




The business should retain the procedures, evidence, and acceptance standards that make AI useful as models improve and economics change.

David C. Robinson

CEO and Co-Founder, Human AI, Inc.

October 2026 · [Lectotype.com](https://lectotype.com)

The standard against which your AI is measured.

Your Business Rules Should Survive a Model Change

A business may change an AI model because another one performs better, costs less, or fits its deployment requirements. The company's approval rules should not have to be reinvented each time. Neither should the procedure for preparing a proposal or the evidence required to accept a report. Those are assets the business has developed by learning how to do its work.

Cohere CEO Aidan Gomez made the ownership argument in September: *"No government or enterprise should have to choose between capable AI and control over their technology."* His statement concerned sovereign AI deployment. The same concern reaches into everyday operations, where a company needs control over the procedures and records surrounding the intelligence it buys. Model access is one part of that operating architecture. [1]

Meta CEO Mark Zuckerberg emphasized another advantage of organizational independence: *"Since the models are open, anyone is capable of testing for themselves as well."* His argument supports access to independent testing. A business also needs its own standard for what those tests should establish, whether it uses an open model or a hosted service. A published benchmark cannot determine whether an agent respects a particular customer's pricing rules. [2]

Lectotype is being developed to preserve that business standard around the model. The organization defines its approved evidence, workflow steps, permissions, and acceptance criteria. A model is assigned the parts of the task that benefit from judgment or synthesis. Known calculations and procedures remain in software. The resulting system gives the business a place to retain what it has learned as the available models change.

Satya Nadella's separation of intelligence supply from authority is especially useful here. He calls for model diversity and independent controls, with the organization determining access and permitted actions. For a client, those principles mean that choosing a model and authorizing a business action are separate decisions. A new model can enter the evaluation process without acquiring the right to rewrite the company's procedure or approve its own release. [3]

Take a workflow that qualifies an inquiry, drafts a response, and prepares a proposal. A more capable model may help interpret an unfamiliar request, while a less expensive model may handle a well-defined drafting step. Duplicate checks, routing, and price calculations can run as ordinary software. The correct arrangement depends on the

work and the acceptance standard. We will test those choices with client cases rather than assume that one model belongs in every step.

The comparison has to include the cost of reaching an acceptable result. A lower price per token may be outweighed by retries, longer review, or repairs. A more expensive model may justify its cost when it resolves cases that otherwise require a specialist. The client's test should hold the inputs, permissions, and review standard constant, then measure accepted outcomes and the work required to achieve them.

Our compiling capability architecture follows that logic by retaining reusable procedures and checks in the surrounding program. Once a team establishes how to gather evidence, calculate a value, or handle a known exception, it can preserve that work as an explicit operating asset. The next model then works within the tested procedure. Learning becomes something the organization can inspect, revise, and carry forward across suppliers.

AI21's Ori Goshen connects enterprise adoption to the system around the LLM and to evaluation of the operating result. That supports a practical purchasing question: which combination of model and workflow meets our standard at an acceptable total cost? The answer may change as models improve. The organization should retain the evidence and procedure needed to make that decision again. [4]

Portability requires an evaluation before a substitute model is released. It may interpret a tool differently, produce different structures, or fail on cases the previous model handled. We will rerun representative work and planted violations, compare review effort and acceptance, and preserve the configuration behind the result. The business owner will then decide whether the new configuration meets the agreed standard.

Our first Lectotype client engagements will make that process concrete on one bounded workflow at a time. We will establish the baseline, identify the model work and software work, and test the choices against the client's requirements. Those engagements will prioritize the integrations, measurements, and release features in our sprint plans. If your business wants the freedom to change models while retaining its operating knowledge, bring a workflow whose success you can define.

Sources

[1] Aidan Gomez. [Cohere and Aleph Alpha sign agreement to become the first transatlantic sovereign AI solution](#). September 16, 2026.

[2] Mark Zuckerberg. [Open Source AI is the Path Forward](#). July 23, 2024.

[3] Satya Nadella. [Models as Insider Risks in the Super Intelligence Era](#). October 10, 2026.

[4] Ori Goshen. [2024 Enterprise AI Forecast: From Pilot to Production](#). January 9, 2024.