

A Critical Infrastructure AI Exit Plan Must Preserve a Safe Service

A Board Decision on Operational Separation

Phil Hatch | Akholi

October 23, 2026

Critical-infrastructure boards should approve an AI supplier only with a feasible route to continue or restrict the service safely at exit. Termination rights and data exports do not establish a safe operating transition.

Define the Safe Service Before the Exit Clause

The board should ask management what happens to the physical or essential service when an AI supplier is no longer available or trusted. The required outcome may be continued operation, a restricted operating mode, or a controlled transition to a safe state. The operator's engineering and safety functions should define that outcome before the commercial team treats exit as complete.

Allied OT guidance calls for safety-aware failure procedures and mechanisms to bypass or replace AI within existing operating controls.¹ It is voluntary security guidance, not a safety certification. The operator must demonstrate a suitable local method for its actual assets and process conditions.

¹ Cybersecurity and Infrastructure Security Agency and Australian Signals Directorate et al., "Principles for the Secure Integration of Artificial Intelligence in Operational Technology," Cybersecurity and Infrastructure Security Agency and Australian Signals Directorate et al., December 4, 2025, secs. 2.4 and 4.2, "Embed Safety and Failsafe Mechanisms," <https://www.cyber.gov.au/publication/principles-for-the-secure-integration-of-artificial-intelligence-in-operational-technology>.

Distinguish advice, planning support, maintenance administration, and consequential control. Exiting a document assistant is different from removing a component on which an operational decision depends. The board should approve the relevant service boundary and retained operator authority rather than a single exit policy for every AI tool.

Identify What the Operator Needs Without the Supplier

The study describes operational knowledge in personal notes and experienced staff, alongside access and coordination difficulties.² These banking accounts do not establish infrastructure safety outcomes. They provide a reason to test whether the operator's own transition plan depends on knowledge or access it does not control.

List the records, configurations, interfaces, permissions, licenses, and skills required for the chosen exit mode. Include outstanding work and supplier-created context. A data export may omit the interpretation, process state, or external right needed to use the information. Management should identify those gaps before the relationship becomes distressed.

A hypothetical maintenance-planning service may export asset recommendations while the supplier retains the rule history and unresolved exceptions. A successor could receive the files without knowing which recommendations were approved or already acted on. The transition record should preserve disposition and authority as well as content.

Commission a Safe Exit Demonstration

Figure 1 proposes four acceptance conditions: defined safe mode, usable transition assets, available operating capacity, and reconciled outstanding work. Test a representative portion of the service with the supplier unavailable to answer routine questions. Record which conclusions depend on cooperation that may disappear during a disputed exit.

Engineering should determine how connections and permissions are withdrawn without creating an unsafe process state. A commercial instruction to switch off the agent is not an engineering procedure. The test should identify sequence, prerequisites, permitted manual intervention, and the person authorized to decide when the transition is safe.

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 4, "Knowledge at Risk of Departure," <https://doi.org/10.2139/ssrn.7164559>; Hatch, *Agentic AI Outsourcing*, Part 3, Financial Institution Readiness, sec. 4.1, "Ability to Provision Access and Identities for Agents," indicator C108.

DOE's Stormbreaker work evaluates configuration-specific AI behavior in critical infrastructure; it does not certify a deployment.³ The analogous procurement lesson is to test the actual exit configuration. Evidence from normal operation or a different site cannot establish that the proposed replacement mode works here.

Figure 1. The safe operational exit decision



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

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

Secure Cooperation and Rights While the Deal Is Negotiable

The agreement should identify transition deliverables, formats, update frequency, access, training, supplier assistance, and the treatment of downstream dependencies. Record which rights survive termination and which require a separate license or authorization. A right to receive a file is weaker than a right and ability to operate the required service.

EU, UK, Swiss, US, and other infrastructure operators face different safety, cyber, data, licensing, and continuity duties. Local engineering and legal specialists should assess the proposed transition. A common group exit method can organize the work while site-specific conditions determine the safe implementation.

Management should distinguish planned exit from sudden supplier failure. Planned assistance may be unavailable in an insolvency, cyber incident, or contractual dispute. The board should see which transition assets and retained capabilities are available independently, and which assumptions remain contingent on supplier cooperation.

³ US Department of Energy, Office of Cybersecurity, Energy Security, and Emergency Response, *CESER Releases New Testbed to Advance LLM and Agentic AI Evaluation for Critical Infrastructure* (US Department of Energy, Office of Cybersecurity, Energy Security, and Emergency Response, July 16, 2026), configuration transfer and limit discovery, <https://www.energy.gov/ceser/articles/ceser-releases-new-testbed-advance-llm-and-agentic-ai-evaluation-critical>.

Keep Exit Capacity Alive During the Contract

A reasonable concern is that maintaining an alternative route reduces the AI business case. The operator should choose a proportionate safe mode and include its cost in the investment decision. It should not remove the capability and continue to describe the original exit plan as feasible.

Rehearse material changes in supplier architecture, site configuration, staff, and operating conditions. Update the transition assets and test whether a successor or retained team can use them. A completed annual document review should not substitute for evidence that the required service can be sustained or safely restricted.

The board should approve the supplier relationship with a current, demonstrated exit capability and explicit residual exposure. The test of exit is the state in which the operator and affected service are left. A signed termination notice closes a commercial relationship; safe operational separation requires its own evidence.

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 safe operational exit decision record linking engineering-defined service state, usable transition assets, independent capacity, and reconciliation of outstanding work. 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

Cybersecurity and Infrastructure Security Agency and Australian Signals Directorate et al.. “Principles for the Secure Integration of Artificial Intelligence in Operational Technology.” Cybersecurity and Infrastructure Security Agency and Australian Signals Directorate et al.. December 4, 2025.
<https://www.cyber.gov.au/publication/principles-for-the-secure-integration-of-artificial-intelligence-in-operational-technology>.

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>.

US Department of Energy, Office of Cybersecurity, Energy Security, and Emergency Response. *CESER Releases New Testbed to Advance LLM and Agentic AI Evaluation for Critical Infrastructure*. US Department of Energy, Office of Cybersecurity, Energy Security, and Emergency Response, July 16, 2026.
<https://www.energy.gov/ceser/articles/ceser-releases-new-testbed-advance-llm-and-agentic-ai-evaluation-critical>.

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