

Grid AI Outsourcing and Delegated Authority

A Utility Board Decision on Bounded Delivery

Phil Hatch | Akholi

October 23, 2026

Utility boards should decide what authority remains with the operator before buying an agentic service. Delegation, behavioral freedom, operating conditions, and intervention are separate decisions that require evidence appropriate to essential service.

Retain the Operator's Authority Explicitly

The utility board should require an authority statement before approving outsourced agentic delivery. The statement should distinguish advice, preparation, and consequential action, identify the operator decisions retained, and define the conditions under which the provider may act. A broad promise of autonomy does not establish a safe or lawful operating scope.

The allied May 2026 guidance recommends low-risk, non-sensitive initial uses and warns against broad access to critical systems. The board should treat that caution seriously. The proposed approach favors bounded support; any expansion toward consequential grid action requires separate engineering, safety, legal, and operator approval.¹

The sourcing decision should begin with the business need and alternatives. Process simplification, deterministic automation, or qualified human support may meet the need with different risks. Management should explain why agentic behavior is useful for the proposed task and which uncertainty it introduces.

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>.

Separate Delegation from Behavioral Freedom

Delegation defines what the provider is authorized to accomplish on the operator's behalf. Behavioral freedom defines how much choice the agent has in pursuing it. A service can have a narrow goal and still reach dangerous tools, or broad information access while remaining unable to act. The board needs both dimensions described.

A hypothetical outage-report assistant could summarize approved records without authority to alter dispatch instructions or equipment settings. Its output still requires appropriate validation before operational reliance. Management should specify which conclusions the assistant may produce, who checks them, and how uncertainty or missing information reaches an operator.

The study describes identity and stopping difficulties in bank engagements. These mechanisms support scrutiny of cross-firm authority and intervention, but do not establish grid safety performance or incident rates. Utility leaders should use them as questions for their own service assessment.²

Approve an Operating Envelope

Figure 1 proposes an authority envelope covering purpose, permitted actions, operating conditions, and intervention. It should identify prohibited actions and the system controls enforcing them. The provider should not be able to expand the envelope through a routine tool or configuration update.

Management should define the conditions that invalidate normal operation: unavailable or stale information, unexpected system state, uncertain identity, failed monitoring, or other task-specific limits. The service should enter the approved restricted state and escalate. It should not improvise an operational workaround merely to satisfy its assigned goal.

The operator's engineering and safety specialists should determine the restricted state. In a physical system, abrupt withdrawal of access or connectivity can have consequences. The board should require evidence of safe intervention without prescribing a universal technical stop method.

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, Co-Management, sec. 2.4, "Speed of Stopping a Misbehaving Agent," indicator C143.

Figure 1. The utility board authority envelope

Decision area	Retained operator decision	Required service evidence
Purpose and action	Approved support and prohibited authority	Enforced task and tool limits
Operating conditions	When normal use becomes unsuitable	Restricted-state and escalation tests
Intervention	Who can act and how safely	Engineering-approved joint rehearsal
Expansion	Additional authority and consequences	New configuration and local approvals

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.

Require Configuration Specific Evidence

The proposed service should be tested with the actual model, instructions, tools, permissions, interfaces, and representative operating conditions. A benchmark on another configuration cannot establish the behavior of this one. Document both successful completion and the enforcement of prohibited actions.

DOE's Stormbreaker testbed emphasizes configuration transfer and limit discovery in power-system and OT evaluation. It demonstrates why a static model score alone is insufficient. The board should require management to apply that evaluation discipline locally, while recognizing that a testbed is not a safety approval for autonomous grid control.³

The provider should disclose human correction and dependencies in the demonstrated result. If people intervene to produce acceptable support, that contribution belongs in the operating and cost model. Management should not accept an autonomy claim that omits the staff needed to maintain the service.

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>.

Preserve Local Duties and Recovery Capability

EU, UK, Swiss, US, and other grid operators face different sector, cybersecurity, data, and safety duties. Management should obtain an applicability assessment for the specific entity and task. A common supplier contract can support a shared authority vocabulary, but cannot replace local operational approval.

The board should distinguish a useful advisory result from permission to act on it automatically. A support agent can perform well at summarization while remaining unsuitable for dispatch or equipment control. Expansion therefore requires new evidence about the consequential action, its operating conditions, and the ability to intervene. The prior support pilot should be treated as evidence within its scope, with the additional authority assessed as a separate decision.

The continuity plan should identify the accepted service level during restriction, the retained staff and tools, and the evidence needed for restart. Reconcile incomplete or uncertain work before resumption. A restored application does not establish that its recommendations or work queue are safe to rely on.

Mayer Brown's agentic governance analysis supports explicit action boundaries and human checkpoints. The utility application proposed here requires engineering-led interpretation of those principles and tested cooperation between operator and provider. General governance commentary does not certify the physical process.⁴

Approve Support Within Demonstrated Limits

The board should approve a defined support scope, funded oversight, stop conditions, and a review date. Management should report the authority granted, test limitations, material dependencies, and unresolved findings. An expansion request should explain the additional consequences and evidence rather than rely on the earlier pilot's approval.

The decision is whether the utility can retain control of essential service while obtaining a demonstrated support benefit. Bounded authority, configuration-specific testing, and a viable recovery design make that decision reviewable. Where those conditions remain unproven, keep consequential grid authority with the approved operator processes.

⁴ Arsen Kourinian, "Governance of Agentic Artificial Intelligence Systems," Mayer Brown, February 24, 2026, "Mitigation Measures," <https://www.mayerbrown.com/en/insights/publications/2026/02/governance-of-agentic-artificial-intelligence-systems>.

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 utility board authority envelope separating delegation, behavioral freedom, operating conditions, and safe intervention. 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>.

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>.

Kourinian, Arsen. "Governance of Agentic Artificial Intelligence Systems." Mayer Brown. February 24, 2026. <https://www.mayerbrown.com/en/insights/publications/2026/02/governance-of-agentic-artificial-intelligence-systems>.

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."