

After the Agent Acts Recover the Service and Repair the Harm

A Joint Method for Consequential Incidents

Phil Hatch | Akholi

October 23, 2026

Delivery managers should plan for actions that cannot be undone by stopping the agent. The incident method must separate containment, correction of operational records, repair of external consequences, and proof that the service is safe to resume.

Start with the Consequence Already Created

The buyer incident lead should establish what happened outside the agent before declaring containment successful. An agent can be stopped while an incorrect payment, customer message, record change, or submitted application remains active. Disabling the execution route prevents further action; it does not prove that earlier effects were reversed.

Mayer Brown's contract analysis connects agentic service scope with governance, performance obligations, and remedies.¹ Managers should turn those provisions into an executable response method. The agreement must identify who can preserve evidence, restrict the service, contact affected parties, authorize corrections, and obtain cooperation from downstream providers.

The response should begin from verified external effects. Record what was attempted, what the receiving system accepted, what remains pending, and which consequences are uncertain. An agent's own summary may assist investigation but should not be the sole evidence used to decide that no harm occurred.

¹ Rohith P. George et al., "Contracting for Agentic AI Solutions: Shifting the Model from SaaS to Services," Mayer Brown, February 17, 2026, secs. 1–6, <https://www.mayerbrown.com/en/insights/publications/2026/02/contracting-for-agentic-ai-solutions-shifting-the-model-from-saas-to-services>.

Separate Containment from Repair

Allied agentic-security guidance emphasizes resilience and human intervention.² Containment limits additional exposure. It may include restricting tools, suspending a workflow, revoking delegated authority, or routing work to an approved fallback. The appropriate action depends on the service; an abrupt shutdown can create a different risk where unfinished work must remain controlled.

Repair addresses effects already created. Some records can be corrected through an authorized reversal. Other actions require a new transaction, customer notification, specialist review, or a formal remedy. The incident team should not describe a compensating action as if it erased the original event. Preserve both in the record.

The study's reports of stopping difficulties and fragmented visibility support examining these distinctions across the firm boundary. They do not establish the loss or legal remedy in any particular incident.³ The practical implication is to assign operational decisions before a dispute about responsibility consumes the available response time.

Use One Consequence and Recovery Record

Figure 1 proposes a joint record linking each affected case to its observed consequence, owner, repair authority, current state, and completion evidence. Include uncertain cases until their disposition is resolved. A total count of restarted jobs cannot show whether the original customers, balances, or submissions were corrected.


For a hypothetical duplicate order, the team should establish whether both requests were accepted, whether fulfillment started, and whether cancellation remains possible. The provider may be able to stop its own queue while only the buyer can contact the recipient or approve a financial adjustment. Those dependencies belong in the response record.

The record should distinguish evidence preservation from restoration. Changing a production record can remove information needed to understand the original action. Agree how to preserve relevant versions, logs, permissions, and communications under applicable access and retention rules before performing a correction that would obscure the sequence.

² Australian Signals Directorate's Australian Cyber Security Centre et al., *Careful Adoption of Agentic AI Services* (Australian Signals Directorate, May 1, 2026), "Manage third-party components," "Identity management," "Resilience," and "Human in the loop," <https://www.cyber.gov.au/publication/careful-adoption-of-agentic-ai-services>.

³ 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 3, Co-Management, sec. 2.4, "Speed of Stopping a Misbehaving Agent in Live Operations," indicator C143, <https://doi.org/10.2139/ssrn.7164559>; Hatch, *Agentic AI Outsourcing*, Part 2, sec. 5.1, "Agent Actions Visible Only Inside Each Firm's Own Tools."

Figure 1. From containment to verified repair



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

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

Make the Buyer and Provider Decisions Explicit

Assign an incident commander with authority appropriate to the event, supported by business, security, legal, privacy, and supplier specialists. State which decisions the provider can make immediately, which require buyer approval, and which cannot wait for ordinary governance meetings. Keep a timed decision log and unresolved questions visible.

A legitimate concern is that a common record could expose sensitive information across firms. Use controlled access and the minimum necessary case detail. EU, UK, Swiss, and US notification, secrecy, privacy, and sector duties need local assessment. Contractual incident notice should support legally required decisions without being mistaken for the legal deadline itself.

Where responsibility or recovery costs are disputed, preserve the claim separately from urgent mitigation. The commercial team can record reservations while operations performs authorized protective work. A liability clause should not become an unplanned prerequisite for contacting affected customers or preventing another consequential action.

Prove Repair Before Returning to Normal

Define resumption conditions for the failed mechanism. Correct the cause, test affected boundaries, reconcile the backlog, and confirm who accepts the remaining uncertainty. Replaying all failed jobs without knowing which actions reached the receiving system can reproduce the original harm.

Review the incident after recovery. Identify whether delayed decisions, incomplete records, unavailable authority, or missing subcontractor cooperation prevented effective repair. Update the operating instructions and rehearse the affected path. The lesson should be a changed ability to respond, not merely a revised incident description.

The manager should close the event only when the service state and the consequential case states are reconciled, required follow-up has owners, and the resumption decision is supported by evidence. A stopped agent is the beginning of control. Completed repair is a separate outcome that the two firms must prove.

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 consequence-and-recovery record separating agent containment, externally verified effects, authorized repair, and evidence-based resumption. 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

Australian Signals Directorate's Australian Cyber Security Centre et al.. *Careful Adoption of Agentic AI Services*. Australian Signals Directorate, May 1, 2026. <https://www.cyber.gov.au/publication/careful-adoption-of-agentic-ai-services>.

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>.

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