Verifiable Governance in Prior Authorization

A Structural Analysis of the Evidentiary Gap in AI-Assisted Utilization Management

A Policy and Infrastructure White Paper

April 2026

Prepared for health plans, providers, state regulators, and federal agencies engaged in the 2026 prior authorization transition

Executive Summary

Prior authorization in American healthcare is undergoing a regulatory and technological transition without precedent in the history of utilization management. The CMS Interoperability and Prior Authorization Final Rule (CMS-0057-F) begins taking effect in 2026. AHIP's June 2025 commitments set industry-wide targets for real-time decisioning and transparency. Multiple states have enacted or introduced AI-in-utilization-management legislation. Plaintiff litigation against AI-assisted denials has accelerated. The Senate Permanent Subcommittee on Investigations published findings on Medicare Advantage denial practices that will shape regulatory attention for years.

All of this unfolds against the reality that roughly 130 million prior authorization determinations are produced annually across the U.S. healthcare system, an increasing share of them shaped by algorithmic tools whose behavior is difficult for any external party to verify. The existing evidentiary infrastructure — policy documents, reviewer attestations, SOC 2 reports, post-hoc audit responses — was designed for an era when decisions were human-made and reviewable case by case. It does not scale to algorithmic decisioning at population volume, and it does not produce the form of evidence that regulators, auditors, courts, and members are increasingly demanding.

This paper examines the structural gap between what current prior authorization governance can prove and what the 2026 regulatory environment requires. It argues that the gap cannot be closed by additional policy documents, better attestations, or more comprehensive audits — not because any of these are wrong, but because they all rely on the same underlying evidentiary mechanism: trust in the assertions of the party being evaluated. When the party is an insurer, a delegated UM vendor, or an AI-automation platform, and the evaluator is a regulator, plaintiff, or member, the evidentiary mechanism is exactly what is in question.

What the transition requires is a form of evidence that does not rely on trust in the party producing it. Such evidence exists — it has existed in cryptographic research literature for more than a decade and is now practical at the latency and volume prior authorization demands. This paper describes what verifiable-governance infrastructure is, what it can and cannot prove, how it interacts with existing governance mechanisms, and what the implications are for health plans, providers, regulators, and the vendor ecosystem that supports all three.

The paper is deliberately neutral on vendor selection and commercial strategy. Its purpose is to make the structural analysis accessible to the decision-makers — in health plans, provider organizations, regulatory agencies, and the research community — who will shape how the transition unfolds over the next three years. The conclusion the paper argues for is not that verifiable governance should be adopted as a particular product; it is that the transition is underway whether any single actor plans for it or not, and

that the organizations and agencies that understand the structural dynamics early will shape the terms on which the transition occurs.

1. The Evidentiary Gap in Prior Authorization

1.1 What current governance can actually prove

Consider a specific prior authorization determination — any one will do. A member's physician submits a request for a specialty therapy. The request enters the payer's utilization management pipeline. It is routed through triage logic. Its clinical data is extracted into structured features. An algorithmic tool produces a recommendation. A human clinician reviews the recommendation. A determination is issued.

Now suppose a regulator, a plaintiff, or the member asks a specific question about this determination: was the algorithmic recommendation influenced by cost features that should not have been considered? Was the model used for this decision the model the payer's documentation says was used? Did the reviewing clinician's credentials actually cover this category of decision at the time of review? Has the record of this determination been altered since it was issued?

Every one of these questions, under the current evidentiary infrastructure, is answered through a combination of policy documents, internal attestations, reviewer testimony, and audit trail reconstruction. The answer, in each case, is produced under pressure, after the fact, by the organization whose behavior is being evaluated. The answer is as credible as that organization is — which is to say, in an adversarial context, often not credible enough to resolve the question.

This is not a criticism of any particular payer, vendor, or clinical reviewer. It is an observation about the structure of the evidence. When the party producing the evidence is the same party whose behavior is being evaluated, and when the evidence takes the form of assertions about what occurred, there is no mechanism by which the evaluator can distinguish a truthful account from a carefully constructed one. The evaluator must either trust the producer or conduct an independent investigation — which, in the volume and technical complexity of modern utilization management, is rarely practical.

1.2 Why the volume of AI-assisted decisioning changes the problem

Traditional utilization management governance was built around individual case review. An auditor could pull a sample of determinations, examine the member record, interview the reviewer, and form an independent judgment. This approach scales poorly — sample sizes are small relative to the population, investigations are expensive, and the conclusions generalize unevenly. But it was adequate when most determinations involved meaningful human judgment and when the number of determinations produced by any single payer was in the hundreds of thousands, not the hundreds of millions.

Two shifts have changed this calculus. The first is volume: AHIP's 2025 industry commitments target 80% of prior authorizations being decided in real time, with leading payer-side vendors reporting that 85% or more of requests are now auto-approved or pre-processed by algorithmic tools. The second is the centrality of AI to decision production. When a model trained on millions of historical cases is the primary mechanism producing a recommendation that a reviewer confirms in seconds, the nature of what governance needs to evaluate has fundamentally changed. It is no longer the reviewer's judgment that carries most of the evidentiary weight — it is the model's behavior, the feature pipeline that produced the model's inputs, and the thresholding logic that turned the model's output into a decision.

None of these things are visible in the governance artifacts the current infrastructure produces. A policy document describes what the model is supposed to do. A SOC 2 report attests that the organization follows its documented processes. A reviewer attestation confirms that the reviewer considered the case. None of these artifacts tell an evaluator anything about what the model actually did on a specific decision — which is the question that matters when the question is asked.

1.3 The specific questions the 2026 environment will demand answers to

The regulatory and legal environment emerging in 2026 will press on the evidentiary gap from multiple directions simultaneously. Each direction asks a specific question that current infrastructure is not structured to answer:

State insurance regulators implementing AI-in-UM legislation (California, Colorado, Texas, New York, and others) are asking whether algorithmic tools are being used in ways consistent with the payer's filed policies. This is fundamentally a question about the relationship between what the payer claims about its algorithms and what the algorithms actually do.

CMS, under CMS-0057-F and its Medicare Advantage oversight authorities, is asking whether AI-driven denials are consistent with coverage determinations and whether the determinations pattern in ways suggestive of improper cost-driven decisioning. Current evidence for either question is derived from aggregate denial statistics, which are suggestive but not probative.

Plaintiff counsel in individual and class-action litigation are asking whether specific determinations, or patterns of determinations affecting their clients, were materially influenced by factors the payer claims do not influence decisions. Discovery currently produces millions of pages of policy documents and internal communications, from which the answer is reconstructed inferentially over years of litigation.

Delegated counterparties — health systems accepting delegated UM from payers, provider-side risk entities, contract administrators — are asking how they can continue to accept delegation when their downstream liability depends on the delegating payer's algorithmic behavior, which they cannot independently evaluate.

Members and their advocates are asking, with increasing public and political force, whether the specific denial they experienced was produced by a process that was supposed to have been structured differently. This is a question that in most cases never gets answered, which is part of what drives the political environment the payer industry now operates in.

In every case, the question being asked is specific, consequential, and structurally difficult to answer using the evidentiary infrastructure that currently exists. The transition the industry is now in the middle of is not principally a transition in technology, though that is part of it. It is a transition in what counts as evidence. That transition is being driven by external pressure, not by industry choice, and organizations that do not adapt to it will find themselves increasingly defending positions with artifacts that no one outside the organization accepts as sufficient.

2. Why Better Attestation Cannot Close the Gap

2.1 The structural limit of self-produced evidence

A common response to the evidentiary gap described above is that the existing attestation framework can be tightened. If SOC 2 reports are not specific enough, more detailed third-party audits can be commissioned. If policy documents are too general, they can be rewritten with more granular commitments. If reviewer attestations lack specificity, they can be augmented with checklist protocols. Each of these responses is reasonable and often appropriate. But each has the same structural limitation: all of them are forms of evidence produced by the organization being evaluated, evaluated against commitments made by the organization being evaluated.

The third-party auditor is retained by the organization, has a commercial relationship with the organization, and produces an opinion constrained by the information the organization provides. Even the best-intentioned audit cannot independently verify what happened inside an AI model on a specific determination without access that the organization may or may not provide, using testing methodology that the organization may or may not permit. The auditor's conclusions rest on what the organization shows them, which is not the same as what actually occurred.

This limitation is not a criticism of auditors or auditing. Financial auditing, safety auditing, and quality auditing all operate within this structural constraint and have developed sophisticated tools for mitigating it — sampling methodology, independent testing, regulatory enforcement backstops, industry-wide standards. These tools work reasonably well in domains where the activity being audited leaves physical or structural traces independent of the organization's own records. They work less well in domains where the activity being audited occurs inside computational systems that can be reconfigured, retrained, or recalibrated without leaving externally observable evidence of the change.

2.2 The specific problem with AI auditability

AI auditability is particularly resistant to traditional attestation mechanisms for several reasons specific to how machine learning systems are developed and deployed.

Models drift. A model committed to production in January 2026 may be retrained in April, further retrained in July, and recalibrated in October. Each version may be nominally the same model but may behave meaningfully differently on similar inputs. A policy commitment made about the January model is not a commitment about the July model, and the organization may not have captured the information needed to reconstruct, months later, whether a specific determination was made by the January version or the July version.

Feature pipelines drift more silently than models. The 147 features used as input to a prior authorization classifier are not raw data — they are the output of an upstream feature-engineering pipeline that normalizes, imputes, bins, and derives. A feature nominally called "clinical severity score" can be silently re-derived to incorporate different underlying inputs without the feature's name, position, or documentation changing. The classifier behaves differently; the attestation remains technically accurate; no external party is positioned to notice.

Thresholding logic is often considered a minor post-processing step but is frequently where cost-sensitivity enters the system. Moving the approve-deny threshold on a continuous confidence score by two percentage points can shift millions of determinations without changing any component of the model the policy documents describe. The thresholding logic is often not covered by the organization's AI-governance attestations because it is not considered part of "the model."

Training data is the single largest source of unobservable bias, and the one most resistant to post-hoc attestation. A model trained on a dataset that correlated cost with clinical severity in its labels will learn that correlation regardless of what policy documents say about cost independence. No amount of post-deployment attestation can verify what the training data actually encoded, unless the training process itself produced evidence at training time that was preserved and is independently verifiable.

Each of these drift vectors is individually addressable by careful organizational practice. Collectively, they make traditional attestation-based AI governance fundamentally incomplete. The gap between what a thorough attestation can cover and what a regulator, plaintiff, or member actually wants to know is not small. It is the difference between "the organization says the model respects cost-independence" and "this specific decision demonstrably did not depend on cost features."

3. What Verifiable Governance Infrastructure Is

3.1 The basic architectural idea

Cryptographic research has, over the past decade, produced a class of techniques generally referred to as succinct proof systems — methods by which a party performing a computation can produce a small

cryptographic certificate attesting that the computation has certain properties, such that any third party can verify the certificate in milliseconds without access to the underlying computation, data, or infrastructure. The mathematical literature refers to these techniques by acronyms including SNARK (Succinct Non-interactive Argument of Knowledge) and STARK (Scalable Transparent Argument of Knowledge). The technical details are beyond the scope of this paper; what matters here is the commercial and evidentiary implication.

Applied to a prior authorization determination, the architecture produces, at the moment the determination is issued, a small cryptographic artifact — typically one to two kilobytes — that attests to specific structural properties of the decision. A third party holding the certificate and the corresponding public verification key can verify the attested properties in roughly 25 milliseconds, without any need for access to the payer's systems, models, or data. The attested properties are not vague claims about process — they are mathematical statements about what did and did not occur in the production of the specific decision the certificate accompanies.

3.2 What the certificate can attest to

A full implementation of verifiable governance infrastructure applied to prior authorization typically attests to four categories of property. Each addresses one of the evidentiary gaps described in Section 1.

Model identity. The certificate binds the decision to a specific, cryptographically committed version of the model that produced it. Any subsequent claim that a different model version produced the decision — an earlier version, a later version, an experimental variant, a manually-adjusted fork — is mathematically refuted by the certificate. This resolves, at the per-decision level, the drift problem described in Section 2.2.

Feature-level independence. The certificate attests that the decision was structurally independent of a declared set of prohibited features — that is, the decision would not have changed if those features had been zeroed. The declared prohibited set is a governance commitment published by the organization and signed cryptographically; the certificate binds each decision to whichever version of the prohibited-features declaration was in force at decision time. This shifts the evidentiary question from "trust the organization's attestation about cost-independence" to "verify mathematically that cost-independence held for this specific decision against the declared registry."

Record integrity. The certificate is chained to prior certificates for the same member, the same model version, and the same workflow, such that no certificate in the chain can be altered or re-ordered without invalidating every subsequent certificate. Allegations that a determination's record was altered after issuance — in response to appeal, litigation, or regulatory inquiry — are refuted by the chain's cryptographic structure.

Reviewer credentialing. For determinations involving human review, the certificate binds the decision to the credential of the reviewing clinician as it existed at decision time. This addresses the increasingly common litigation allegation that decisions attributed to one credential level were in fact produced under a different credential level, or that credentials had lapsed or been narrowed at the relevant time.

3.3 What the certificate cannot attest to

Verifiable governance infrastructure addresses structural properties of a decision's production. It does not — and this is the most important caveat in the entire paper — address clinical appropriateness, coverage correctness, or substantive quality of the determination.

A perfectly valid certificate can accompany a determination that is clinically indefensible. The certificate proves that the specific model ran, that specific features were excluded, that the record has not been altered, and that the reviewer credential was valid. It does not prove that the model's clinical criteria were medically appropriate, that the training data was representative of the member population, that the decision served the member's medical interest, or that the organization's policy was consistent with generally accepted clinical practice. Those are substantive clinical and policy judgments that the cryptography was not designed to evaluate and cannot.

This limitation is not a failure of the architecture. It is a clarification of what problem the architecture solves. The structural-evidentiary gap described in Section 1 is a real and important problem, but it is distinct from the substantive clinical-quality problem that also exists in prior authorization. Verifiable governance addresses the first; it does not pretend to address the second. Organizations that adopt the architecture should be clear-eyed about this boundary, and observers evaluating such adoptions should be equally clear-eyed. A payer that deploys verifiable governance is not thereby absolved of the obligation to make clinically defensible decisions; it is merely making the structural properties of those decisions externally verifiable.

3.4 The completeness problem

One further limitation deserves explicit treatment. The feature-independence proof attests to independence from the features the organization has declared as prohibited. It cannot attest to independence from features the organization has not declared. If a legitimately problematic feature is omitted from the prohibited-features registry, the certificate will validate without catching it.

This is a real limitation, but it is also more usefully viewed as a feature of the architecture than a flaw. Current attestation-based governance allows organizations to make vague commitments about what their algorithms do and do not consider, with no mechanism for external parties to verify the completeness of those commitments. Verifiable governance forces the commitments into a specific, versioned, publicly-posted artifact — the prohibited-features registry — that the cryptography then enforces at every

decision. The registry itself becomes a governance artifact that regulators, members, plaintiffs, and researchers can evaluate and criticize.

If the registry is too narrow — if it excludes only direct cost features and omits demographic proxies, litigation-propensity proxies, or other variables that correlate with payer economics — that narrowness is now visible, commentable, and contestable in a way that no policy document currently is. The rigor of the governance scales with the rigor of the registry, and the registry's rigor is now a matter of public record rather than private assertion.

The same point applies to a related concept: the clinical-confounder declaration. Clinical severity correlates with cost in ways that are medically legitimate — sicker patients require more expensive care — and a proof requiring pure independence from cost would be clinically indefensible. Verifiable governance therefore includes a mechanism for declaring a set of clinical-confounder features through which cost-correlation is permitted. The same transparency logic applies: the confounder set is published, signed, and versioned, and its breadth is subject to the same external evaluation as the prohibited-features registry. A confounder set drawn too broadly becomes a loophole; a confounder set drawn with appropriate clinical rigor becomes a defensible operational choice. The architecture does not make that judgment; it makes the judgment visible to parties who can.

4. Implications for Different Actors in the System

The transition to verifiable-governance infrastructure does not affect all actors in the prior authorization system equally or in the same way. The implications differ substantially depending on the actor's current evidentiary posture and the competitive dynamics of their segment.

4.1 Health plans

For payers, verifiable governance creates a structural differentiation opportunity that is largely unavailable through other infrastructure investments. A payer that publishes a rigorous prohibited-features registry and produces verifiable certificates for each determination can make commitments about algorithmic behavior that competitors relying on attestation cannot match. The commitments have evidentiary force in regulatory inquiry, in litigation discovery, in delegated-UM arrangements, and in member-facing communications.

The commitments are also durable in a way that marketing claims are not. Once a registry is published and certificates begin flowing, any divergence between the registry's commitments and observed behavior becomes mathematically detectable by any third party. This eliminates the possibility of quiet drift between stated policy and operational reality — which is, depending on the payer's strategic posture, either a significant competitive advantage or a significant constraint.

For payers whose current market position depends on algorithmic behavior they would not want to defend in a public registry, the architecture is unattractive and likely to be adopted only under regulatory compulsion. For payers whose algorithmic behavior is, or could become, consistent with a rigorous registry, the architecture offers a differentiation that is hard for competitors to match — particularly for incumbent competitors whose margins depend on opacity they cannot publicly commit to relinquishing.

This asymmetry — verifiable governance favors challengers over incumbents in segments where incumbency depends on opacity — is one of the more commercially significant features of the transition. It shapes who is likely to adopt first, who will adopt defensively, and who is likely to resist until regulation forces the question.

4.2 Providers and provider organizations

For providers, the transition primarily affects their relationship to payer-side AI-assisted decisioning. A verifiable certificate accompanying a determination does not change whether a specific request was approved or denied; it changes what information is available to the provider — and to the provider's appeal, legal, and administrative operations — when a determination is contested.

Current provider appeals operate against an evidentiary fog. The provider knows the determination; they may have access to the determination's cited criteria; but they have no independent mechanism for evaluating whether the determination was produced consistent with the payer's stated policies. An appeal is therefore constructed against the payer's version of what occurred, with the payer controlling both the record and the interpretation.

Certificates change this dynamic at the margin. A provider can verify that a determination was produced by the declared model, respected the declared prohibited features, and was reviewed by a credentialed clinician — or they can observe the absence of such certificates and draw inferences accordingly. In the longer term, provider-facing tools can aggregate certificate patterns across populations, identifying systemic issues more efficiently than case-by-case appeals.

Provider organizations operating at risk — accountable care organizations, risk-bearing provider groups, vertically integrated health systems — face a more acute version of the same question. Their financial exposure depends on payer-side algorithmic behavior they cannot currently evaluate. Verifiable governance makes that behavior partially evaluable, changing the risk-allocation economics of risk-bearing arrangements in ways that are likely to affect contract negotiations.

4.3 State and federal regulators

Regulators are perhaps the actor whose relationship to the evidentiary gap is most directly affected by verifiable governance. The regulatory toolkit in prior authorization has historically been constrained by the difficulty of independently evaluating algorithmic behavior at scale. Enforcement actions have been episodic, resource-intensive, and often limited to the most egregious cases.

A regulatory environment in which payers produce verifiable certificates for determinations transforms the economics of oversight. Certificates are small; verification is fast; the artifacts are inherently machine-readable. A state insurance department or CMS office can, in principle, verify the full certificate record of a Medicare Advantage plan's annual determinations in a computational workflow rather than through document-review audits that take months.

The regulatory implications are twofold. First, regulators can set specific, machine-verifiable expectations — a prohibited-features registry that must include certain minimums, a confounder set that must satisfy certain transparency requirements, a chain-integrity commitment that must cover certain record classes. Second, enforcement becomes substantially more efficient, because violations can be detected computationally rather than through case-by-case audit.

Neither of these implications makes regulation automatic or uncontested. Rigorous regulatory standards for verifiable governance require regulators to develop technical capacity that most agencies currently lack. The appropriate breadth of prohibited-features registries is a policy question that regulators, industry, and public interest advocates will negotiate, with substantial disagreement likely. But the transition moves the regulatory conversation from "can we even evaluate this" — which has been the binding constraint for a decade — to "what should we evaluate and how," which is a fundamentally more productive question.

4.4 Members and member advocates

For members, the transition's implications are indirect but potentially significant over time. Individual members are unlikely to verify certificates on their own determinations — the technical capacity is not consumer-level — but member advocacy organizations, plaintiff's counsel, and investigative journalists can and will. As certificates propagate through the system, the pool of people positioned to answer specific questions about specific determinations grows from "the payer's internal compliance staff" to "any party the member chooses to share the certificate with."

The political implications of this shift are larger than the mechanical ones. Prior authorization has become a politically charged issue in part because specific denial stories resonate in public discourse in ways that aggregate denial statistics do not. A political environment in which member-side organizations can rigorously analyze specific determinations — as opposed to relying on anecdotes and statistics — is a substantially different environment for payers to operate in. Whether this is a better environment or a worse one depends on the observer's values; that it is a different environment is not in question.

4.5 Vendors and the AI-automation ecosystem

The prior authorization technology ecosystem — AI-automation vendors, delegated UM operators, clinical decision support providers, payment integrity platforms — sits in a particularly exposed position in the transition. These vendors sell capabilities whose value depends on payers trusting them with

consequential decisions. Trust has historically been established through references, case studies, and the vendor's internal compliance infrastructure. These mechanisms are not going away, but they are becoming insufficient on their own.

A vendor that produces verifiable certificates alongside its algorithmic outputs offers its payer customers something more than a capability — it offers an independently-verifiable commitment that the capability operates within the payer's governance constraints. This is directly commercially valuable to vendors in enterprise sales cycles, where AI-governance review has become a common procurement gate. It is also defensively valuable in the vendor's own exposure to regulatory and litigation attention, which has been rising.

The ecosystem-level dynamics favor vendors that integrate verifiable governance early. Certificates are compositional — a payer's full-stack governance posture aggregates certificates from multiple vendors, and a vendor that cannot contribute verifiable evidence to that aggregate will be increasingly disadvantaged in enterprise deals. This is not a regulatory requirement; it is a commercial consequence of how downstream buyers will construct their own governance postures as the category develops.

5. The Path Forward

5.1 What adoption actually looks like

Verifiable governance infrastructure is practical at the scale prior authorization operates. Certificate generation adds roughly 800 milliseconds to a determination at commercial volumes — well within existing decision-engine latency envelopes — and certificate size is roughly 1.4 kilobytes per determination, which is negligible relative to the member record itself. Verification on the third-party side runs in approximately 25 milliseconds. These numbers are engineering realities, not projections; they are achievable with currently-available hardware and software infrastructure.

Adoption at organizational scale is less a technical question than a governance and commercial one. A payer implementing verifiable governance must first develop the prohibited-features registry and confounder declarations — a process that typically involves ethics, compliance, clinical, legal, and actuarial leadership, and that surfaces internal questions that many organizations have not previously forced themselves to answer explicitly. The cryptographic infrastructure is straightforward to integrate once these declarations exist; the declarations themselves are the organizational work.

This is a feature of the transition rather than a bug. Organizations that adopt verifiable governance are forced to make their governance posture explicit and durable. Organizations that cannot make such posture explicit — whether because internal agreement does not exist, or because the implicit posture would not survive being written down — will not be able to adopt the infrastructure meaningfully regardless of their technical capacity. The adoption question is therefore a question about organizational readiness for governance legibility, not a question about technology procurement.

5.2 The likely sequencing

Based on the structural dynamics described above, a plausible adoption sequence is visible even without specific forecasting. Early adopters will be organizations whose commercial position is strengthened by governance legibility — challenger payers, AI-automation vendors competing for enterprise payer contracts, risk-bearing providers seeking differentiation in delegated arrangements, and specialized UM operators positioning against incumbents. These organizations have incentive alignment with the transition and face the lowest organizational friction in adoption.

Mid-stage adopters will include regulatory-forward organizations operating in states with aggressive AI-in-UM legislation, national payers with Medicare Advantage books of business facing CMS scrutiny, and organizations whose board-level AI governance commitments have outpaced their operational ability to enforce them. These organizations will adopt because doing so reduces their regulatory and reputational risk, even when doing so is commercially neutral or slightly negative.

Late adopters — potentially forced adopters — will be organizations whose current commercial position depends on evidentiary opacity they cannot publicly commit to relinquishing. For these organizations, adoption is a strategic cost rather than a strategic benefit, and the timing of adoption will be driven by external compulsion rather than internal choice. In payer markets specifically, this likely means the transition is completed across industry rather than within any single firm, with competitive pressure and regulatory enforcement both contributing.

5.3 What actors should do

This paper does not make vendor recommendations. It concludes with observations about what each major actor in the system should consider doing to navigate the transition intelligently, regardless of which specific vendor or implementation they engage.

Health plans should treat the adoption question as a strategic question, not a compliance question. The decision to adopt verifiable governance is materially different from the decision to implement a new audit framework or to update policy documents. It is closer in character to a decision about business-model positioning — whether the organization's competitive advantage is consistent with, or dependent on, evidentiary opacity. Plans that answer this question consciously will navigate the transition better than plans that allow it to be answered by default.

Provider organizations should begin developing internal capacity to evaluate certificates in their own workflow. This does not require deployment of cryptographic infrastructure on the provider side — verification is lightweight and can be integrated into existing provider IT — but it does require organizational decisions about how certificate-based evidence will be used in appeals, contract negotiations, and risk-sharing arrangements. Providers that enter these conversations with a clear view

of what certificates enable will negotiate more effectively than providers who encounter the concept mid-dispute.

State and federal regulators should begin building technical capacity for certificate evaluation. This is a long-horizon capacity-building question, not a short-term procurement question. The agencies that will be most effective in the transition are those that begin now, not those that wait until regulated entities begin producing certificates and the agencies discover they cannot evaluate them. The technical capacity required is modest but specific — it is not a skill set current regulatory agencies typically have in-house, and building it takes time.

Vendors and AI-automation platforms should evaluate where in their product roadmap verifiable governance integration sits. For vendors selling into regulated decision environments, this is not primarily a feature question — it is a question about the durability of the commercial position the vendor currently occupies. A vendor whose competitive position rests on capabilities that can only be attested, not verified, is a vendor whose position is eroding even if they do not notice the erosion yet.

Researchers, public-interest organizations, and member advocates should study the transition as it unfolds. The decisions being made now — about registry breadth, confounder declarations, enforcement standards, adoption incentives — will shape the operational meaning of verifiable governance for years. Advocacy input during the formative period is more effective than advocacy input after standards have solidified.

6. Conclusion

Prior authorization is in the middle of an evidentiary transition. The transition is being driven by external pressure — regulatory, legal, political, and commercial — and is proceeding whether any single actor in the system plans for it or resists it. Verifiable-governance infrastructure is not the cause of the transition; it is one of the tools by which the transition is being made possible. The tool will be used, unevenly and imperfectly, over the next several years.

The organizations and agencies that understand the structural dynamics of the transition early — what verifiable governance can and cannot prove, how it interacts with existing mechanisms, what adoption requires organizationally — will shape the terms on which the transition occurs. The ones that do not will find the terms shaped for them, and will adapt later and at greater cost.

This paper has argued that the evidentiary gap in prior authorization is real, that it cannot be closed by improving the attestation-based infrastructure currently in use, and that the cryptographic techniques required to close it are now practical at industry scale. It has also argued that verifiable-governance infrastructure does not resolve the substantive clinical-quality questions in prior authorization — those remain clinical, legal, and ethical questions that human judgment must answer — and that the

architecture's value depends on the rigor of the governance commitments it makes verifiable, not on the cryptography itself.

The transition is, in the end, a transition from governance-by-assertion to governance-by-verification. It is uncomfortable for organizations whose assertions have been their shield. It is empowering for organizations whose behavior is consistent with the assertions they would make if forced to make them durably. And it is, regardless of any individual organization's posture, the direction the policy, regulatory, and commercial environment is moving — with or without their participation.

The question for each actor in the system is not whether the transition will occur. It is whether they will shape it, or be shaped by it.

About this paper

This white paper is prepared as a neutral analysis of the structural evidentiary questions facing prior authorization governance in the 2026 regulatory environment. It does not endorse specific vendors or implementations. Readers evaluating verifiable-governance infrastructure for their organizations should conduct diligence on specific offerings and should engage their own clinical, legal, compliance, and actuarial leadership in the registry and confounder-set design process. The paper's conclusions are the author's and should not be attributed to any organization whose work is referenced or described.