Evidence architecture · Enterprise AI commercialization

Why AI Governance Is a Revenue Strategy, Not a Compliance Cost

How evidence architecture converts validation, governance, and monitoring into decision-ready commercial infrastructure.

Olga Lavinda, PhD · Lawrence Meche, PhD

Health AI LLC · March 2026 · revised July 17, 2026

The governance artifact is not the policy. It is the traceable chain from claim → requirement → evidence → decision → owner → response.

AI loses commercial value at the handoffs between model, workflow, authorization, coverage, procurement, and monitoring. Evidence architecture preserves decision-grade evidence across those handoffs.

Governance does not guarantee revenue. It makes value traceable, reviewable, reusable, and defensible.

PERMANENT IDENTIFIER

doi.org/10.5281/zenodo.21403849

constat.dev/whitepaper

SCOPE

Healthcare and automotive

Comparison, not regulatory equivalence

0 Executive summary

The model is only one part of the commercial system.

Abstract

Enterprise AI use is widespread, but enterprise-level financial impact remains limited. McKinsey's 2025 survey found that 39% of respondents attributed any EBIT impact to AI, while nearly two-thirds said their organizations had not begun scaling AI across the enterprise.[2] RAND's interview research identifies recurring breakdowns in problem definition, data, infrastructure, organizational capability, and stakeholder alignment.[1]

This report addresses the mechanism behind that gap. We define **evidence architecture** as the requirements, records, responsibilities, validation, and monitoring that allow a system to answer the distinct questions posed by regulators, payors, procurement teams, operators, customers, and executives. Governance does not guarantee revenue. It makes value traceable, reviewable, reusable, and defensible.

Contribution

The argument is not merely that governance matters. It is that **evidence loss at decision handoffs** is a commercial failure mode—and that prospective evidence reuse is the operating mechanism by which governance can create leverage.

Performance is not portability

A strong model can remain commercially inert if its evidence cannot survive the next review, buyer, setting, or change.

Connected does not mean identical

Authorization and coverage/payment can share evidence, but they ask different questions and apply different standards.

Reuse is the leverage point

One governed record can support multiple decisions without pretending that one approval answers every downstream need.

Late evidence is expensive evidence

Missing baselines, endpoints, ownership, and monitoring plans often cannot be reconstructed after deployment.

RAND and McKinsey establish context, not causality. Neither source proves that governance alone produces return; this report advances a bounded operating thesis rather than a universal ROI claim.

1 The productive function of governance

Governance creates value when it produces decision-ready evidence.

Governance is often framed as control: a policy, committee, approval gate, or restriction. Those mechanisms can matter, but the enterprise does not receive value from their existence alone. The productive layer is the record they create—intended-use requirements, acceptance criteria, source provenance, validation results, human-oversight design, change history, monitoring thresholds, and named accountability.

Decision systems served by a reusable evidence record

Decision-maker	Question that must be answered	Reusable evidence artifact
Regulator	Is the device safe and effective for its intended use?	Requirements, validation, labeling, risk controls, change rationale
Payor	Is the service beneficial, covered, coded, and payable in this context?	Clinical utility, outcomes, comparator, code, setting, policy context
Procurement	Can this system be operated safely and accountably here?	Controls, security, workflow testing, monitoring and response plan
Operator	What should I do, what are the limits, and when do I escalate?	Human factors, intended use, limitations, thresholds, playbook
Executive	Is value real, durable, and proportionate to risk?	Outcome measures, cost, adoption, incidents, ownership, trend

These questions are related, but no single artifact answers all of them. Evidence architecture is therefore interface design between decision systems. It makes the differences explicit early enough to plan for them, while preserving shared provenance so that the organization does not repeatedly rebuild the same factual record.

Where the economic leverage appears

The revenue mechanism is indirect but concrete: earlier discovery of missing assumptions; less duplicated evidence work; faster and more consistent review; clearer procurement readiness; measurable post-deployment outcomes; and a record that can support change without erasing what was previously known.

> **Evidence architecture is interface design between decisions—not extra documentation around the model.**

That chain cannot rescue a product without clinical utility, workflow fit, technical performance, responsible leadership, or a viable payment pathway. It can prevent an otherwise viable product from losing value because the evidence required by its next decision-maker was never designed, captured, or retained.

2 The authorization-to-value gap

One device. Two connected decision gates.

FDA's principal role is evaluating device safety and effectiveness. After authorization, payors and providers make separate decisions about coverage, payment, use, and recommendation. FDA notes that evidence submitted for authorization may not always overlap with the evidence payors need, contributing to delayed or denied coverage.[3]

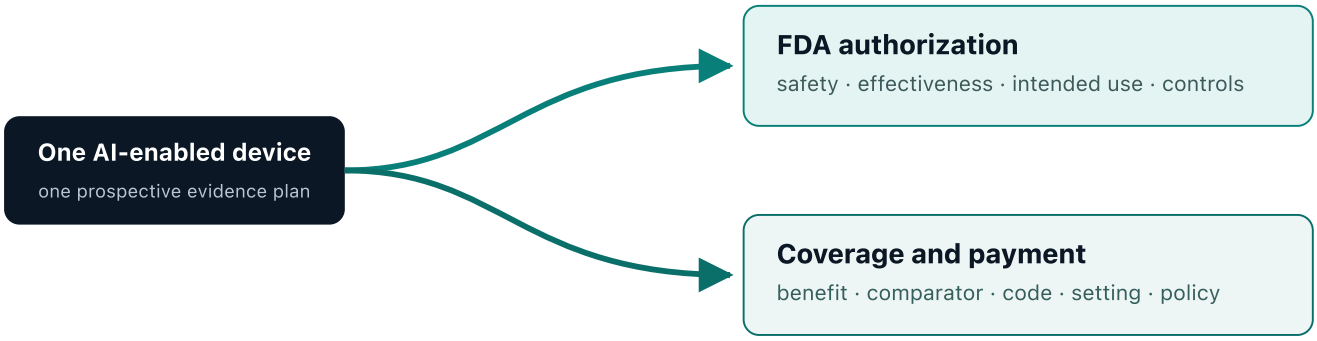

Evidence can overlap; the decision standards do not collapse into one another. Source: FDA payor-communication guidance.[3]

Authorization question

Does the record support safety and effectiveness for the intended use, with appropriate risks, controls, and labeling?

Coverage/payment question

Does the service provide the required benefit in context, with an applicable code, setting, policy, and payment mechanism?

FDA's Early Payor Feedback and Parallel Review pathways exist because earlier coordination can help sponsors plan evidence for regulatory and coverage decisions.[3] The strategic error is not failing to make one study satisfy every audience. It is reaching authorization before learning that a necessary downstream endpoint, comparator, baseline, or economic measure was never collected.

RAPID makes the same commercial logic explicit. Announced by FDA and CMS in April 2026 for certain Breakthrough Devices, the pathway brings regulatory and Medicare coverage actors together earlier so evidence generated for FDA review may also support coverage decisions.[4] It is a targeted pathway, not universal reimbursement for AI.

CMS payment remains specific to service, code, setting, practitioner, and coverage policy. The 2026 Physician Fee Schedule is a payment rule—not a universal reimbursement pathway for "AI."[7]

> **The validation strategy fails commercially when it reaches the intended approval but leaves the next decision impossible to make.**

3 Two sectors, one structural problem

Evidence must survive deployment, change, and review.

Healthcare and automotive are not regulatory equivalents. They are useful comparators because both place adaptive or software-mediated systems into high-consequence workflows, where performance claims must remain bounded by intended conditions and supported after release.

Cross-sector evidence architecture—not legal or regulatory equivalence

Healthcare · medical-device AI	Automotive · automated driving	Shared structural pattern
FDA authorization for an intended use	Vehicle-safety and UNECE regulatory regimes	**Regulatory gate**
Coverage, payment, provider adoption, and procurement	Insurance, operator, procurement, and market adoption	**Commercial gate**
Validation, real-world evidence, and postmarket monitoring[5][6]	Operating-domain evidence, field events, software changes, and corrective action[8]	**Runtime evidence**
Intended use, population, clinical setting, and workflow	Defined operating conditions, users, environment, and limits	**Bounded performance**

FDA's 2025 draft guidance proposes a total-product-life-cycle approach for AI-enabled device software. It is nonbinding draft guidance and should not be represented as a universal postmarket mandate.[5] FDA's final real-world-evidence guidance separately addresses when real-world data may be fit to support device regulatory decisions.[6] In automotive, UNECE's WP.29 adopted automated-driving regulatory instruments in June 2026.[8]

The common signal is operational, not doctrinal: decision-makers increasingly need traceable claims, defined operating conditions, validation, change history, and a response when real-world behavior departs from expectation.

> **High-consequence AI is not governed once. Its evidence must remain coherent as the system, environment, and decision context change.**

The comparison does not imply a shared regulator, liability standard, payment model, or evidentiary threshold. It identifies a shared enterprise architecture problem.

4 The revenue arithmetic of governance

Governance changes economics by changing evidence flow.

The phrase “revenue strategy” should not be read as a claim that documentation causes sales. The defensible mechanism is more specific: evidence architecture can alter the cost, speed, consistency, and optionality of the decisions required to commercialize and sustain AI.

How evidence architecture can create commercial leverage

Value lever	Operating mechanism	Boundary of the claim
Less rework	Shared requirements, provenance, and validation records reduce repeated reconstruction across reviews	Only when artifacts are trustworthy, current, and fit for the new decision
Earlier clarity	Distinct regulator, payor, buyer, and operator questions are mapped before endpoints become unrecoverable	Planning cannot guarantee authorization, coverage, or purchase
Faster review	Claims can be traced to source evidence, owner, decision, and change history	Review speed still depends on reviewer, scope, quality, and novelty
Stronger procurement	Limitations, controls, workflow testing, monitoring, and response ownership are visible	Evidence cannot substitute for product fit, security, or acceptable economics
Durable value	Runtime outcomes and incidents remain connected to the original intended use and acceptance criteria	Monitoring reveals performance; it does not manufacture benefit

Commercial leverage = evidence reuse + earlier decision clarity + measurable outcomes – rework – evidence latency

Decision heuristic, not an empirical financial formula

McKinsey's highest-performing cohort was substantially more likely to report fundamental workflow redesign.[2] That is consistent with the operating thesis here but does not prove it. Workflow, technical quality, leadership, data, economics, and evidence architecture interact; governance is one part of the causal system.

What this thesis rejects

- That every failed AI dollar can be attributed to absent governance.
- That an authorized product is automatically covered, billable, or commercially viable.
- That documentation is an automatic legal defense or a substitute for safe design.
- That one approval artifact can answer every downstream decision unchanged.

> **Governance does not monetize the model. It prevents the evidence needed to monetize, operate, and defend the system from disappearing between decisions.**

5 From thesis to operating model

RIGOR: a reusable evidence lifecycle.

Health AI's RIGOR framework organizes evidence work into five connected domains. Its purpose is not to collapse regulatory, commercial, and operational standards into one checklist. It is to preserve a shared, inspectable record while making each decision's distinct criteria explicit.

Evidence Architecture
The commercial principle: evidence survives decisions

→

RIGOR
The operating method

→

Constat
The evidence system for FDA AI devices

Requirements

Define intended use, users, decisions, failure costs, success criteria, and the evidence each decision-maker will need before architecture or model selection.

Implementation

Record data provenance, design decisions, model changes, human oversight, limitations, and the rationale connecting the system to its workflow.

Governance

Assign owners, review rights, escalation paths, approval gates, and change-control responsibilities. A policy without operating ownership is not a control.

Operational proof

Test with representative users, populations, settings, and workflows. Technical performance and operational utility answer different questions.

Runtime monitoring

Predefine drift, safety, equity, workflow, and outcome signals; preserve the source record; and specify who responds when thresholds are crossed.

The practical advantage is prospective reuse. A monitoring plan designed after launch cannot recover every missing baseline. A coverage strategy designed after authorization may discover that the necessary endpoint was never measured. Evidence planning does not make requirements identical; it makes their differences visible soon enough to act.

Constat: evidence architecture in practice

Constat operationalizes this evidence architecture for FDA AI/ML-enabled medical devices. It connects source-quoted premarket records, postmarket signals, reimbursement pathways, and compliance evidence in one queryable system. The browser console and Constat MCP expose the same evidence base to people and AI agents while keeping denominators and source context near the claim.

constat.dev · constat.dev/connect · healthai.com/rigor

Scope boundary: RIGOR is the authors' method; Constat is its FDA-device implementation. This report describes the operating thesis; it does not claim experimental proof of a universal governance-return effect.

6 References

Primary sources and source notes

1. **RAND Corporation (2024).** *Why AI Projects Fail.* Interview-based analysis; “more than 80%” is explicitly qualified as an estimate. RAND_RRA2680-1.pdf
2. **McKinsey & Company (2025).** *The State of AI in 2025: Agents, Innovation, and Transformation.* mckinsey.com/…/the-state-of-ai
3. **U.S. Food and Drug Administration (2026).** *Medical Device Coverage Initiatives: Connecting with Payors.* fda.gov/…/payor-communication-task-force
4. **U.S. Food and Drug Administration (2026).** *CMS and FDA Announce RAPID Coverage Pathway to Accelerate Patient Access to Life-Changing Medical Devices.* fda.gov/…/rapid-coverage-pathway
5. **U.S. Food and Drug Administration (2025).** *Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations.* Draft; not for implementation; nonbinding recommendations. fda.gov/…/lifecycle-management
6. **U.S. Food and Drug Administration (2025).** *Use of Real-World Evidence to Support Regulatory Decision-Making for Medical Devices.* Final guidance. fda.gov/…/real-world-evidence
7. **Centers for Medicare & Medicaid Services (2025).** *Calendar Year 2026 Medicare Physician Fee Schedule Final Rule.* cms.gov/…/cms-1832-f
8. **United Nations Economic Commission for Europe (2026).** *Working Party on Automated/Autonomous and Connected Vehicles—Recent Activities.* unece.org/…/connected-vehicles

Suggested citation

Lavinda, O., & Meche, L. (2026). *Why AI Governance Is a Revenue Strategy, Not a Compliance Cost* (revised July 17, 2026). Health AI LLC. https://doi.org/10.5281/zenodo.21403849

About the authors

Olga Lavinda, PhD is Founder and CEO of Health AI LLC and developer of the RIGOR evidence lifecycle. Her work spans AI literacy, evidence-grounded institutional adoption, and high-consequence AI systems. lavinda.ai

Lawrence Meche, PhD is Chief Technology Officer of Health AI LLC. His work spans industrial inspection, scientific computing, and validation methods for AI used in critical infrastructure.

> **Author disclosure.** The authors are principals of Health AI LLC, which develops the RIGOR method and Constat product discussed in this report. This report is descriptive industry analysis, not legal, regulatory, reimbursement, coding, coverage, or investment advice.