

The First Twelve Weeks of Agentic Outsourcing Review

A Board Programme for Evidence and Action

Phil Hatch | Akholi

October 23, 2026

The board should require a twelve-week programme that identifies material agentic engagements, tests consequential control assumptions, and converts findings into restrictions, remediation, or a justified continuation decision. The timetable is a proposed management discipline, not a regulatory deadline.

Commission a Review That Can Change Delivery

The board should ask management to establish where provider agents can affect money, customers, regulated work, sensitive information, or critical operations. The first objective is an actionable inventory of material exposure. A comprehensive technical census can follow, but it should not delay scrutiny of the most consequential services.

Deloitte's research describes agent deployment advancing faster than governance maturity in many surveyed organizations.¹ It reports business and technology leaders' perspectives, not an audit of this organization. The board should use it as a reason to test its own evidence rather than adopt an assumed failure rate.

Give the review access to contracts, operating teams, configurations, records, and relevant provider personnel. Name an executive sponsor with authority to resolve cross-functional issues. Without access and decision power, a twelve-week exercise can produce a report while leaving the service unchanged.

¹ Andy Bayiates, "Business and IT Leaders Report AI Agents Are Scaling Faster Than Their Guardrails," Deloitte Insights, April 24, 2026, survey scope and governance discussion, <https://www.deloitte.com/us/en/insights/topics/emerging-technologies/ai-agents-scaling-faster.html>.

Weeks One to Three Establish the Exposure

Identify material engagements, responsible buyer owners, provider delivery leaders, approved agent scope, downstream dependencies, and current restrictions. Reconcile the inventory to procurement, technology, data, and operations records. Mark uncertainty rather than fill gaps with a general supplier assurance.

The study reports governance-function departures and manual alterations across its sampled engagements.² Those accounts do not establish the condition of the board's own portfolio. They identify two priority tests: whether intended controls operate and whether human repair is visible.

Select services by consequence and uncertainty. Include an apparently successful engagement if management cannot explain how its outcomes are verified. Obtain immediate action on a serious known exposure; the programme should not postpone a necessary restriction until the final board meeting.

Weeks Four to Eight Test the Operating Reality

Figure 1 proposes a sequence from scope to evidence, correction, and board decision. For priority services, reconstruct consequential cases, observe approval and intervention, test a controlled interruption, and reconcile claimed performance with accepted outcomes and human work.

Distinguish documentary adequacy from operating effectiveness. A contract clause, process diagram, and completed review can coexist with an unusable escalation route. Record the actual finding, evidence limitation, consequence, responsible owner, and proposed action rather than compress all findings into one maturity score.

EY's governance research emphasizes recurring evidence of control effectiveness.³ The proposed review applies that principle to a bounded portfolio exercise. Use specialists appropriate to the questions, and preserve independence where delivery sponsors would otherwise assess their own important claims.

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 5, “The Five Manual Alteration Methods” and “Disclosure of Manual Alteration,” table 5.8.

³ Richard Jackson and John McLain, “AI Governance Has Entered Its Next Phase: Closing the Confidence Gap,” EY, September 15, 2026, US survey population and chapters 4–5, https://www.ey.com/en_us/insights/assurance/ai-governance-has-entered-its-next-phase-closing-the-confidence-gap.

Figure 1. The proposed twelve week review programme



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

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

Weeks Nine to Twelve Resolve and Decide

Management should classify findings into issues that require immediate restriction, funded remediation, further investigation, or accepted residual exposure. Define acceptance criteria for closure and retest significant corrections. A new policy or training record alone should not close an issue caused by unavailable authority or incompatible tools.

EU, UK, Swiss, US, and other obligations should be assessed by entity, activity, and jurisdiction. Counsel and relevant control functions should determine whether findings require a legal, regulatory, contractual, customer, or incident response. The review itself should not label every operational departure a breach.

Bring the board a concise portfolio view: material services, evidence-supported conclusions, unresolved limitations, actions taken, and decisions requiring board attention. Include the cost and operating implications of remediation so that control requirements and the delivery business case remain connected.

Turn the Programme into Continuing Oversight

A reasonable objection is that twelve weeks is too short to audit a large portfolio. The programme should establish priorities, test the most material uncertainties, and create an accountable continuation plan. It should not claim comprehensive assurance where the work completed cannot support it.

Require periodic evidence on material changes, interventions, exceptions, accepted outcomes, and overdue remediation after the programme ends. Set triggers for renewed review when the agent's authority, model, process, data, or provider chain changes. Keep the portfolio record current enough to guide decisions.

The board's desired result is a defensible account of what can continue, what must change, and what remains uncertain. A twelve-week timetable is useful only if it produces those decisions. The programme should end with operating action and sustained ownership, not merely a completed governance exercise.

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 proposed twelve-week board review programme linking portfolio uncertainty to operating tests, corrective decisions, and continuing oversight. 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

Bayiates, Andy. "Business and IT Leaders Report AI Agents Are Scaling Faster Than Their Guardrails." Deloitte Insights. April 24, 2026. <https://www.deloitte.com/us/en/insights/topics/emerging-technologies/ai-agents-scaling-faster.html>.

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>.

Jackson, Richard, and John McLain. "AI Governance Has Entered Its Next Phase: Closing the Confidence Gap." EY. September 15, 2026. https://www.ey.com/en_us/insights/assurance/ai-governance-has-entered-its-next-phase-closing-the-confidence-gap.

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