screening infrastructure, the EHR system at the patient's own institution performs the eligibility assessment internally and generates a binary signal for each inclusion and exclusion criterion. A patient with the required diagnosis

generates a signal of `has_diagnosis=1`. A patient in the required age range generates `in_age_range=1`. A patient without the specified contraindications generates `no_contraindication=1`. These binary witnesses are the inputs to a ZK eligibility circuit. The circuit processes the witnesses and produces a proof that all inclusion criteria are met and no exclusion criteria are triggered.

The trial site receives the proof and an eligibility outcome. It does not receive the patient's diagnosis, age, contraindication details, or any other underlying clinical data. The proof is cryptographically binding: a trial site that receives a valid eligibility proof can be confident that the patient's EHR system, attesting to the protocol's criteria, determined that the patient meets those criteria. The proof cannot be forged, and it cannot be reverse-engineered to reveal the underlying clinical data.

> **Circuit implementation omitted from public documentation.** The specific ZK circuit used to combine eligibility witnesses and produce the eligibility proof is not described here. Healthcare technology teams integrating AffixIO's API interact with the proof output and the verification interface; the circuit internals are not required for integration.

The patient benefits in a way that the conventional approach cannot provide: they can seek trial participation across multiple sites and multiple trials without disclosing their condition to multiple screening coordinators and research databases. A patient with a sensitive diagnosis, such as an early-stage cancer, a psychiatric condition, or an HIV-related illness, can test their eligibility for available trials without that disclosure being recorded in trial screening logs outside their primary care institution. For conditions where disclosure itself carries stigma or social consequences, this is not a minor technical detail. It is a meaningful improvement in patient dignity and data control.

> **Key distinction:** The ZK eligibility approach does not change who performs the eligibility assessment. The patient's own EHR system, governed by the institution's existing HIPAA controls, does that work. What changes is what leaves that environment: a proof and an outcome, rather than the patient record.

SECTION 4

Use Case 2: Clinician Credential Verification Across Jurisdictions

Credentialling is a necessary but often cumbersome part of clinical operations. Before a physician can see patients at a hospital, prescribe controlled substances, or join a telehealth platform, their credentials must be verified: medical degree, training completion, board certification, state medical licence, DEA registration for prescribing, and absence of sanctions or disciplinary actions. In the United States, medical licensing is state-based, which means a physician licensed in California is not automatically authorised to practice in New York, and a physician providing telehealth consultations to patients in multiple states may need to maintain licences in each of those states or rely on the Interstate Medical Licensure Compact where it applies.

The conventional credentialling process involves collecting copies of credential documents, submitting verification requests to state licensing boards, the National Practitioner Data Bank (NPDB), the Federation of State Medical Boards (FSMB), and speciality boards, and maintaining a database of verified credential records. This process is time-consuming, taking weeks or months at hospitals that have not modernised their credentialling workflows, and it creates a database of sensitive credential details that is attractive to fraudsters. Synthetic identity attacks and credential fraud in healthcare credentialling have increased as electronic health records have made the value of compromised clinical credentials more apparent.

ZK credential proofs

A ZK-based credentialling approach works from a similar principle to the clinical trial eligibility case. State licensing boards, speciality boards, and the DEA hold the authoritative source data for clinical credentials. Rather than sharing that data with each credentialling request, those authorities can issue binary eligibility signals: `licensed_in_state=1`, `speciality_certified=1`, `no_active_sanctions=1`, `dea_registration_current=1`. A ZK circuit combines these signals into a cross-jurisdictional credential proof. A hospital or

telehealth platform verifying the proof receives confirmation that all required credentials are valid; it does not receive the licence numbers, certificate identifiers, or registration details themselves.

This approach matters particularly for telehealth, where the expansion of cross-state practice has outpaced the credentialling infrastructure. A telehealth platform operating in 30 states currently needs to verify and maintain credentials for physicians across 30 licensing jurisdictions, with different renewal cycles and different data formats for each. A ZK credential proof issued by each relevant state board and presented by the physician to the platform each time credentials are checked reduces this to a verification step rather than a data collection step. The platform confirms validity without accumulating a database of raw credential data.

The same approach extends naturally to cross-border clinical portability. The global clinical staffing and medical commerce market is estimated at $16 billion and involves significant movement of clinicians across national boundaries. A physician trained and licensed in Germany practising in the UK, or a nurse certified in the Philippines working in Australia, faces a credentialling process that involves sharing credential documents with multiple foreign institutions and verification agencies. ZK credential proofs issued by the authoritative body in the country of training or primary licence could provide a portable, verifiable, privacy-respecting credential that travels with the clinician rather than being repeatedly disclosed to each new institution. The EUDI Wallet infrastructure in the EU provides a natural transport layer for this kind of credential in European jurisdictions.

> **Note on implementation:** Achieving this vision requires licensing boards and speciality bodies to participate in a proof-issuing infrastructure. AffixIO provides the cryptographic attestation layer; engagement with licensing authorities on issuing binary signals is the responsibility of the healthcare organisations and platforms implementing the solution.

SECTION 5

Use Case 3: Auditing AI Treatment Recommendations Without PHI Access

Clinical AI systems that make or support treatment recommendations create a new kind of regulatory and governance challenge. They must be accurate, but they must also be auditable: when a recommendation is questioned, whether by a clinician, a patient, a hospital risk committee, or a regulator, there must be a record that allows the recommendation to be traced back to its inputs, its model version, and the logic applied. Without such a record, neither the hospital nor the patient can understand what the AI actually did, and the accountability that regulators and professional standards bodies expect cannot be discharged.

The conventional approach to clinical AI auditability is to log the patient record that was provided to the AI system alongside the recommendation it produced. This approach is thorough in one sense: it preserves a complete record of the AI's inputs and outputs. But it creates a substantial PHI problem. The audit log is now a parallel store of patient health information, accessible to auditors, compliance officers, legal teams, and regulators, all of whom now need to be managed as PHI-accessing parties under HIPAA. A PHI-bearing audit log expands the number of people and systems with legitimate access to patient data for purposes that are not directly related to the patient's care. This is, at minimum, in tension with the Minimum Necessary standard. In practice, healthcare audit logs have been the source of some of the most damaging healthcare data breaches, precisely because they aggregate PHI from many patients in a single accessible location.

ZK governance records for AI recommendations

A ZK audit approach generates a governance record for each AI recommendation that contains all of the elements an auditor needs, without containing any patient data. For each recommendation, the governance record includes: a hash of the recommendation itself, a reference to the patient's ZK eligibility proof (confirming that the patient met the relevant eligibility criteria, without identifying who the patient is or what those criteria were in detail), a hash of the model version that generated the

recommendation, a timestamp, and a ZK proof of computation confirming that the recommendation was generated by that model version for a patient holding that eligibility proof. The governance record is anchored in a SHA-256 Merkle tree and signed with an ML-DSA-65 post-quantum signature.

An auditor presented with this governance record can verify: that the recommendation was generated by a specific, identified model version; that the patient for whom it was generated held a valid eligibility proof at the time; that the recommendation has not been altered since it was generated; and that the governance record itself has not been tampered with. What the auditor cannot determine from the governance record alone is the patient's identity, their diagnosis, or the specific clinical parameters that informed the recommendation. PHI does not enter the audit infrastructure at all.

This structure is directly relevant to FDA requirements. The FDA's AI/ML-Based Software as a Medical Device (SaMD) Action Plan, and the more detailed guidance documents that have followed it through 2024 and 2025, require clinical AI systems to maintain audit trails that support post-market surveillance and real-world performance monitoring. The intent is clear: regulators need to be able to examine AI system behaviour over time, identify performance drift, and trace adverse events back to specific model versions and inputs. ZK governance records satisfy this intent without creating a PHI store in the audit infrastructure. For radiology AI compliance, oncology AI, and other clinical AI applications where SaMD classification is likely, this distinction is significant.

> **On FDA SaMD classification:** Whether a specific clinical AI system constitutes a medical device under FDA jurisdiction depends on its intended use and the claims made about it. Healthcare organisations deploying clinical AI should obtain formal regulatory classification advice from a qualified regulatory affairs specialist. This paper discusses the audit trail dimension of SaMD compliance, not the classification question.

SECTION 6

Use Case 4: Cross-Border Clinical Portability

Patients move. They travel internationally, emigrate, retire abroad, and seek specialist treatment in countries with expertise in their particular condition. When they do, they face a practical problem that the digital health industry has been grappling with for decades: their health records are held by institutions in one jurisdiction, and the healthcare providers they encounter in another jurisdiction have no straightforward way to access or verify their clinical history. The standard response to this problem, asking patients to carry paper copies of their records or to request that records be faxed or emailed between institutions, is both unreliable and creates significant PHI exposure at every transfer point.

FHIR (Fast Healthcare Interoperability Resources) provides an internationally adopted standard for the structure and exchange of health data, and it has become the foundation of most modern health data portability initiatives. But FHIR addresses the format and transport of health records; it does not address the privacy question of what should be shared when a patient presents to a foreign healthcare provider. Sharing a patient's complete FHIR record with a walk-in clinic in a foreign country because the patient has a headache is not a proportionate response. Sharing nothing at all because the records are HIPAA-protected and the foreign clinic has no BAA in place is clinically unsafe.

ZK eligibility proofs as portable clinical credentials

A ZK-based approach to clinical portability treats health eligibility status as something a patient can carry and present selectively, rather than something that must be transferred wholesale from one institution to another. A patient living with a managed chronic condition could carry a set of ZK eligibility proofs: `condition_managed=1`, `current_prescription_active=1`, `vaccinated_per_schedule=1`. These proofs are issued by the patient's primary care institution, signed with ML-DSA-65, and anchored in a Merkle tree. The patient presents them to a healthcare provider in any jurisdiction. The provider verifies the proof against the issuing institution's public key and receives the eligibility outcome: yes, this patient is currently managing a

chronic condition and has an active prescription. The provider does not receive the patient's medical record, the name of the condition, the specific prescription, or the institution that issued the proof.

In the European Union, the eIDAS 2.0 EUDI Wallet regulation provides an infrastructure framework for exactly this kind of credential portability. The EUDI Wallet is designed to carry verifiable credentials issued by authoritative bodies, presentable selectively across EU member states. ZK health eligibility credentials fit naturally into this framework: a patient with a EUDI Wallet could carry their ZK clinical proofs alongside their identity documents and present them to healthcare providers anywhere in the EU. In the UK, the NHS App is developing credential functionality that could serve a similar purpose. For cross-border deployments involving the US and the EU, the FHIR standard provides a common data model that can underpin the binary witness extraction process on both sides of the Atlantic.

For the estimated $16 billion global clinical staffing and medical commerce market, similar logic applies to professionals rather than patients. A clinical researcher participating in a multi-country trial, or a healthcare worker providing services across national borders, can carry ZK credential proofs that are valid across jurisdictions without repeatedly disclosing the underlying credential documents to each institution they engage with. The receiving institution verifies the proof; it receives confirmation of validity rather than a copy of the credential.

SECTION 7

The ZK Architecture for Healthcare

The four use cases described above all use the same underlying architectural pattern, which is worth making explicit. Understanding how the layers fit together helps healthcare technology teams assess where their existing infrastructure interfaces with an attestation approach and what changes are needed to implement it.

Three-layer architecture

The architecture separates into three distinct layers. The first is the **clinical data layer**: the EHR systems, PACS (picture archiving and communication systems), laboratory information systems, and other repositories where PHI actually lives. This layer is governed by the healthcare organisation's existing HIPAA controls. BAAs cover the vendors and service providers who have access to it. Access controls restrict who can query it. The clinical data layer is not changed by adopting a ZK attestation approach; it remains exactly as it is, governed by the existing compliance posture.

The second is the **attestation layer**: the component that extracts binary eligibility signals from the clinical data layer, constructs ZK proofs from those signals, and passes the proofs to whoever needs to verify them. This is where AffixIO's API operates. The attestation layer receives binary witnesses, not patient records. The witnesses are processed in memory to generate the proof and are not retained. PHI does not enter the attestation layer; only the binary outputs of eligibility assessments do. Because those binary outputs may be derived from PHI, AffixIO operates as a Business Associate under HIPAA for healthcare deployments, and a BAA is required.

The third is the **governance layer**: the infrastructure where proofs are anchored, signed, and made available for verification. Each proof is anchored in a SHA-256 Merkle tree. The Merkle root is signed with an ML-DSA-65 signature, a post-quantum lattice-based signature algorithm selected for its resistance to quantum computing attacks. The governance layer provides the tamper-resistance and long-term verifiability that audit and compliance processes require. Governance records in this layer do not contain PHI; they contain proof references, model version hashes, recommendation hashes, and timestamps.

Where AffixIO fits

AffixIO operates in the attestation layer and the governance layer. It does not access the clinical data layer. The EHR adapter, which is the component that queries the EHR system, extracts binary witnesses, and passes them to AffixIO's API, sits at the boundary between the clinical data layer and the attestation layer. That adapter is typically implemented by the healthcare organisation or by the EHR vendor; it is the component that requires the most

careful HIPAA assessment, because it is the component that touches PHI directly. AffixIO's role begins after the binary witnesses have been extracted and the PHI has been left behind.

This positioning matters for the compliance analysis. AffixIO does not need access to diagnosis codes, patient names, medication lists, or any other PHI in order to generate a valid eligibility proof. What it receives is a set of binary flags and a commitment to the criteria being assessed. What it produces is a cryptographic proof that the flags satisfy the criteria. The proof cannot be reversed to reveal the flags, and the flags cannot be reversed to reveal the patient data they were derived from. The chain of inference from the proof back to the patient record is broken at the attestation layer boundary.

SECTION 8

HIPAA Security Rule and the Proof-Based Compliance Record

The HIPAA Security Rule (45 CFR Part 164, Subparts A and C) establishes the framework of safeguards that covered entities and business associates must implement to protect electronic PHI (ePHI). The Security Rule organises its requirements into three categories: administrative safeguards, physical safeguards, and technical safeguards. While a full Security Rule compliance programme is beyond the scope of this paper, the technical safeguards are directly relevant to understanding how ZK governance records fit into a HIPAA-compliant architecture.

Technical safeguards and audit controls

The Security Rule's technical safeguard requirements include access controls (45 CFR 164.312(a)(1)), which require mechanisms to allow only authorised persons to access ePHI; audit controls (45 CFR 164.312(b)), which require hardware, software, and procedural mechanisms to record and examine activity in information systems that contain or use ePHI; integrity controls (45 CFR 164.312(c)(1)), which require mechanisms to authenticate that ePHI has

not been improperly altered or destroyed; and transmission security (45 CFR 164.312(e)(1)), which requires technical security measures to guard against unauthorised access during transmission.

ZK governance records address the audit control requirement in a distinctive way. A conventional audit log satisfies this requirement by recording who accessed what ePHI, when, and for what purpose. That log itself contains ePHI references, which means the audit infrastructure becomes part of the ePHI compliance perimeter. A ZK governance record, by contrast, records that a clinical AI interaction occurred, that it involved a patient holding a specified eligibility proof, that the AI model used was a specific verified version, and that the recommendation produced has not been altered since it was generated, all without the governance record itself containing ePHI. The integrity and verifiability that the audit control requirement seeks are provided, while the PHI exposure that conventional audit logs create is avoided.

The Minimum Necessary standard as a technical property

The Minimum Necessary standard at 45 CFR 164.502(b) is typically treated as a policy requirement: covered entities must establish policies and procedures to identify the persons or classes of persons who need access to PHI, and to limit requests for PHI to what is reasonably necessary. In practice, implementation often relies on role-based access controls and training, which are policy instruments rather than technical ones. A determined or negligent employee with appropriately broad access can still access more PHI than is necessary for their role.

A ZK attestation architecture implements the Minimum Necessary standard at the technical level rather than the policy level. If the attestation layer is designed so that it receives only binary witnesses and not underlying patient data, it is technically impossible for the attestation layer to access more PHI than the minimum necessary for the eligibility determination. The constraint is structural, not behavioural. This is a meaningful improvement over a policy-only implementation, where the guarantee of minimum necessary access depends on human compliance rather than technical architecture. It does not replace the policy requirement, which still applies to all parties in the system, but it adds a technical enforcement layer.

> **45 CFR 164.502(b) in practice:** Healthcare organisations adopting ZK attestation should document in their HIPAA policies how the attestation layer implements the Minimum Necessary standard technically, and confirm that the EHR adapter component follows the same principle. The policy documentation reinforces the technical control and supports the organisation's risk assessment.

SECTION 9

HITECH, the 21st Century Cures Act, and FDA AI Guidance

HIPAA does not operate in isolation. Healthcare technology teams navigating clinical AI compliance need to understand how three additional regulatory frameworks interact with HIPAA and with a ZK attestation approach: the Health Information Technology for Economic and Clinical Health (HITECH) Act, the 21st Century Cures Act, and the FDA's evolving guidance on AI and machine learning in medical devices.

HITECH and the expansion of HIPAA obligations

The HITECH Act, enacted as part of the American Recovery and Reinvestment Act of 2009, significantly strengthened HIPAA's enforcement regime. HITECH increased civil monetary penalties for HIPAA violations to a maximum of $1.9 million per violation category per year (figures updated for 2023 inflation adjustments). More significantly for clinical AI deployments, HITECH extended HIPAA's direct obligations to business associates. Before HITECH, business associates were contractually obligated to comply with HIPAA through the BAA, but enforcement action could only be taken against covered entities. After HITECH, the OCR can pursue enforcement actions directly against business associates. Any AI vendor or attestation layer provider that handles PHI, or signals derived from PHI, in connection with a healthcare organisation's operations is in scope for direct HITECH enforcement.

HITECH also introduced the Breach Notification Rule (45 CFR Part 164, Subpart D), which requires covered entities and business associates to notify affected individuals, the Secretary of HHS, and in some cases the media, when a breach of unsecured PHI occurs. The healthcare data breach notification regime has become one of the most significant practical drivers of PHI minimisation strategies, because the notification and remediation costs of a large breach are substantial, reputationally damaging, and in many cases avoidable through better data architecture. A clinical AI architecture that limits PHI to the clinical data layer and uses ZK proofs for everything downstream reduces the number of systems that would trigger breach notification requirements if compromised.

The 21st Century Cures Act and information blocking

The 21st Century Cures Act, enacted in 2016 and with its information blocking provisions entering into force through 2020 and 2022, takes a different angle on health data: it is primarily concerned with ensuring that patients can access and use their own health data, and that providers and health IT vendors do not engage in practices that unreasonably restrict such access. The Act's information blocking prohibition, enforced by the Office of the National Coordinator for Health Information Technology (ONC), prohibits covered actors from interfering with the access, exchange, or use of electronic health information except under defined exceptions.

Clinical AI teams need to be aware that a ZK attestation approach that gives patients portable ZK eligibility proofs could, if implemented thoughtfully, be a pro-patient extension of the 21st Century Cures Act's intent: it gives patients a verifiable credential they control, derived from their own health data, that they can present to healthcare providers without requiring full record transfer. Conversely, an attestation layer that is designed in a way that limits patient access to their own eligibility proofs could create information blocking concerns. Healthcare organisations should ensure that their implementation of ZK attestation includes a patient access path for the proofs generated about them.

FDA AI/ML guidance for Software as a Medical Device

The FDA's regulatory framework for AI/ML-Based Software as a Medical Device (SaMD) has developed substantially through the early 2020s. The 2021 AI/ML-Based SaMD Action Plan established the FDA's intent to develop a risk-based regulatory framework for AI systems that learn and adapt over time. Subsequent guidance documents, including the 2024 marketing submission guidance for predetermined change control plans, have moved those intentions into detailed requirements. The core obligations for SaMD using AI/ML include: a predetermined change control plan that defines what kinds of model updates are permitted without additional FDA review; real-world performance monitoring to track whether the AI system's performance in deployment matches its performance in validation; and audit trails that allow regulators to reconstruct the system's behaviour at any point in time.

ZK governance records are well-suited to supporting the audit trail requirement for AI treatment recommendation audit and similar SaMD applications. Each governance record anchors a recommendation to a specific model version hash, confirming which model generated which recommendation, in a tamper-resistant form that regulators can verify without accessing patient records. When a model is updated as part of a predetermined change control plan, the model version hash in the governance layer changes, creating a clear record of when the change occurred and which recommendations were generated by which model version. Healthcare organisations implementing clinical AI under FDA SaMD oversight should discuss with their regulatory affairs team how ZK governance records integrate with their overall SaMD audit trail strategy, rather than treating governance records as a standalone compliance solution.

SECTION 10

Implementation Patterns

The architecture described in earlier sections can be implemented in several different configurations depending on the clinical AI deployment context. The three patterns below represent the most common scenarios encountered in healthcare technology programmes.

Pattern 1: EHR adapter model

In this pattern, a thin adapter layer sits between the EHR system and AffixIO's API. The adapter is the component that has direct access to the EHR and that contains the logic for extracting binary eligibility witnesses from FHIR resources. When an eligibility check is required, the adapter queries the relevant FHIR resources (Condition, MedicationRequest, Observation, and so on), evaluates each criterion against the data returned, generates a binary witness for each criterion, passes the witnesses to AffixIO's API, and discards the source data. The adapter never retains PHI beyond the transaction that requires it. The adapter vendor, if different from the EHR vendor, requires its own BAA with the healthcare organisation.

The EHR adapter model is the most flexible pattern and is appropriate for custom clinical AI integrations where the healthcare organisation controls the eligibility logic. It requires careful design of the adapter's data handling: the adapter code should be reviewed to confirm that source FHIR data is not logged, cached, or transmitted beyond the adapter process. A security review of the adapter should be part of the BAA negotiation and HIPAA risk assessment process.

Pattern 2: Clinical trial platform integration

In this pattern, a trial management system calls AffixIO's eligibility proof endpoint at the screening stage of the trial. The trial management system knows the protocol's inclusion and exclusion criteria; the EHR system knows whether the patient meets them. The adapter extracts the binary witnesses from the EHR, passes them to AffixIO's API, and the trial management system receives a proof reference that is stored in the trial record in place of patient data. When the trial record is audited, the auditor sees that patient X (identified by an internal trial identifier, not by name or medical record number) had a valid eligibility proof at the time of screening. The patient's diagnosis and clinical parameters are not recorded in the trial management system.

This pattern is relevant to any clinical trial operating under a BAA-based data governance arrangement with multiple hospital sites. It reduces the PHI footprint of the trial management system substantially, which simplifies the

HIPAA compliance posture of the trial sponsor and reduces the breach notification surface if any component of the trial management infrastructure is compromised.

Pattern 3: Telehealth credential chain

In this pattern, a credentialling service calls AffixIO at provider onboarding. The credentialling service receives binary credential validity signals from licensing boards, the NPDB, and speciality bodies, and passes them to AffixIO's API to generate a credential proof. The proof is attached to each patient encounter in the telehealth platform as evidence that the treating clinician held valid credentials at the time of the encounter. The telehealth platform does not store the clinician's licence numbers, DEA registration number, or board certificate identifiers; it stores the proof reference.

This pattern simplifies the telehealth platform's obligations around credential data management. A telehealth platform that does not hold raw credential data cannot expose it in a breach. The platform's credential audit trail, which is a regulatory requirement for most telehealth operators, is satisfied by the proof references rather than by a database of credential documents.

BAA coverage for all patterns

All three patterns require a BAA between the healthcare organisation and AffixIO, covering AffixIO's role as a Business Associate in the attestation layer. The BAA should specify the permitted uses and disclosures of the binary witness data and the proof outputs, the security requirements that AffixIO applies to the attestation layer, and the obligations of both parties in the event of a breach. Healthcare organisations should also confirm that their EHR vendor and any adapter vendor are covered by appropriate BAAs, and that their cloud infrastructure provider has a HIPAA-eligible services agreement in place for any infrastructure used in the attestation layer.

SECTION 11

Known Limitations

This paper would not be complete without an honest account of what ZK attestation for clinical AI does not solve. The approach is genuinely useful, but it is not a universal answer to healthcare data privacy, and overstating what it provides would be a disservice to healthcare technology teams making implementation decisions.

The garbage-in problem

A ZK eligibility proof proves that a computation was performed correctly. It does not validate the underlying data that the computation operated on. If a patient's diagnosis code is recorded incorrectly in the source EHR, perhaps because of a data entry error, a miscoding, or a pending update that has not yet been confirmed, the binary witness extracted from that code will be incorrect. The ZK proof will then attest, with full cryptographic validity, to an eligibility determination that is factually wrong. The proof is correct about the computation; the computation is incorrect about the patient's clinical status.

This limitation means that ZK attestation cannot replace the quality controls that healthcare organisations should be applying to their source EHR data. Diagnosis code accuracy, medication list currency, and lab value validity all need to be maintained through existing clinical and administrative processes. ZK attestation works downstream of those processes; it does not improve them. For clinical trial eligibility in particular, where an incorrect inclusion determination could expose a patient to an unsuitable protocol, this limitation has direct patient safety implications that the trial sponsor and IRB need to consider.

Complex eligibility criteria

Clinical eligibility criteria are often more nuanced than binary flags. An exclusion criterion of "not currently taking medications that interact with the study drug" involves a dynamic list of contraindicated medications that may be updated as new interaction data emerges. An inclusion criterion of "adequate renal function" typically involves a calculated value (eGFR) derived from multiple lab results, with a threshold that may depend on other patient

parameters. Mapping these complex criteria to binary witnesses requires sophisticated adapter logic and careful translation of the protocol's language into computable rules. That translation process is itself a source of potential error and needs to be validated by clinical experts before deployment.

The EHR adapter is therefore not a simple technical component: it encodes clinical judgment about how to translate protocol criteria into computable queries. Healthcare organisations implementing this approach should ensure that adapter logic is reviewed and approved by clinical informatics specialists and, where the application is a regulated clinical trial, by the trial sponsor's medical team. AffixIO's role in this process is to provide the cryptographic attestation layer; the clinical validity of the criteria mapping is the responsibility of the healthcare organisation and the trial sponsor.

Regulatory acceptance is not yet established

ZK proofs are cryptographically well-founded, and the approach described in this paper is technically sound. However, regulatory acceptance of ZK-based clinical records as a primary compliance record is not yet established in most healthcare jurisdictions. The FDA has not issued guidance specifically addressing ZK governance records for SaMD audit trails. NHS England and the EMA have not published positions on ZK-based clinical portability credentials. Healthcare organisations adopting this approach should engage with their relevant regulatory body, whether FDA, NHS England, EMA, or another authority, before treating ZK governance records as their sole compliance record for any regulated activity. The most prudent approach during this period is to use ZK governance records as one component of a broader compliance record that also includes conventional documentation, rather than as a replacement for it.

BAA coverage of the full supply chain

The BAA requirement extends to the full supply chain of parties that handle PHI or PHI-derived data in connection with a clinical AI deployment. The EHR adapter vendor, the cloud infrastructure provider running the adapter, AffixIO as the attestation layer provider, and any subprocessors that AffixIO uses in delivering its service all need to be covered by appropriate BAAs. Healthcare organisations should map the full data flow from the EHR system through to

the governance layer and confirm BAA coverage at each step before going live with a clinical AI deployment. Gaps in BAA coverage are one of the most common findings in HIPAA audits of healthcare technology deployments.

SECTION 12

Conclusion

Healthcare AI is not slowing down. The clinical value of AI-assisted diagnostics, trial matching, treatment support, and drug discovery is too significant for the healthcare industry to set aside because of compliance complexity. But the compliance complexity is real, and it is not going away either. HIPAA, HITECH, the 21st Century Cures Act, and FDA SaMD guidance together create a demanding regulatory environment for any AI system that touches patient data, and the penalties for getting it wrong, financial, reputational, and in terms of patient harm, are serious.

ZK patient attestation offers a structural way to address the most fundamental part of that complexity: the need to share PHI in order to make clinical determinations. By separating the eligibility determination, which requires PHI, from the proof of that determination, which does not, ZK attestation allows clinical AI workflows to operate at arm's length from the patient records they depend on. Trial eligibility can be confirmed without disclosing diagnoses to trial sites. Clinician credentials can be verified without accumulating credential databases. AI treatment recommendations can be audited without creating PHI-bearing audit logs. Patient health eligibility status can cross jurisdictional borders without triggering the full apparatus of international health data transfer.

The approach is not a replacement for HIPAA compliance. Covered entities and business associates still need BAAs, workforce training, policies and procedures, risk assessments, and all of the other elements that a mature HIPAA compliance programme requires. What ZK attestation provides is a technical architecture that implements the Minimum Necessary standard at a structural level, reduces the PHI footprint of clinical AI deployments, and generates tamper-resistant governance records without adding to the stock of PHI in the system. It is a HIPAA-native architecture rather than a

workaround, and it grows more valuable as clinical AI deployments become more complex and the regulatory scrutiny of healthcare data handling intensifies.

AffixIO's attestation and audit layer is one component in a larger ecosystem that includes EHR systems, clinical AI models, trial management platforms, telehealth infrastructure, and the regulatory frameworks that govern all of them. Healthcare technology teams using AffixIO's API take responsibility for the EHR adapter that feeds it and the clinical AI system that consumes its proofs. AffixIO provides the cryptographic attestation and governance layer with the robustness and auditability that healthcare applications require. The rest of the architecture, including the clinical judgments, the regulatory conversations, and the BAA structure, is the shared responsibility of the healthcare organisation and its technology partners.

Related white papers

- WP-006: PII-Free KYC
- WP-009: Privacy-Preserving Age Verification
- WP-008: ZK GDPR Article 25
- WP-002: Post-Quantum Attestation

Frequently asked questions

Does ZK patient attestation replace HIPAA compliance?

No. HIPAA obligations apply to covered entities and business associates regardless of technical architecture. ZK attestation implements the Minimum Necessary standard technically and reduces PHI exposure, but covered entities still need BAAs, workforce training, policies and procedures, and risk assessments. The attestation layer is one component of a compliant architecture, not a shortcut around the compliance programme.

What is a Business Associate Agreement and does AffixIO require one?

A BAA is a contract required by HIPAA between a covered entity and any vendor that handles PHI on its behalf. Because AffixIO processes binary eligibility signals derived from PHI in the attestation layer, a BAA is required for

any healthcare deployment. AffixIO's healthcare BAA covers the permitted uses and disclosures of the binary witness data and proof outputs, the security measures applied in the attestation layer, and breach notification obligations.

Can ZK proofs satisfy FDA audit trail requirements for AI/ML SaMD?

ZK governance records can form part of a compliant audit trail for AI/ML Software as a Medical Device. Each recommendation generates a tamper-resistant record linked to a model version hash and anchored in a Merkle tree signed with ML-DSA-65. Healthcare organisations should confirm with their regulatory affairs team that the specific record format satisfies their SaMD submission requirements and engage with the FDA on how ZK governance records fit within their predetermined change control plan before relying on them as the primary audit record.

WP-006 · WP-009 · WP-002

WP-025 · June 2026

▶ About
▶ Solutions
▶ Legal
▶ Trust & Security
Contact
truth layer | yes | no | proof