

Bounded Authority and Recovery in Infrastructure Agent Services

An Assurance Offer for Critical Infrastructure Providers

Phil Hatch | Akholi

October 23, 2026

Infrastructure providers should sell a defined support service with explicit authority limits and demonstrated recovery. The offer must explain its dependencies and evidence without implying that cybersecurity guidance certifies autonomous critical control.

Sell a Service the Buyer Can Approve

The provider's chief executive and assurance lead should define the authority and recovery obligations in an infrastructure agent offer before marketing autonomy. The buyer must preserve essential service and may be unable to accept an open-ended action space. A credible offer explains the support task, prohibited actions, dependencies, and conditions under which it can operate.

The allied May 2026 guidance recommends low-risk, non-sensitive uses and strong limits on access. Providers should reflect that scope in the offer rather than use general agent-security language to imply authorization for critical control. Any proposal to expand toward consequential operations requires separate engineering, safety, legal, and customer approval.¹

A bounded offer can still create commercial value. Documentation assistance, approved-information retrieval, and other carefully selected support activities may reduce effort without

¹ 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>.

granting control authority. The provider should demonstrate the specific outcome and retained human duties rather than claim a general maturity level.

Document the Operating Assumptions

The service specification should identify the environment, information sources, interfaces, tools, identities, and human decisions on which performance depends. It should also state what happens when a source is missing, stale, contradictory, or inaccessible. An agent should expose an unmet prerequisite instead of inventing a completion.

The study's reported access and observability problems show why these assumptions deserve commercial attention. A service may appear to work because people repair blocked access or incomplete execution. Banking evidence does not establish infrastructure failure rates, but it does support asking providers to identify necessary human work and evidence of actual completion.²

The provider should distinguish tested conditions from intended capability. A narrow demonstration with clean inputs cannot support an unrestricted production claim. State the workload, configuration, exceptions, and constraints represented by the evidence so the buyer can judge its relevance.

Build a Service Assurance Dossier

Figure 1 proposes four parts of a service assurance dossier: authority, dependencies, test evidence, and recovery. The dossier should be maintained for the configuration offered to the buyer. It should explain the provider's responsibilities and the operator prerequisites, with access to supporting evidence under appropriate protection.

The provider should include negative tests demonstrating prohibited actions are blocked. Positive task completion alone cannot prove the authority boundary. Testers should examine both the ordinary agent permissions and the administrative routes that could add tools, delegate work, or change the configuration.

Use configuration-specific evaluation. DOE's Stormbreaker approach emphasizes varying instructions, tools, environments, and other components to understand limitations. That

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 3, Financial Institution Readiness, sec. 4.1, "Ability to Provision Access and Identities for Agents," indicator C108, <https://doi.org/10.2139/ssrn.7164559>; Hatch, *Agentic AI Outsourcing*, Part 3, AI Technology, sec. 4.1, "Accuracy of the Tools' Own Success and Failure Reports," indicator C172.

supports a provider test plan that records the configuration and operating conditions. It does not turn a successful evaluation into safety certification for the customer's site.³

Figure 1. The bounded agent service assurance dossier

Dossier part	Provider commitment	Buyer acceptance evidence
Authority	Defined task and prohibited actions	Enforced limits and negative tests
Dependencies	Visible prerequisites and human work	Site-specific dependency review
Evaluation	Configuration and conditions disclosed	Representative completion evidence
Recovery	Priced cooperation and readiness	Rehearsed safe state and restoration

Author's proposed operating framework. Adapt to the service and applicable jurisdiction; not an empirically validated standard.

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

Promise Recovery That Operations Can Demonstrate

The recovery commitment should state who can restrict the service, what evidence is preserved, how outstanding work is reconciled, and what support remains available. The operator's engineering team should determine the safe state. The provider should demonstrate cooperation with that design rather than assume that immediate isolation is always the correct response.

Recovery capacity needs pricing. Skilled staff, current procedures, replacement access, and test environments cost money even when no incident occurs. The proposal should identify the included readiness and response duties, the limits of surge capacity, and any buyer prerequisites. Avoid an unqualified recovery promise supported only by an on-call email address.

The commercial terms should preserve incident cooperation while cause and liability are investigated. Mayer Brown's contract analysis connects performance dependencies, governance, liability, and exit. Applied to this service, the provider should make the

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>.

cooperation and evidence needed for recovery explicit rather than leave them to later negotiation.⁴

Govern Expansion as a New Offer

Adding a tool, changing an operational interface, or enabling a consequential action can change the nature of the service. The provider should assess the change, retest the relevant boundaries, and obtain the buyer's approval before expanding. A general improvement clause should not become permission for new physical-process authority.

The dossier should also support refusal of an unsuitable request. If a buyer asks the provider to connect a support agent to a control interface without the required assessment, the provider should document the additional authority and evidence needed. Offer an alternative that preserves the approved boundary. This gives commercial teams a concrete basis for negotiating scope instead of leaving engineers to resist an autonomy promise already made in the sales process.

Global providers should maintain a common technical description and jurisdiction-specific conclusions. EU, UK, Swiss, US, and other operators face different sector and safety requirements. The offer should state which local assessments the provider supports and which remain the operator's responsibility, without claiming that one national guideline authorizes every deployment.

The assurance lead should control sales claims against the dossier. Every claim of completion, resilience, or autonomy should have a stated scope and evidence. If the evidence is limited to a support environment, the proposal and demonstration should retain that limit.

Win Approval with Evidence and Clear Limits

Before signature, commercial, engineering, security, and assurance leaders should approve the same dossier and service specification. Identify untested assumptions, required operator decisions, priced readiness, and excluded authority. Resolve inconsistencies between the demonstration, statement of work, and operating procedures.

A buyer can make a defensible decision when the provider explains exactly what it will do and how the parties recover when it fails. The provider should compete on a demonstrated support outcome and credible cooperation, with any expansion dependent on new evidence and approval.

⁴ Rohith P. George et al., "Key Contract Issues in Agentic AI Implementation and Integration Deals," Mayer Brown, June 16, 2026, "Liability," "Performance commitments and dependencies," and "Exit and transition," <https://www.mayerbrown.com/en/insights/publications/2026/06/key-contract-issues-in-agentic-ai-implementation-and-integration-deals>.

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, 2026.⁵

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 provider service assurance dossier tying a bounded infrastructure offer to configuration-specific evidence and priced recovery. 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., Veronica R. Glick, Arsen Kourinian, et al.. "Key Contract Issues in Agentic AI Implementation and Integration Deals." Mayer Brown. June 16, 2026. <https://www.mayerbrown.com/en/insights/publications/2026/06/key-contract-issues-in-agentic-ai-implementation-and-integration-deals>.

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>.

⁵ Hatch, *Agentic AI Outsourcing*, Introduction, "Executive Summary" and "Evidentiary Basis and Its Limits."