

Put the Approval Gate Before the Irreversible Action

A Manager Method for Timely Authorization

Phil Hatch | Akholi

October 23, 2026

Delivery managers should place approval before the point at which the business consequence becomes difficult or impossible to reverse. The control must bind the approval to the exact action and prevent execution when approval is missing, expired, or no longer applicable.

Find the Point That Changes the Outside World

The service manager should identify the event after which an agent action creates a consequential external effect. Sending a message, releasing a payment, changing access, submitting an application, and deleting a record can have different reversal limits. The approval point should be designed around that event rather than placed wherever the workflow tool offers a review screen.

NIST's identity analysis distinguishes delegated entitlements from broad credential sharing.¹ A valid identity can still attempt an action beyond the intended business approval. The buyer and provider therefore need a separate authorization decision tied to the action that will be executed.

For a hypothetical supplier-payment service, approval of a draft invoice should not authorize a later payment with a changed account or amount. The manager should identify which changes invalidate the approval and which system prevents execution until the required authority is restored.

¹ Bill Fisher and Ryan Galluzzo, "Back to the Future: Why Agentic AI Needs a Strong Identity Foundation," National Institute of Standards and Technology, August 27, 2026, "Credential Sharing" and "Static and Long-lived Credentials," <https://www.nist.gov/blogs/cybersecurity-insights/back-future-why-agentic-ai-needs-strong-identity-foundation>.

Bind Approval to a Specific Proposed Action

Record the object, destination, amount or scope, relevant evidence, policy version, approving authority, and validity period. The receiving execution mechanism should be able to establish that the approval applies to the current action. A generic approval flag in a separate system creates a gap if the proposal can change afterward.

The study reports review constraints and joint decisions that could not keep pace with agent output.² Managers should test whether the actual approval route can operate within the service deadline. The response to delay should be a safe restriction or designed alternative, not an implicit permission to proceed.

Distinguish business approval from technical checks. A human may decide that an action is appropriate while automated controls enforce destination, aggregate limits, and expiry. Neither layer should silently waive the other. A reviewer's authority and the technical control's scope should be explicit.

Test the Gate as an Enforcement Boundary

Figure 1 proposes four tests: missing approval, changed proposal, expired authority, and bypass through another route. Observe the receiving system, not only the agent's explanation. The service should demonstrate that the consequential action cannot occur outside the approved conditions through the tools and delegates available to it.

Include retries and concurrent work. An approval intended for one action should not authorize several submissions because the first acknowledgment was delayed. Preserve the relation between the approved instruction, execution attempt, receiving result, and any cancellation. Uncertainty about an earlier attempt should trigger reconciliation before another consequential submission.

BCG's control-plane analysis includes runtime enforcement at the execution boundary.³ Applied to outsourcing, both firms should agree where that boundary exists and who can change it. A buyer approval screen is insufficient if a provider administrator can bypass it through a different tool or credential.

² 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>; Hatch, *Agentic AI Outsourcing*, Part 2, sec. 5.2, "Joint Decision Speed Mismatched to Agent Speed."

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>.

Figure 1. The irreversible action approval test



Author's proposed method; adapt to the service and applicable jurisdiction.

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

Design the Exception Route Before Production

Define what happens when the approver is unavailable, information is incomplete, or a deadline is approaching. The service may stop, escalate to an authorized alternate, or use a narrower preapproved route. Each option should preserve the governing authority and leave a record. Urgency alone should not create a new mandate.

EU, UK, Swiss, and US services may have activity-specific approval, professional, financial, privacy, and record duties. The manager should obtain the relevant applicability conclusions. A contractual approval can support those duties but cannot substitute for a decision that must be made by a particular authorized person or institution.

Where an action is reversible only within a short interval, record that interval and the conditions for using it. A theoretical reversal option is not a preventive control if the organization cannot detect the problem, decide, and act before the option expires.

Approve the Gate Only After a Joint Demonstration

The practical objection is that prior approval can eliminate the speed benefit. Classify actions by consequence and uncertainty. Narrow preauthorization with enforceable limits may support fast routine work, while material changes require a fresh decision. The manager should measure the resulting service performance rather than apply the same manual gate everywhere.

Before release, demonstrate permitted execution and each relevant failure condition. Record bypass routes, residual uncertainty, and the person accepting restrictions. Reopen the assessment when tools, delegates, workflows, or business authority change. An old test should not be carried forward into a materially different action path.

The approval gate is effective when it precedes the consequence, applies to the action executed, and cannot be bypassed within the available service authority. The manager should approve production on that evidence. A review record created after an irreversible action may support investigation, but it should not be represented as prevention.

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 a cross-firm approval acceptance test binding authority to the actual consequential action and testing changes, expiry, retries, and alternate execution routes. 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

Fisher, Bill, and Ryan Galluzzo. “Back to the Future: Why Agentic AI Needs a Strong Identity Foundation.” National Institute of Standards and Technology. August 27, 2026. <https://www.nist.gov/blogs/cybersecurity-insights/back-future-why-agentic-ai-needs-strong-identity-foundation>.

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>.

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.”