




A generated proposal becomes a business commitment when it leaves your system. Lectotype puts the authority to release it into the workflow.

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.

What Happens When You Push Send

A sales agent has prepared a proposal and is ready to send it. The customer's name is correct, the language reads well, and the offer looks reasonable. Before it leaves the company, someone still has to establish whether the price is approved, the claims are supported, and the sender has authority to make the offer. That moment is where AI capability becomes a business commitment.

OpenAI CEO Sam Altman and chief scientist Jakub Pachocki describe the governing objective in their joint essay: *“Powerful systems must remain safe, aligned with human intent, and subject to human control.”* Businesses have to turn that objective into an operating procedure. For a proposal, human control includes deciding which commitments an agent may make and which require a person's approval. It also includes controlling the tool that sends the message. [1]

At Lectotype, we begin before the model drafts the proposal. The workflow establishes the user's identity, the purpose of the task, the relevant business rules, and the evidence that may be used. Approved product claims and current prices become part of a bounded task. A model can then help compose the response within that context. The release decision remains part of the surrounding procedure.

This order follows the principle Satya Nadella sets out in his essay on models as insider risks. He writes that *“the controls that govern what a model can access and what actions it can take must sit outside the model.”* A model's willingness to follow an instruction is useful, but the business needs the action itself to pass through its authorization mechanism. That is why the sending integration belongs in the governance design. [2]

Suppose the customer asks for a discount above the sales team's permitted limit. The agent may recognize a persuasive commercial reason for granting it, and the draft may be excellent. The workflow should still require the named approver before release. If approval is missing, the disposition is a request for that approval, with the proposed terms preserved for review. The client defines the threshold and the people authorized to cross it.

The same procedure needs to handle evidence that has changed. A product description may be superseded, a customer record may be incomplete, or a price schedule may have expired. The workflow should identify the unresolved condition and route it to the appropriate person. Deterministic checks are useful for known requirements, while claims that require judgment need a defined review method. Each condition has to be resolved before the business accepts the proposal.

NVIDIA's agent security research explains why enforcement must survive changes in the route an agent takes. An agent may use a different tool, write code, or delegate a task, and each path can carry the same unauthorized action. Their recommendation is to enforce policy outside the agent's control and cover permitted intermediaries as well as direct transfers. A proposal workflow therefore needs a controlled route to the customer, with credentials scoped to the authorized work. [3]

The client should also be able to inspect the record after release. Which product information supported the claims? Which price version applied, and who authorized the exception? What was actually delivered, and when? Those are operating facts that should be recorded by the system. An agent's explanation can assist a review, but the preserved evidence and execution record must carry the decision. [2]

Microsoft AI CEO Mustafa Suleyman describes advanced AI as “*carefully calibrated, contextualized, within limits.*” A sales workflow gives those words a practical meaning. The agent can take on preparation and follow-up while the organization retains the boundaries around its commitments. Over time, the business can expand the permitted work as the workflow meets its acceptance standard and the responsible owner approves the change. [4]

We are developing Lectotype's first commercial engagements around bounded workflows such as lead response, follow-up, and proposal preparation. With each client, we will agree the rules, connect the required systems, and test ordinary work alongside planted violations. The engagement will expose the integration and reporting work that belongs in our next sprints. Deployment will proceed through the agreed acceptance and release decisions.

The demonstration we want a client to see includes a proposal that is allowed to proceed and one that must stop for approval. Both should leave a record that the sales leader can understand. Bring us a workflow where your team already knows what has to be true before a message goes out. We will use that standard to define what happens when you push send.

Sources

[1] Sam Altman and Jakub Pachocki. [Built to benefit everyone: our plan](#). June 8, 2026.

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

[3] NVIDIA AI Security. [Assume Misalignment: Securing AI Agents Beyond Prompt Injection](#). Accessed October 11, 2026.

[4] Mustafa Suleyman. [Towards Humanist Superintelligence](#). November 6, 2025.