

Authorizing Outsourced Infrastructure Agent Identities

A Joint Access Method for Utility Managers

Phil Hatch | Akholi

October 23, 2026

Infrastructure managers should authorize provider agents through task-specific identities and tested revocation. Initial use should remain limited to approved low-risk support, with operational engineering defining how access can be withdrawn safely.

Authorize a Defined Support Task

The utility's vendor manager and identity lead should require a separate access decision for every provider agent. An infrastructure support contract does not authorize an agent to inherit a technician's broad permissions. The decision should name the support task, data, tools, environment, operating conditions, duration, and responsible firms.

The allied May 2026 guidance recommends low-risk, non-sensitive initial tasks and warns against broad or unrestricted access, especially to critical systems. That caution is directly relevant to sourcing. Security guidance should not be presented as permission for autonomous control of a physical process.¹

Operational technology engineers should define the boundary between the support service and consequential operations. A read-only analysis tool can still create risk if its output drives an operator decision without adequate validation. The access review should consider both what the agent can change and how its information is used.

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>.

Give the Agent an Attributable Identity

The identity record should distinguish the provider organization, agent principal, sponsoring buyer owner, and administrator able to change the agent. Record the approved grant, credential lifetime, allowed delegation, and renewal owner. A named employee's account should not become an informal substitute because the agent onboarding route is slow.

The study's banking respondents described shared accounts, staff credentials, and delayed provisioning. These are mechanisms worth testing in infrastructure support, not evidence of their prevalence in utilities. The operational lesson is to provide an authorized identity route and a visible fallback when provisioning cannot be completed.²

NIST's identity and authority concept paper identifies questions about authorization and auditing for software agents. Its status is exploratory project development, not certification. Managers should require direct evidence that the proposed identity can be authenticated, constrained, monitored, and withdrawn in their own environment.³

Use a Joint Access Approval Record

Figure 1 proposes a record completed by the vendor manager, IAM lead, OT security, process owner, and provider. Each permission should be tied to the approved support task. The record should distinguish development, testing, and production and identify any route from the support environment toward operational systems.

Require explicit denial of prohibited actions. A hypothetical maintenance-document assistant may retrieve approved procedures but cannot alter work orders, change equipment settings, or issue instructions to control systems. Test these limits through the actual tool interfaces rather than rely solely on instructions in the agent's prompt.

The provider should identify who can add tools, alter permissions, or deploy a new configuration. Those administrative powers can change the effective authority even when the agent's ordinary credential remains unchanged. Buyer approval should be required for changes that expand the approved exposure.

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>.

³ National Institute of Standards and Technology, "New Concept Paper on Identity and Authority of Software Agents," National Institute of Standards and Technology, February 5, 2026, project status and identification, authorization, auditing, and non-repudiation questions, <https://www.nist.gov/news-events/news/2026/02/new-concept-paper-identity-and-authority-software-agents>.

Figure 1. Approve and withdraw infrastructure agent access

Record element	Buyer decision	Joint test
Task and environment	Approved low-risk support scope	No route to prohibited action
Identity and grant	Principal sponsor permission and expiry	Attribution and least privilege
Administrative change	Who can expand tools or access	Change approval and denied excess
Revocation	Safe withdrawal and outstanding work	Sessions delegates and queues reconciled

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.

Test Revocation Without Creating an Unsafe State

Revocation should reach active sessions, delegated credentials, queued tasks, and cached access. The engineering team should determine which access can be removed immediately and which operations require a controlled transition. Managers should not prescribe abrupt network isolation without assessing its effect on the physical service.

Use a safe representative test environment and record its limits. Confirm that the agent cannot continue prohibited activity after revocation and that necessary human operations remain available. Reconcile outstanding work before reissuing authority, so uncertainty does not produce duplicate maintenance instructions or unsupported state assumptions.

DOE's Stormbreaker testbed illustrates why configuration and environment matter in critical-infrastructure AI evaluation. It varies components and conditions to examine capability limits. The proposed access test applies that principle locally; a testbed result elsewhere does not establish that this supplier configuration is safe.⁴

⁴ 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>.

Maintain Local and Contractual Accountability

The contract should require current identity inventories, access evidence, incident cooperation, and notice of material changes. The provider should obtain equivalent cooperation from subcontractors with relevant access. A buyer that cannot discover who is acting cannot reliably determine whether its authority boundary remains intact.

Reconcile the approved identity record against live grants before acceptance. An expired test credential, emergency account, or unused integration token can provide a route outside the proposed support scope. The provider and IAM owner should identify such grants and confirm their disposition. Where access must remain for another legitimate service, document the separate purpose and enforcement boundary rather than assume that removing the agent's primary credential ends every relevant permission.

EU, UK, Swiss, and US infrastructure operators have different sector, cybersecurity, data, and safety obligations. The service owner should obtain a local applicability decision and connect it to the common access record. Neither a foreign security guideline nor a supplier certification replaces those duties.

Review credentials against actual use. Remove obsolete grants, investigate unexplained access, and confirm the renewal owner remains in role. An inventory that is not reconciled to live systems can preserve permissions after the task or relationship has ended.

Release Only the Access Demonstrated by the Test

The manager should approve the smallest permission set that completes the authorized support task. Record failed denial or revocation tests as release conditions, assign corrective owners, and retest before expanding access. Where engineering evidence is incomplete, keep the service outside consequential control paths.

The result should be an identity record that explains who acts, why the access exists, how it is limited, and how it ends. That record gives both firms a workable access process while preserving the operator's responsibility for safe essential service.

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 joint infrastructure agent access record with engineering-led safe revocation and explicit support-only limits. 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>.

National Institute of Standards and Technology. "New Concept Paper on Identity and Authority of Software Agents." National Institute of Standards and Technology. February 5, 2026. <https://www.nist.gov/news-events/news/2026/02/new-concept-paper-identity-and-authority-software-agents>.

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."