

Human Oversight Has a Capacity Limit

A Board Decision on Review That Can Operate

Phil Hatch | Akholi

October 23, 2026

Boards should require evidence that the people assigned to oversee outsourced agents can perform the required review at the volume and speed of delivery. Naming a human reviewer does not establish an effective control.

Ask Whether the Review Can Be Performed

The board should require management to demonstrate the capacity behind each material human-oversight commitment. The control depends on qualified people receiving useful evidence with enough time and authority to intervene. A diagram showing a person between the agent and the outcome is not proof that those conditions exist.

The study reports review capacity mismatched to agent output and routine approval at human review points.¹ These are vendor-team accounts, not a verified count of failed bank controls. They justify an immediate management test of whether the organization's own oversight design can operate under actual delivery conditions.

The issue is particularly acute when a provider can increase output faster than either firm can increase qualified review. A review obligation that worked during a small pilot may become nominal at production volume. The board should ask what changes when volume rises, exceptions cluster, or a reviewer is unavailable.

¹ Phil Hatch, *Agentic AI Outsourcing at Tier 1 Financial Institutions: Full Study Report*, version 1.8.2 (Akholi, forthcoming manuscript dated November 1, 2026), Part 2, secs. 4.2–4.3, “Human Review Capacity Against Agent Output Volume” and “Routine Approval of Agent Output at Human Review Points,” <https://doi.org/10.2139/ssrn.7164559>.

Specify the Judgment the Person Must Make

Management should distinguish approval before action, review of exceptions, sampling after action, and continuing professional collaboration. These serve different purposes and permit different exposure. Sampling after a consequential event cannot be described as a control that prevents that event.

Stanford's selected enterprise cases describe oversight as a context-dependent design choice, not simply a sign of immature automation.² The board should therefore request a rationale tied to the consequence of error, required judgment, recoverability, and applicable obligations. It should not seek one automation percentage for all services.

For each review point, state what the person must inspect, what evidence is available, how long the task takes, and what happens when the person disagrees or cannot decide. A reviewer who can only accept an output, without access to its basis or authority to stop the action, is performing a different function from the one the board may assume.

Test Capacity Under Representative Demand

Figure 1 proposes a capacity decision record. Measure arrivals, review effort, available qualified time, backlog age, and escalation demand. Use representative difficult cases and peak conditions, including correlated failures. The purpose is a management decision about feasible coverage, not a statistical exercise for directors.

A hypothetical team receiving 600 cases that each require two minutes of review faces 20 hours of review work before breaks, investigation, or other duties. Two reviewers with six usable hours each cannot complete that obligation in the period. The arithmetic is illustrative; management must obtain the service's actual workload and effort evidence.

Do not solve a capacity gap merely by shortening the permitted review time. Test whether the judgment remains reliable and whether unresolved cases accumulate. The provider's assurance dashboard should distinguish completed review, deferred work, sampled work, and automatic approval so the buyer can see which control operated.

² Elisa Pereira et al., *The Enterprise AI Playbook: Lessons from 51 Successful Deployments*, Stanford Digital Economy Lab, April 2026, methodology and chapters 1–3 and 8; purposive successful-case sample, https://digitaleconomy.stanford.edu/app/uploads/2026/03/EnterpriseAIPlaybook_PereiraGraylinBrynjolfsson.pdf.

Figure 1. The human oversight capacity decision

Capacity element	Board question	Management evidence
Judgment	What must the reviewer decide	Review purpose and required evidence
Demand	Can qualified time cover the work	Arrivals effort peaks and backlog
Authority	Can the person intervene in time	Tested restriction and escalation
Scale	When must the service slow or stop	Capacity trigger and funded response

Author's proposed operating framework. Adapt to the service and applicable jurisdiction; not an empirically validated standard.

Source: Phil Hatch, proposed framework developed for this paper.

Choose a Feasible Oversight Design

Management can narrow agent authority, reduce volume, improve evidence presentation, add qualified capacity, redesign the workflow, or change the role of human review. These choices have different cost and risk consequences. The decision should identify which error the review is intended to prevent and whether another control can perform that function more reliably.

BCG's governance analysis emphasizes runtime enforcement rather than relying solely on manual signoffs.³ Technical boundaries can reduce the work people must review, but they do not automatically replace professional or legally required decisions. Test the boundary and preserve human authority where the service requires it.

EU, UK, Swiss, and US obligations differ by sector and activity. Management should identify where a qualified person must make or review a decision and what evidence that person needs. A group staffing ratio cannot substitute for local professional, privacy, conduct, or safety requirements.

³ Shaheer Rizvi et al., "Enterprise AI Control Plane: The CIO's Guide to Governing and Accelerating AI Agents," Boston Consulting Group, August 14, 2026, "Technical Components of the EACP" and "Lessons Learned," <https://www.bcg.com/publications/2026/how-cios-govern-ai-agents-at-scale>.

Fund the Capacity Before Approving Scale

An objection is that maintaining reviewers could consume the automation benefit. That may be true for a particular design and should appear in the business case. Management should then change the design or limit its scope, rather than approve a control that cannot be performed and count the unprovided work as a saving.

The committee should receive the oversight purpose, tested capacity, peak constraints, backlog limits, intervention authority, and action plan for any shortfall. Require a trigger that restricts the service when oversight demand exceeds safe capacity. Expansion should reopen the capacity decision before the additional work arrives.

Human oversight is a funded operating capability. The board should approve it on demonstrated judgment and throughput, with its limits visible. A named reviewer becomes a credible control only when the person can understand the case, make the required decision, and act before the consequence exceeds the approved boundary.

Evidence and Source Note

The study is a nonrepresentative cross-section of 850 vendor delivery-team respondents on 150 engagements serving Tier 1 financial institutions, collected during January–June 2026. Neither banks nor vendors corroborated the accounts. It establishes reported experience, not causation, legal breach, or market prevalence. Its application here is a mechanism to investigate, not an estimate of conditions in other industries. The author-supplied version 1.8.2 is dated November 1, 2026 and is cited as a forthcoming manuscript consulted October 23, 2026.⁴

This is an analytical decision paper based on the supplied study and selected primary public sources checked through October 23, 2026. The original contribution is an executive oversight capacity decision record linking required judgment, demonstrated workload, intervention authority, and limits on service expansion. The proposed method has not been independently validated. Legal requirements depend on the entity, activity, jurisdiction, and effective provisions; supervisory observations, research, and professional commentary have distinct authority.

Bibliography

Hatch, Phil. *Agentic AI Outsourcing at Tier 1 Financial Institutions: Full Study Report*. Version 1.8.2. Akholi. Forthcoming manuscript dated November 1, 2026. Author-supplied manuscript consulted October 23, 2026. <https://doi.org/10.2139/ssrn.7164559>.

Pereira, Elisa, Alvin Wang Graylin, and Erik Brynjolfsson. *The Enterprise AI Playbook: Lessons from 51 Successful Deployments*. Stanford Digital Economy Lab. April 2026. https://digitaleconomy.stanford.edu/app/uploads/2026/03/EnterpriseAIPlaybook_PereiraGraylinBrynjolfsson.pdf.

Rizvi, Shaheer, Zeeshan Shah, Sarvesh Singh, Nicolas Tanaka, Francois D'Agostini, and Jazz Tobaccowalla. "Enterprise AI Control Plane: The CIO's Guide to Governing and Accelerating AI Agents." Boston Consulting Group. August 14, 2026. <https://www.bcg.com/publications/2026/how-cios-govern-ai-agents-at-scale>.

⁴ Hatch, *Agentic AI Outsourcing*, Introduction, "Executive Summary" and "Evidentiary Basis and Its Limits."