

Redesigning Agentic Outsourcing

Optimize the Whole Delivery Method

Phil Hatch | Akholi

October 23, 2026

Agentic AI changes the inputs, decisions, review, exceptions, and acceptance surrounding its assigned task. Buyers and providers should map and redesign the complete delivery method, then prove its business outcome before expanding autonomy.

Approve the Whole Delivery Method Before Expanding Autonomy

Agentic outsourcing requires a delivery method designed for the complete business function. Replacing one process step with an agent can accelerate execution while leaving the service constrained by poor inputs, delayed buyer decisions, insufficient review, or unresolved exceptions. Leadership should therefore judge the accepted business outcome across both firms and their dependencies. A faster agent task does not establish that the redesigned service can meet its obligations.

An illustrative invoice service explains the distinction. An agent extracts and checks invoice details, but unmatched items wait for the buyer to confirm the supplier or purchase order. Provider staff repair inputs and pursue decisions while a dashboard counts completed checks. The unresolved queue can grow even as the reported task becomes faster. Improving the function requires changes to source quality, decision authority, response time, exception ownership, and acceptance. This example illustrates the mechanism; it is not a measured study case.

Executives should commission a joint delivery-model audit, appoint a buyer executive accountable for the functional outcome, and require a current-state map, an approved design, and bounded operating evidence before further expansion. The first month should establish the material gaps and conditions for continued operation; design and trial gates should then reflect the function's complexity and risk. The audit must examine buyer obligations alongside provider execution because neither firm controls the complete service alone.

The Study Shows Why Task Substitution Is Insufficient

The banking study records requests to reproduce existing workflows rather than reconsider them. On willingness to redesign processes for agentic AI, 429 of 804 valid respondent ratings, or 53.36 percent, were 3 or below on the 10-point scale. An individual contributor in technology engineering and operations stated, “Nobody re-engineered the process before handing it to the agent.” The evidence identifies a readiness problem to investigate; it does not establish that unchanged workflows always fail.¹

The difficulty extends beyond an obsolete process diagram. Respondents described outdated documentation, local variants, and providers reconstructing procedures by observing buyer staff. Some also reported problems after human checks were removed. The contracted sequence could differ from the work sustaining delivery, and removing a step could remove a necessary safeguard. The audit must establish the work and controls required for the agent-enabled service before deciding what to retain or replace.²

Deloitte’s August 2026 research identifies fragmented systems, poorly understood processes, and entrenched practices as obstacles to agentic adoption. Its sample comprises US leaders at organizations at least piloting agentic AI. Read alongside the banking accounts, it strengthens the case for examining the operating method without establishing a universal failure rate or a causal guarantee that redesign will succeed.³

Map the Function from Demand to Accepted Outcome

The map should begin with the result the buyer must accept. In the invoice example, that result is a valid invoice ready for authorized payment, with exceptions resolved or explicitly assigned. Classifying the invoice is an intermediate activity. Define the population served, completion rule, quality, timing, control obligations, justified exclusions, and downstream consequences. These conditions establish a stable boundary for both operational design and economic comparison. The UK Financial Conduct Authority similarly expects firms to map the people, processes, technology, facilities, information, and third-party dependencies supporting

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, “Process Documentation, Maturity, and Redesign” and “Automation of Processes as They Stand Rather Than Redesign,” <https://doi.org/10.2139/ssrn.7164559>.

² Hatch, *Agentic AI Outsourcing*, Part 2, “Process Documentation and Reported Reconstruction,” “Process Standardization and Repeatability,” and “Inconsistency of Processes Across Business Units and Regions”; Hatch, *Agentic AI Outsourcing*, Part 2, “Process Documentation, Maturity, and Redesign” and “Automation of Processes as They Stand Rather Than Redesign.”

³ China Widener et al., “AI Agents Are Only the Beginning: The Path to Agentic Transformation,” Deloitte Insights, August 12, 2026, “Most Processes Aren’t Ready for AI Agents” and “Methodology,” <https://www.deloitte.com/us/en/insights/industry/technology/path-to-agentic-transformation.html>.

important business services. For an in-scope financial firm, the map serves an existing resilience expectation as well as the proposed redesign.⁴

Trace representative items through both firms, including rejected, reopened, and failed work. Reconcile the contracted process with system events and staff walkthroughs, identifying variants across products, locations, and shifts. For a service spanning Switzerland, the United Kingdom, the European Union, and the United States, the map should attach each applicable local acceptance, record, or authorization requirement to the affected cases and responsible owner. This distinguishes a necessary local variation from an undocumented workaround. Include spreadsheets, portal retrieval, repeated runs, repairs, and informal decisions: the study's process-alteration accounts make omitted human work a priority for investigation.⁵

Each consequential handoff needs a shared record of the input's authoritative source, meaning, version, permissions, and defect owner. The record should also define acceptable completion, the person authorized to decide an exception, the response deadline, and the controlled action if that deadline is missed. Reconcile the buyer and provider versions. A connected interface is usable only when the receiving party understands the data and can act on it.

Measure execution separately from waiting for decisions, access, batch windows, review, and exception resolution. The study distinguishes a quick conversational response from the formal decision delivery requires. A response timestamp can therefore conceal the unresolved dependency. Recording the authorized decision and its operational effect makes the handoff testable.⁶

The completed map should support a baseline for accepted volume and complexity, quality, elapsed time, full cost, and recovery. Include buyer and provider effort, review, rework, technology, and supporting services. Separate transition expenditure from ongoing cost while retaining both in the investment decision. PwC's shared-services analysis similarly considers transition, technology and data, operations, governance, and capabilities alongside labor economics.⁷

⁴ Financial Conduct Authority, "Outsourcing and Operational Resilience," January 9, 2020, last updated July 14, 2026, "Operational Resilience and Third Party Providers," <https://www.fca.org.uk/firms/outsourcing-and-operational-resilience>.

⁵ Hatch, *Agentic AI Outsourcing*, Part 5, "Manual Alteration of Financial Institution Processes" and "Reasons for Manual Alteration of Financial Institution Processes."


⁶ Hatch, *Agentic AI Outsourcing*, Part 2, "Speed of Financial Institution Response to Issues Raised by the Engagement" and "Joint Decision Speed Mismatched to Agent Speed."

⁷ Namit Kapoor et al., "Agentic-First GBS: Moving Beyond Labor Arbitrage," PwC, May 2026, "How to Get Started: 5 Steps to Take Right Now," <https://www.pwc.com/us/en/services/consulting/business-transformation/intelligence-over-arbitrage-agentic-gbs.html>.

Design Execution and Authority Across Both Firms

The future design should connect service definition, execution, acceptance, and adaptation through shared evidence. Figure 1 places buyer and provider duties in the same flow and links the supporting data, models, tools, and providers to delivery and recovery. Defects and material changes return to accountable owners. This is a proposed operating design, requiring service-specific testing.

Figure 1. The delivery method reaches across both firms and their dependencies



Proposed design. Shared records and change acceptance connect the complete service.

Source: Author's proposed application of the study's process and joint-delivery evidence and the cited operating-model literature; not a measured study process or validated design.

Whole-method redesign requires examining each consequential activity, not changing every step. Retain work that serves the outcome or a necessary obligation, remove duplication where its purpose can be met more effectively, and select human judgment, deterministic automation, assistance, or agent execution accordingly. A Harvard-led experiment published in 2026 with BCG consultants found that AI assistance improved some tasks but worsened performance on another. The result concerns assisted knowledge work, not autonomous outsourcing; it provides a reason to test task allocation rather than assume that similar-looking work is equally suitable for agents.⁸

Control design requires the same discrimination. BCG's process-redesign guidance anticipates removing some checks aimed at human uncertainty, whereas the study records difficulties after human checks disappeared. These sources concern different designs and do not establish a general rule about fewer checks. Preserve each necessary control purpose and

⁸ Fabrizio Dell'Acqua et al., "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality," *Organization Science* 37, no. 2 (2026): 403–4, <https://doi.org/10.1287/orsc.2025.21838>.

prove that its replacement addresses the errors and consequences the proposed service can produce.⁹

Decision authority and qualified capacity must be designed together. McKinsey's September 2026 blueprint links redesigned workflows to explicit decision logic and corresponding changes in roles.¹⁰ In a two-party service, that requires specifying what agents may execute, what needs approval, and who may change instructions, restrict operation, or authorize restart. Each cross-firm exception needs an available owner and a response commitment matched to the potential consequence, with controlled action when no decision arrives. Budget review against ordinary and peak demand, including limits or rerouting when capacity is exceeded. The design is incomplete if its required decisions cannot be made within the operating window.

Exceptions and fallback belong within the normal method. Test inconsistent inputs, uncertain outputs, unavailable approvers, and failed models, tools, or interfaces. Recovery requires usable records, permissions, people, and capacity, together with the case ownership needed to return interrupted work to a controlled state. The fallback should preserve the same necessary acceptance and control purposes as the primary route; a nominal recovery mechanism that bypasses them cannot establish service readiness.

Make the Agreement Implement the Operating Design

The operating agreement should express the same service boundary as the delivery map. Mayer Brown's February 2026 analysis proposes adapting outsourcing provisions to agentic services, including defined tasks and delegated authority, outcome-based service levels, and usable audit rights. Its relevance is contractual: platform availability alone cannot specify the performance of an agent-delivered business function. The parties must negotiate obligations that match their control over the service.¹¹

Translate those obligations into operating practice. Align buyer response duties, provider responsibilities, human contribution, evidence access, pricing assumptions, and acceptance. Identify changes to models, prompts, tools, rules, or supporting providers that require testing

⁹ Eric Jesse et al., "Scaling AI Requires New Processes, Not Just New Tools," Boston Consulting Group, January 20, 2026, "Four Key Decisions for an End-to-End Process Transformation" and "How to Meet the Change Management Challenge," <https://www.bcg.com/publications/2026/scaling-ai-requires-new-processes-not-just-new-tools>; Hatch, *Agentic AI Outsourcing*, Part 2, "Process Documentation, Maturity, and Redesign" and "Automation of Processes as They Stand Rather Than Redesign."

¹⁰ McKinsey & Company, "Stacking the Odds: A Blueprint for Successfully Scaling Agentic AI," September 14, 2026, "Designing a Rigorous Blueprint," <https://www.mckinsey.com/capabilities/operations/our-insights/stacking-the-odds-a-blueprint-for-successfully-scaling-agentic-ai>.

¹¹ Rohith P. George et al., "Contracting for Agentic AI Solutions: Shifting the Model from SaaS to Services," Mayer Brown, February 17, 2026, secs. 1, 3, and 5, <https://www.mayerbrown.com/en/insights/publications/2026/02/contracting-for-agentic-ai-solutions-shifting-the-model-from-saas-to-services>.

and approval. Notice must reach the people running the affected service, and they must have a decision on materiality and an executable fallback. Contract language cannot compensate for an approval route that the operating teams cannot use.

Recalculate capacity and cost for the complete flow before accepting the design. Faster execution can increase review and decision demand; a lower agent task cost can coexist with higher rework or buyer effort. Compare accepted outcomes with the baseline on equivalent scope, complexity, quality, and demand. Fund retained review and recovery, and expose work transferred outside the provider's measured step. These are conditions of the redesigned service, not incidental costs to be excluded from its case.

Prove the Service Before Authorizing Further Scale

The audit should assign every material gap to a buyer or provider owner and distinguish prerequisites for operation from improvements suitable for later controlled releases. Missing authority for consequential decisions, unacceptable customer exposure, unusable evidence, or unworkable fallback should constrain expansion until resolved or adequately controlled. Figure 2 links the proposed map, design, trial, and expansion decisions to the evidence executives should require.

Figure 2. Evidence gates for the redesigned delivery method

Decision gate	Evidence required	Expansion condition
Map actual delivery	Actual flow, material variants, hidden human work, dependencies, and functional baseline.	Buyer and provider agree the evidence and expose unresolved gaps.
Accept the future design	Assigned authority, data/access readiness, review capacity, exceptions, fallback, and material-change route.	Critical prerequisites are met or adequately controlled.
Prove the bounded service	Accepted outcomes, full cost, time, traceability, and recovery under ordinary and difficult conditions.	Accountable buyer accepts evidence reviewed independently of delivery.
Expand and adapt	Monitored queues, manual effort, outcomes, changes, and tested rollback.	Each increase remains within proven operating conditions.

Proposed gates, not a universal schedule or validated method. The buyer owns the functional outcome; both firms own their assigned changes.

Source: Author's proposed audit, design, and acceptance sequence. Gate duration and operating tolerances must be set for the function; the sequence is not a validated intervention or universal deployment schedule.

Agree trial acceptance criteria before execution: the business outcome, case mix and exclusions, tolerable errors and delays, evidence completeness, recovery limits, and observation period. State which findings block expansion and who may accept a remaining limitation. Keep rejected cases, reruns, and manual queues within a consistent outcome denominator. If historical evidence is incomplete, reconstruct the baseline transparently and label uncertainty; a target is not evidence of earlier performance.

Run the complete service within defined permissions, volumes, customer populations, and rollback conditions. Compare accepted outcomes, full cost, elapsed time, traceability, exception recovery, and resilience with the baseline. Include peak arrivals, delayed decisions, inconsistent inputs, model or tool changes, and unavailable dependencies. Straightforward demonstration cases cannot establish performance under the conditions most likely to expose the method's limits.

The transition requires separate proof because old and new routes may coexist. Assign ownership of each item, detect duplicate or missing processing, and specify how interrupted work returns to a controlled state. Rehearse the handover with buyer decision makers, provider operators, and reviewers using their real permissions and escalation routes. A fallback must absorb the expected backlog, not merely be technically available. Include parallel-running effort in the investment and staffing plan.

The accountable buyer executive should accept the service on evidence reviewed independently of delivery and correction. Record unresolved limitations and the conditions under which approval remains valid. Expand in bounded increments while monitoring outcomes, queues, manual effort, and recovery. A new dependency or material change should trigger reassessment, and the feedback record should close with an accepted correction to the delivery method.

Require Proof of the Business Result Across Both Firms

Authorize the joint audit and redesign with an accountable outcome owner and dated evidence returns. Within the first month, require the delivery map, baseline, material gaps, and proposed operating conditions; approve subsequent gates according to risk and complexity. Further scale should depend on evidence that both firms can deliver the accepted business result at an acceptable cost, quality, control level, and recovery capability. That evidence requires examination of the complete method, including the buyer decisions and human work on which agent execution depends.

Evidence and Source Note

The study records 850 vendor delivery-team respondents on 150 Tier 1 financial-institution engagements during January–June 2026 in a nonrepresentative, single cross-section. Neither banks nor vendors were approached for corroboration. Its mechanisms inform questions for other sectors; its percentages do not establish their prevalence. The author-supplied version 1.8.2, dated November 1, 2026, was consulted as forthcoming material alongside public sources checked through October 23, 2026. The proposed map and gates require function-specific testing.¹²

Bibliography

- Dell’Acqua, Fabrizio, Edward McFowland III, Ethan Mollick, et al. “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality.” *Organization Science* 37, no. 2 (2026): 403–423. <https://doi.org/10.1287/orsc.2025.21838>.
- Financial Conduct Authority. “Outsourcing and Operational Resilience.” January 9, 2020. Last updated July 14, 2026. <https://www.fca.org.uk/firms/outsourcing-and-operational-resilience>.
- George, Rohith P., Joe Pennell, Brad L. Peterson, and Oliver Yaros. “Contracting for Agentic AI Solutions: Shifting the Model from SaaS to Services.” Mayer Brown. February 17, 2026. <https://www.mayerbrown.com/en/insights/publications/2026/02/contracting-for-agentic-ai-solutions-shifting-the-model-from-saas-to-services>.
- 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>.
- Jesse, Eric, Zeeshan Shah, and Rajeev Singh. “Scaling AI Requires New Processes, Not Just New Tools.” Boston Consulting Group. January 20, 2026. <https://www.bcg.com/publications/2026/scaling-ai-requires-new-processes-not-just-new-tools>.
- Kapoor, Namit, Mukund Kumar, and Fuad Abdelhadi. “Agentic-First GBS: Moving Beyond Labor Arbitrage.” PwC. May 2026. <https://www.pwc.com/us/en/services/consulting/business-transformation/intelligence-over-arbitrage-agentic-gbs.html>.
- McKinsey & Company. “Stacking the Odds: A Blueprint for Successfully Scaling Agentic AI.” September 14, 2026. <https://www.mckinsey.com/capabilities/operations/our-insights/stacking-the-odds-a-blueprint-for-successfully-scaling-agentic-ai>.
- Widener, China, Laura Shact, David Jarvis, and Sayantani Mazumder. “AI Agents Are Only the Beginning: The Path to Agentic Transformation.” Deloitte Insights. August 12, 2026. <https://www.deloitte.com/us/en/insights/industry/technology/path-to-agentic-transformation.html>.

¹² Hatch, *Agentic AI Outsourcing*, Introduction, “Evidentiary Basis and Its Limits.”