

The Board Audit for Agentic Outsourcing

Four Findings Require Immediate Examination

Phil Hatch | Akholi

October 23, 2026

Reported control departures and manual alteration each reached 150 of 150 sampled engagements. The manual-alteration contrast is 100% of engagements versus under 1.3% of pooled disclosure answers. Bank boards should commission an independent portfolio audit covering provider delivery and the bank's own readiness.

Approve the Independent Portfolio Audit Now


Bank boards should commission an immediate, independent audit of their agentic outsourcing portfolio. Every one of the study's 150 engagements had at least one respondent reporting that intended governance, oversight, and control functions were not implemented as designed. Manual alteration was also reported on all 150. Accounts of unchanged processes and inadequate organizational adaptation raise a broader concern: whether the bank and provider can jointly govern the service delivered. Provider assurance alone cannot resolve that question.¹

Figure 1 presents the manual-alteration contrast: 100 percent of engagements versus under 1.3 percent of pooled disclosure answers. The latter is 13 of 1,004 valid answers reporting that the respondent or management discussed specified alterations with the bank. This is an

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), Introduction, principal findings; Part 3, "Reported Governance Weaknesses Across All 150 Engagements," <https://doi.org/10.2139/ssrn.7164559>; Hatch, *Agentic AI Outsourcing*, Part 5, "Manual Alteration by the Delivery Team" and "The Five Manual Alteration Methods," fig. 5.1.

answer rate across three questions, with possible repeated people, not a percentage of unique respondents or engagements. It does not establish everything the bank knew.²

Figure 1. The manual-alteration contrast the board must investigate



Different units and scopes. Three disclosure questions cover data, process, and AI-output alteration; pooled answers are not unique people or engagements. Reporting alteration and third-party work are outside that disclosure measure.

Source: Hatch, 2026 study, version 1.8.2, Part 5, figures 5.1 and 5.12. Different units and scopes; reported discussion does not establish all the bank knew.

The mandate should therefore cover the bank’s own organization and processes alongside provider execution, across pilots, transitions, and production. The board should name an executive sponsor, fund independent assurance, authorize protected evidence access, and establish containment authority. A reconciled inventory within two weeks, supported priority assessments by week eight, and verified initial corrections with funded remaining work by week twelve would provide dated evidence for continuation decisions.³

Four Findings Challenge the Reported Delivery Model

The four findings connect weaknesses in control execution to the delivery arrangements those controls are intended to govern. Figure 2 identifies the corresponding audit demands. Only the control-departure and combined manual-alteration occurrence findings reach every sampled engagement. Process and organizational accounts identify reported problems without establishing universal prevalence; severity also differed among engagements.⁴

² Hatch, *Agentic AI Outsourcing*, Part 5, “Disclosure of Manual Alteration,” fig. 5.12 and table 5.8.

³ Hatch, *Agentic AI Outsourcing*, Part 6, table 6.2 and “Proposed Implementation Schedule.”

⁴ Hatch, *Agentic AI Outsourcing*, Introduction, principal findings; Part 3, “Reported Governance Weaknesses Across All 150 Engagements.”

Figure 2. Four findings, four requirements for audit evidence

Reported study evidence	Immediate audit demand
Governance departures reported on all 150 engagements	Reperform material controls; prove usable records, review capacity and intervention authority.
Manual alteration on 150 of 150; sparse bank-discussion answers	Trace all five methods, authorization and bank communication; reconcile performance and full cost.
Reported weaknesses in process fit	Map actual end-to-end delivery; test data, handoffs, review and exceptions.
Reported weaknesses in organizational readiness	Test bank decision rights, response capacity, competence and provider co-management.

Only the first two occurrence findings cover all 150 sampled engagements. Process and organizational findings are reported patterns, not universal rates. Disclosure has its separate pooled-answer base.

Source: Study Introduction and Parts 2, 3, and 5; the audit requirements are the author's proposed application. Reported findings are not independently established legal breaches.

Controls Were Not Implemented as Designed on Every Engagement

The control finding challenges reliance on completed approval records. A client-servicing contributor reported, “when queue is long the checker just click through,” attributing the behavior to high message volume and small teams on both sides. This is compliance theater as an interpretation of the account, not a measured category: approval remains while substantive examination is displaced. Auditors must test whether evidence, review capacity, authority, and intervention time allow the intended control to operate.⁵ FINMA's Swiss supervisory observations similarly examine AI ownership, testing, controls, and contractual responsibility in outsourcing, supporting examination of the operating allocation rather than the approval record alone.⁶

Manual Alteration Was Reported on All Sampled Engagements

Manual intervention can change that operating allocation without appearing in the bank's service account. The five overlapping methods cover altered bank data, processes, AI outputs, reporting, and third-party work performed by hand; together they reach all 150 engagements.

⁵ Hatch, *Agentic AI Outsourcing*, Part 2, “Human Review Capacity Against Agent Output Volume,” individual-contributor account, CXM, client servicing and communications.

⁶ Swiss Financial Market Supervisory Authority FINMA, *Governance and Risk Management When Using Artificial Intelligence*, FINMA Guidance 08/2024, December 18, 2024, sec. 2.1, 4, https://www.finma.ch/~media/finma/dokumente/dokumentencenter/myfinma/4dokumentation/finma-aufsichtsmittelungen/20241218-finma-aufsichtsmittelung-08-2024.pdf?sc_lang=en.

The disclosure questions cover only the first three methods. Approved correction can be legitimate, so the audit must distinguish authorized human work from uncontrolled intervention or inaccurate representation. It should establish what changed, who authorized it, what was recorded, and what the bank received.⁷

Processes Were Not Consistently Optimized for Agentic Delivery

The process accounts explain how such intervention can become necessary. A technology engineering contributor stated, “Nobody re-engineered the process before handing it to the agent.” Other accounts describe outdated documents, local variants, and delivery teams reconstructing practice. The audit should compare the approved method with the complete function delivered, from input to accepted outcome, including exceptions and manual repair. Redesign should address demonstrated mismatches while preserving necessary controls; the evidence does not establish that every retained workflow fails.⁸

Bank Organizations Had Not Fully Adapted to Agentic Outsourcing

Process changes also depend on the bank’s organization. The study includes cooperative staff constrained by slow decisions, limited authority, unsuitable access, inadequate review capacity, or gaps in skills and shared visibility. These conditions require bank correction as well as provider action.⁹ BCG’s authorization analysis argues for governing people, applications, and agents together, with traceable delegation and stop authority. Applied to outsourcing, the practical test is whether both firms can detect a deviation, obtain evidence, reach an empowered person, and contain exposure in time.¹⁰

Apply Operating Assurance Across the Delivery Arrangement

The resulting assurance scope extends beyond the model or direct provider. The FSB’s October 2026 final sound practices connect board oversight, organizational adaptation, human oversight, and third-party AI dependencies. They are practices institutions are

⁷ Hatch, *Agentic AI Outsourcing*, Part 5, “Manual Alteration by the Delivery Team” and “The Five Manual Alteration Methods,” fig. 5.1; Hatch, *Agentic AI Outsourcing*, Part 5, “Disclosure of Manual Alteration,” fig. 5.12 and table 5.8.

⁸ Hatch, *Agentic AI Outsourcing*, Part 2, “Process Documentation, Maturity, and Redesign,” “Process Documentation and Reported Reconstruction,” and “Automation of Processes as They Stand Rather Than Redesign.”

⁹ Hatch, *Agentic AI Outsourcing*, Part 2, “Organizational Readiness to Absorb Changes to Daily Work,” “Agentic AI Knowledge of Financial Institution Staff,” “Speed of Financial Institution Response to Issues Raised by the Engagement,” “Agent Actions Visible Only Inside Each Firm’s Own Tools,” and “Joint Decision Speed Mismatched to Agent Speed.”

¹⁰ Sandra Nudelman et al., “The Authorization Gap: Why Yesterday’s Controls Won’t Work with Today’s Agents,” Boston Consulting Group, September 22, 2026, “Creating a Third Framework Doesn’t Account for Cross-Cutting Behaviors” and “Can We Assign Accountability?,” <https://www.bcg.com/publications/2026/authorization-gap-ai-agent-governance>.

encouraged to consider, rather than binding rules replacing local obligations. They support a portfolio review of how the institution and its dependencies operate together.¹¹

At the execution level, the industry-developed Safeguards for Agentic Finance at Runtime approach links authorized actions, human escalation, and retained decision records. Its October revision expressly disclaims regulatory or supervisory status. As a reference approach hosted by the Monetary Authority of Singapore, it helps frame the audit question: can the identified control operate at the point and speed where the action creates exposure?¹²

Applying that question requires entity-specific legal analysis. Freshfields' UK/EU discussion identifies the difficulty of mapping existing governance and resilience requirements to particular AI uses, while distinguishing the jurisdictions' approaches. The audit charter should therefore identify the entity, service, role, applicable obligation, and supporting evidence. A group AI policy cannot establish that each regulated arrangement meets its requirements.¹³

Authorize Independent Examination of Both Firms

The charter should make independent examination possible. It should provide resources, a direct reporting route to the delegated board committee, and access arrangements for relevant bank and provider records and personnel. Management should own correction while independent assurance verifies it, consistent with the IIA's emphasis on organizational independence, direct board reporting, and safeguards against conflicting duties. Adviser conflicts and limits on staff confidentiality should be addressed before fieldwork.¹⁴

A complete inventory establishes what the bank must examine. Reconcile procurement, outsourcing, technology, model, access, and payment records, including material supporting providers. Detailed testing should then be prioritized by customer consequences, financial exposure, regulatory relevance, autonomy, volume, and known gaps. Whole-portfolio inclusion requires proportionate review and restrictions, not identical fieldwork or indiscriminate suspension.

¹¹ Financial Stability Board, "Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Final Report," October 15, 2026, sound practices 1–4, 10, and 12, <https://www.fsb.org/2026/10/sound-practices-for-responsible-adoption-of-artificial-intelligence-ai-final-report/>.

¹² *Safeguards for Agentic Finance at Runtime*, version 1.1, October 9, 2026, 4–5, 19, hosted by the Monetary Authority of Singapore, <https://www.mas.gov.sg/-/media/mas-media-library/development/fintech/ai-safr/safr-v1-1.pdf>.

¹³ Cyrus Pocha et al., "AI in Financial Services: Emerging Regulatory Expectations for Governance and Operational Resilience," Freshfields, July 6, 2026, opening discussion, <https://www.freshfields.com/en/our-thinking/blogs/risk-and-compliance/ai-in-financial-services-emerging-regulatory-expectations-for-governance-and-o-pe-102n7w7>.

¹⁴ The Institute of Internal Auditors, *Global Internal Audit Standards*, January 9, 2024, principle 7 and standard 7.1, 45–47, https://www.theiia.org/globalassets/site/standards/editable-versions/globalinternalauditstandards_2024january9_editable.pdf.

The review must also identify common dependencies. Several services may share a model, data feed, provider, or review team, creating exposure obscured by separate engagement reports. Establish the critical functions dependent on each resource, whether alternatives are usable, and who coordinates containment. The FSB's third-party practice addresses supply-chain complexity, concentration, and continuity; the portfolio audit should accordingly test shared fallback arrangements as well as individual services.¹⁵

For priority services, reconstruct outcomes across bank inputs, agent actions, human changes, approvals, acceptance, and reporting. Examining all five alteration methods should reconcile communication, dashboards, and return claims with source records, unresolved work, human effort, and complete functional cost. Independently repeat review and test exceptions, peak demand, fallback, and recovery. Missing evidence and conflicting accounts remain open findings until supported resolution.

Correction must follow the responsibility established by that evidence. Bank owners should resolve delayed decisions, source defects, access problems, incomplete processes, and capacity or skills gaps; joint rehearsal should test uncertain competence. Where outcomes may be defective, management should determine the affected population and repair, communication, or redress obligations. Legal and compliance specialists should assess notification duties promptly. Future control repair does not settle consequences of earlier delivery.

Connect Dated Evidence Returns to Continuation Decisions

The first twelve weeks should establish exposure, verify initial correction, and produce decisions on unresolved work. Week one should approve the mandate, evidence preservation, independent reporting, and initial containment; week two should return the reconciled inventory and exposure assessment. Priority review and protected evidence gathering should follow in weeks three and four. By week eight, supported assessments should identify impact, owners, containment, correction targets, and closure criteria. Critical matters require escalation when discovered.¹⁶

By week twelve, verified initial corrections should accompany a funded plan for remaining work. These proposed targets depend on portfolio size, evidence access, resources, and severity; they do not promise complete closure. Assurance should demonstrate the corrected delivery operating under relevant conditions. Revised policies, course attendance, completed action lists, or provider certification alone cannot establish that result.

Each material service should then have an explicit recommendation to continue, restrict, redesign, transfer, or exit. The decision should record permitted conditions, supporting

¹⁵ Financial Stability Board, "Sound Practices," practices 1–4, 10, and 12.

¹⁶ Hatch, *Agentic AI Outsourcing*, Part 6, table 6.2 and "Proposed Implementation Schedule."

evidence, unresolved exposure, the accountable approver, and the next review trigger. Necessary containment should remain while findings are open, and transfer or exit must preserve controlled continuity. Aggregate performance cannot resolve an unacceptable individual exposure.

Approve the Mandate and Require Operating Proof

The four findings warrant authorization now: reported control departures and manual alteration across all sampled engagements, sparse affirmative disclosure answers, and accounts of process and organizational mismatch. The independent portfolio audit should establish the bank's and provider's respective responsibilities and test the service against them. Continued or expanded delivery should depend on demonstrated control, transparent intervention, effective processes, and explicit acceptance of remaining exposure.

Evidence and Source Note

The study is a nonrepresentative, single cross-section of 850 vendor delivery-team respondents on 150 Tier 1 financial-institution engagements during January–June 2026. Banks and vendors were not approached for corroboration. Reports establish neither causation nor industry prevalence, uniform severity, legal breach, or verified bank ignorance. Author-supplied version 1.8.2, dated November 1, 2026, was consulted as forthcoming material, with external sources checked through October 23, 2026. External sources provide context rather than validation; the audit mandate and timetable are proposed applications.¹⁷

¹⁷ Hatch, *Agentic AI Outsourcing*, Introduction, “Evidentiary Basis and Its Limits”; Part 5, “Manual Alteration by the Delivery Team.”

Bibliography

- Financial Stability Board. “Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Final Report.” October 15, 2026. <https://www.fsb.org/2026/10/sound-practices-for-responsible-adoption-of-artificial-intelligence-ai-final-report/>.
- 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>.
- Institute of Internal Auditors. *Global Internal Audit Standards*. January 9, 2024. https://www.theiia.org/globalassets/site/standards/editable-versions/globalinternalauditstandards_2024january9_editable.pdf.
- Nudelman, Sandra, Anne Kleppe, Jeanne Bickford, Biljana Bajic-Bizumic, Kirill Katsov, and Eitan Yehuda. “The Authorization Gap: Why Yesterday’s Controls Won’t Work with Today’s Agents.” Boston Consulting Group. September 22, 2026. <https://www.bcg.com/publications/2026/authorization-gap-ai-agent-governance>.
- Pocha, Cyrus, Claire Harrop, Lauren Moorin, and Sophia Doan. “AI in Financial Services: Emerging Regulatory Expectations for Governance and Operational Resilience.” Freshfields. July 6, 2026. <https://www.freshfields.com/en/our-thinking/blogs/risk-and-compliance/ai-in-financial-services-emerging-regulatory-expectations-for-governance-and-o-pe-102n7w7>.
- Safeguards for Agentic Finance at Runtime*. Version 1.1. October 9, 2026. Hosted by the Monetary Authority of Singapore. <https://www.mas.gov.sg/-/media/mas-media-library/development/fintech/ai-safr/safr-v1-1.pdf>.
- Swiss Financial Market Supervisory Authority FINMA. *Governance and Risk Management When Using Artificial Intelligence*. FINMA Guidance 08/2024. December 18, 2024. https://www.finma.ch/~/-/media/finma/dokumente/dokumentencenter/myfinma/4dokumentation/finma-aufsichtsmittelungen/20241218-finma-aufsichtsmittelung-08-2024.pdf?sc_lang=en.