

Keep a Human at the Console: Oversight and Liability for Agentic AI

Funding review capacity and allocating responsibility for bounded utility agent workflows

Authored by Chennakeshav Adya

Independent Researcher

October 2026

ABSTRACT

A utility approving an agentic AI workflow needs a defined action boundary, qualified human review and records connecting each request to its authorised outcome. This paper proposes an investment and contracting framework for that decision. Its initial application is a low-risk information-technology workflow that prepares maintenance records using approved, non-sensitive information. Direct plant control and safety decisions are excluded. Eight primary references establish the guidance context; the proposed commercial framework requires utility-specific engineering, security and legal review.

Three figures and three tables organise action rights, evidence readiness, review costs and staged commercialisation. A fictional investment assumes USD 600,000 initial expenditure, USD 500,000 annual gross cash benefit, USD 180,000 platform expenditure and USD 60,000 separate evidence expenditure. Human review costs vary with daily request volume under explicit time and staffing assumptions. Five-year discounted cash cases demonstrate the effect of review workload on affordability. The illustration contains no observed productivity results, incident probabilities, supplier prices or estimates of a particular utility's value.

The recommendation is to fund oversight as an operating commitment and approve permissions at the individual action level. Acceptance requires demonstrable execution boundaries, reconstructable records and a workable human review process. Contractual allocation should follow verified responsibilities and the applicable legal position. Human approval alone establishes neither operational safety nor immunity from liability.

JEL Classification: G31, L94, L95, M15, O33, K12

Keywords: agentic AI, utilities, human oversight, authorisation, operational technology, audit evidence, liability allocation, investment appraisal, workflow governance, maintenance planning

1 The oversight decision

A utility considering an agentic information service should approve the work it intends to delegate, the rights needed to perform it and the people responsible for reviewing the result. The investment decision includes continuing review effort. A proposal that costs software and integration while leaving reviewer availability unresolved has an incomplete operating budget. Define the permitted workflow before estimating its economic contribution.

The initial application examined here prepares draft maintenance records from an approved extract of non-sensitive information. It operates outside the plant control environment. An authorised employee checks the proposed record and decides whether to enter it into the existing maintenance process. No automated change to equipment settings, safety systems or operational control logic is included. The utility must independently assess whether the proposed information and connections satisfy that boundary.

The framework connects an action-rights record, a human review plan, preserved execution evidence and contractual responsibilities. Each component answers a different approval question. Permission determines the actions available to the service. Staffing determines whether the required reviews can actually occur. Evidence supports investigation of what happened. Contracts establish obligations whose legal effect needs qualified review.

An investment committee should receive a specific choice. State the workflow, approved users, information perimeter, proposed expenditure and acceptance requirements. Include the option of financing a limited evaluation while withholding production permission. A successful demonstration supports conclusions only about its tested conditions. Additional access, users or operational dependencies require a fresh decision with the relevant evidence and costs.

2 Define the agent and its permitted environment

The May 2026 joint guidance published by Australia's cyber security authority describes agentic systems as models connected to tools, external data, memory and planning workflows. Its principal scope is LLM-based agents in information-technology environments. It recommends low-risk, non-sensitive tasks and restricted access.[3] Those boundaries guide the example developed in this paper. They provide no endorsement of autonomous utility plant operation.

Write the service boundary in operational terms. Name the source information, permitted tool, destination record and human approver. A draft preparation service might read a sanitised equipment catalogue and propose administrative fields for a work order. The service should have no credentials for the control network and no permission to issue equipment instructions. The final maintenance action remains governed by the utility's existing engineering and work-authorisation arrangements.

Identify dependencies that could widen the boundary. A connection to a general document repository can expose information outside the approved extract. A tool that accepts arbitrary destinations can enable unintended writes. A workflow allowing new agents to inherit privileges can change the original delegation. The evaluation should describe these possibilities and test the restrictions intended to contain them.

Preserve an effective version of the boundary. Record the accountable owner, approved environment and date of approval. Changes to tools, data sources or destinations should produce an identifiable successor version. The investment review can then connect costs and benefits to a service whose rights are known. A general description such as maintenance automation supplies insufficient detail for permission design, acceptance testing or contractual allocation.

3 Read current guidance within its actual scope

NIST's final operational-technology security guide addresses systems interacting with the physical environment and recognises their performance, reliability and safety requirements.[1] The joint December 2025 AI-in-OT principles discuss operator involvement, logging, monitoring and the risks of active influence on control.[2] These documents provide context for a utility's technical review. Their existence establishes no certification of the workflow proposed here.

NIST's AI Risk Management Framework is voluntary guidance for managing AI risk.[4] Its agent identity concept announced in February 2026 is a proposed area of work addressing identity, authorisation and related security questions.[5] NIST's September 2026 update identifies software development workflows as the initial implementation use case.[6] Those publications should retain their development status; they supply no demonstrated utility deployment outcome.

The UK cyber security code addresses AI security across development and operation.[7] NIST's agent standards initiative describes work on interoperability, protocols, security and identity.[8] These sources support asking suppliers for specific evidence about the system offered. Neither an initiative announcement nor a code description determines the enforceability of a procurement clause or the acceptability of an operational use.

The proposed records and economic cases below are original analytical contributions. Their suitability requires application-specific assessment. The framework uses no published accident frequency or assumed insurance recovery to justify adoption. It separates the decision to evaluate a bounded administrative workflow from decisions involving physical assets. Engineers, operators, security reviewers and legal advisers should define the conditions for any actual deployment.

4 Establish an action rights record

Create one record for every tool operation permitted to the agent. Describe the object it can read or change, the initiating user and the conditions requiring human approval. Use explicit destinations and operation types. A service permitted to prepare a draft should lack rights to mark the maintenance task complete or authorise its execution. Those rights belong to separately approved processes.

Table 1. Proposed action rights and approval record

Operation	Proposed permitted scope	Required authority
Read information	Approved non-sensitive extract	Data owner approval
Prepare draft	Specified administrative fields	Workflow owner approval
Change approved record	Excluded from initial service	Separate change decision
Authorise maintenance	Excluded from agent permissions	Existing utility process
Invoke another agent	Excluded from initial service	New delegation review
Affect plant controls	Excluded from paper's example	Independent engineering assessment

Table 1 is an original proposed allocation. It contains no observed configuration or assertion that these restrictions have been implemented. For an actual system, include the relevant account, environment and tool specification. Verify that the effective permissions match the written record. A contractual promise about limited access needs evidence from the deployed configuration and tests.

Make the record understandable to the person granting authority. Show the exact proposed change before approval, including destination and affected fields. Record rejection and modification as valid outcomes. Avoid automatically broadening permissions when a request fails. An unresolved action should return to the named workflow owner. The investment budget should include the time needed to administer this record and review material amendments.

5 Evaluate evidence readiness

Evidence readiness can be organised as a sequence of decisions. Begin with a defined scope, test the rights granted, rehearse the human review process and reconstruct completed requests from preserved records. Expansion follows a separate review of the proposed additional work. Figure 1 expresses this sequence qualitatively. Its positions have no calibrated relationship to technical maturity, incident probability or commercial conversion.

Figure 1. Proposed oversight evidence readiness sequence

At scope definition, identify the information perimeter and actions excluded. Rights testing should demonstrate permitted and rejected operations using the actual tool interface. Review rehearsal should establish whether authorised staff understand the information presented and can return a request for correction. Reconstruction should connect the request, tool invocation, approval and observed destination outcome without depending on the agent's own narrative.

An incomplete step should produce a specific decision. If a reviewer cannot see the proposed destination, revise the interface and repeat that test. If execution evidence cannot be preserved under the supplier's terms, resolve

access and retention before commitment. Assign the relevant expenditure and owner. A broad readiness percentage would obscure the particular condition still unresolved.

Treat expansion evidence separately. Increasing request volume changes staffing needs. Adding another destination changes permission and audit requirements. Introducing sensitive information changes the security assessment. The evaluation should state which records remain reusable and which tests require repetition. The proposed sequence supports a controlled investment discussion; it supplies no universal deployment timetable and cannot replace the utility's own technical acceptance procedures.

6 Authorise each tool invocation

The joint agentic guidance identifies risks from privileges assessed only at start-up and then reused for subsequent invocations.[3] A proposed utility service should therefore be evaluated at the point where it requests a specific tool action. The authorisation decision needs the current identity, permitted operation, destination and relevant approval state. A prior session approval should not silently confer rights over a changed request.

Separate the agent's request from the component enforcing permission. The proposed architecture places policy enforcement at the tool boundary under the utility's control. The model may prepare a request; the executing service checks whether that request falls within approved rights. The implementation must be tested for its actual interface, credentials and error behaviour. This description establishes a design requirement, not evidence that a product satisfies it.

Preserve denied requests as reviewable events with appropriate access restrictions. They can help identify ambiguous scope or attempted privilege escalation. A denied operation should not trigger an unrestricted substitute tool or an agent-created account. Define the permitted response to denial before deployment. The human reviewer needs sufficient context to distinguish a correct rejection from a configuration defect.

Budget for maintaining the policy. New tools and revised schemas can alter the meaning of an invocation. The workflow owner should approve changes against the action record, and technical reviewers should repeat the relevant tests. Procurement terms should identify who supplies change notices, who implements restrictions and who preserves the resulting evidence. Ongoing permission administration is part of the service's operating cost.

7 Give human review an executable purpose

Specify what the human reviewer is expected to decide. In the example, the reviewer checks the draft's destination, required fields and consistency with approved source information. The reviewer can approve administrative entry, amend the draft or reject it. Approval grants no permission to perform maintenance or alter plant operation. Those actions require the utility's separate work procedures.

Provide the reviewer with the original information used and a clear presentation of the proposed changes. Distinguish values copied from the extract from values proposed by the model. Missing information should remain visible. A plausible sentence should not be accepted as evidence that a required technical condition exists. The reviewer must have access to the records needed for the actual decision.

Record the scope of the approval. Identify the person, time, approved version and destination. If the proposed request changes after review, invalidate the earlier approval and seek a fresh decision. The system should be tested for this behaviour. A confirmation screen and a logged button press alone supply limited evidence about whether the reviewer saw the same request later executed.

The utility should assess qualifications and authority independently. A person able to review administrative completeness may lack authority to make engineering decisions. Training should cover the permitted task, refusal conditions and escalation route. Review effectiveness requires evaluation under actual working conditions. This paper proposes no universal review time and no assurance that adding a human approval step makes an application safe.

8 Fund review capacity and exception handling

Prepare a workload estimate from requests requiring approval, expected review effort and the work generated by exceptions. Separate ordinary review from investigations, training and permission changes. For an actual proposal, collect evidence from a bounded evaluation and describe the sample. Until that evidence exists, identify workload estimates as management assumptions. Retain the assumptions alongside the commercial budget.

Design the service around available qualified capacity. A request arriving outside staffed hours needs a defined treatment, such as remaining pending for the next authorised review. It should not gain broader permission because a reviewer is unavailable. If time-sensitive work is proposed, the utility must independently determine suitable coverage and operational arrangements before including it in the service boundary.

Average annual hours provide an initial cost comparison. Staffing also depends on concurrent requests, absences, review difficulty and the need for independent escalation. A fractional annual workload cannot establish the actual roster. The model later in this paper deliberately separates a cost-equivalent calculation from any claim about minimum staffing. A deployment plan needs evidence relevant to peak workload and qualifications.

Set a maximum pending workload and an escalation owner. Review that limit during the evaluation using the actual workflow. When capacity is exceeded, additional requests can remain queued or the service can suspend draft generation under approved rules. Price the effects on usable output and expected cash benefit. Increasing automated volume without funding the required review can leave the investment case internally inconsistent.

9 Preserve reconstructable evidence

An incident review needs records showing the request, information used, permission decision, human approval and destination outcome. The agent's explanation of its behaviour should be preserved where useful, with its evidential limitations identified. Independent tool and destination records are needed to establish what action actually occurred. Access and retention requirements should be agreed before production use.

Table 2. Proposed minimum evidence for reconstructing a request

Evidence item	Required connection	Review purpose
Request identifier	User and workflow version	Establish initiating scope
Source snapshot	Approved extract and timestamp	Identify information used
Tool request	Operation and destination	Inspect proposed action
Permission result	Policy version and decision	Review effective authority
Human decision	Reviewer and approved version	Establish approval scope
Destination result	Outcome and recorded changes	Verify execution

Table 2 sets out proposed records; their legal admissibility and retention period require separate assessment. Protect sensitive records through appropriate access controls. Preserve a method for detecting unauthorised alterations, and test retrieval before acceptance. The commercial agreement should identify which party holds each record and how the utility can obtain it during a dispute or supplier exit.

Define what happens when logging fails. A service that continues executing after losing required evidence presents a different control condition from the one approved. The initial workflow should have an agreed response, including suspension of affected operations. Test that response with the technical team. Evidence preservation has a continuing cost, including storage, review, retrieval and the maintenance of access rights.

10 Control changes to models tools and credentials

A model update can change output behaviour while leaving the visible workflow title unchanged. A tool update can change available operations or the interpretation of fields. Track these changes against the approved scope. Procurement terms should establish notice, evaluation and rollback arrangements appropriate to the service. Acceptance of one version should have an explicit relationship to later versions.

Maintain separate identities for the agent service and human reviewers. Review credential lifetime, storage and revocation with security specialists. Shared accounts can make it difficult to distinguish an automated action from an employee action. NIST's identity concept identifies relevant authorisation and audit questions.[5] It remains a concept for further work and does not certify a particular identity implementation.

The initial service should exclude unreviewed delegation to other agents. If a supplier proposes delegation, identify each downstream identity, inherited right and tool connection. Test the complete path and revise the record of responsibilities. A supplier's description of a single agent may conceal several components with different operators and evidence locations. The commercial agreement should address the components actually used.

Allocate the costs of change. A mandatory supplier update may require repeat testing, reviewer training and revised evidence storage. Identify who pays those costs and whether the utility can delay or reject the change. Review continuity arrangements if the earlier version becomes unavailable. These obligations affect both annual expenditure and the usable term of the investment.

11 Rehearse interruption and degraded operation

Define an interruption procedure suitable for the permitted administrative workflow. Staff should know how to suspend draft generation, revoke the relevant credential and continue the established manual process. An interrupted information service should leave pending records in an identifiable state. The utility's existing operational procedures remain authoritative for maintenance and plant decisions.

Test interruption at several points in a request. A request may have been drafted, approved or partially written when the service becomes unavailable. Record which outcomes the destination system can report reliably. Avoid describing a reversal as successful without checking the affected record. Some administrative changes may require a corrective entry rather than a simple rollback.

The December 2025 OT guidance discusses oversight, monitoring and restoration considerations.[2] Application to the example still requires a technical design suitable for its actual systems. A supplier's generic recovery statement supplies insufficient evidence about the maintenance platform and credentials involved. Ask for a controlled rehearsal and preserve the observed results, including unresolved exceptions.

Include continuing rehearsal expenditure in the operating plan. Staff turnover and tool changes can make an earlier procedure incomplete. Assign responsibility for keeping the manual process usable and reviewing interruption evidence. The economic model should account for the service term during which these obligations remain funded. No interruption procedure presented here establishes a guarantee against service disruption or physical consequences.

12 Define acceptance tests before procurement

Acceptance criteria should be specific enough to determine whether the purchased service meets the approved boundary. Include permitted and prohibited requests, changed destinations, revoked credentials, incomplete source records and unavailable reviewers. Test normal completion and exceptions. Preserve the starting configuration, request, observed result and responsible reviewer for each test.

Use information suitable for the evaluation. The initial example assumes sanitised, non-sensitive inputs. The utility must verify that classification and the adequacy of the sanitisation. A test using artificial records can examine permission enforcement without establishing performance on actual workload. Describe that limitation in the acceptance report and identify the additional evidence needed for production approval.

Review evidence retrieval as a separate acceptance criterion. Select completed and rejected requests and reconstruct them from independent records. Confirm that the retrieval path remains available to the utility. If a record exists only in a supplier interface, assess contractual access and export rights. An investment committee should see any dependency on continuing supplier cooperation.

The procurement team should connect acceptance criteria to payment milestones and remedial obligations, subject to legal review. State who witnesses the tests and who can reject an incomplete result. Repeating failed tests needs an agreed cost allocation. The framework supplies no universal pass percentage; the utility's qualified reviewers should determine the criteria for its chosen workflow and environment.

13 State the fictional economic assumptions

The following arithmetic is a fictional illustration of the cost of human oversight. It uses no utility observations or supplier quotations. Initial expenditure is USD 600,000. Annual gross cash benefit is USD 500,000, platform expenditure is USD 180,000 and separate evidence expenditure is USD 60,000. The last amount covers non-staff evidence services and storage; human review expenditure is calculated separately to avoid counting the same cost twice.

The illustration assumes 250 working days each year and three minutes of review for every request. Continuing training, permission administration and exception review add 400 hours annually. Productive annual capacity is

assumed to be 1,800 hours per full-time equivalent at a loaded annual cost of USD 90,000. The resulting cost equivalent is USD 50 per hour. These values are modelling choices with no claim of market representativeness.

The investment horizon is five years and the discount rate is ten percent. Initial expenditure occurs at time zero; equal annual net receipts occur at each year end. The illustration excludes taxes, financing, working capital, inflation, implementation delay, residual value and incident losses. No probability or insurance recovery is assigned to an adverse event. A real investment needs assumptions supported by relevant evidence.

Gross benefit is held constant across request-volume cases solely to isolate review-cost sensitivity. That assumption should be reconsidered for an actual service. More requests could change usable output, review difficulty or cash conversion. The cases establish arithmetic relationships under stated conditions. They provide no productivity forecast, valuation of a utility or recommendation to adopt a particular agent product.

14 Compare request volume and human review cost

Three fictional volumes are tested: 80, 160 and 240 requests per working day. Annual ordinary review hours equal daily volume multiplied by 250 working days and three minutes, divided by sixty. Add 400 continuing hours to each case. Total annual hours are 1,400, 2,400 and 3,400. At the assumed USD 50 hourly cost, annual review expenditure is USD 70,000, USD 120,000 and USD 170,000.

Table 3. Fictional review workload and annual cash cases in USD thousands


Daily requests	Annual review hours	Review cost	Annual net cash
80	1400	70	190
160	2400	120	140
240	3400	170	90

Annual net cash equals the assumed USD 500,000 gross benefit less USD 180,000 platform costs, USD 60,000 evidence costs and the applicable review cost. Table 3 uses thousands of dollars for costs and cash; hours remain unscaled. The calculation has no additional provision for peak coverage. Its fractional staffing equivalent describes annual cost arithmetic and supplies no roster or minimum staffing conclusion.

For a real proposal, compare the ordinary request with exceptions separately. An incorrect source field may require several reviews and investigation. A difficult request may need a different qualification. Record the evaluation evidence supporting each estimate and avoid treating three minutes as a mandated review limit. Necessary review work should determine the budget. Financial affordability must be assessed using the oversight conditions the utility actually approves.

15 Benchmark complete annual recovery cost

Compare fictional annual gross benefit with recurring costs and recovery of initial expenditure over the assumed term. The five-year annuity factor at ten percent is approximately 3.790787. Dividing USD 600,000 by that factor produces an annual cost-equivalent capital recovery amount of approximately USD 158,278. This is an analytical comparison and includes no supplier profit margin or financing quote.

Figure 2. Fictional annual costs including human review and initial expenditure recovery


Adding platform costs, separate evidence costs and review costs produces recurring totals of USD 310,000, USD 360,000 and USD 410,000. Including the cost-equivalent capital recovery produces approximate annual thresholds of USD 468,278, USD 518,278 and USD 568,278. Under the constant fictional USD 500,000 gross benefit, the lower-volume case exceeds its threshold and the other cases fall below theirs.

Figure 2 compares those components without assigning a probability to any case. The threshold changes if the service term, initial cost or discount assumption changes. A supplier-funded structure would also need its actual payment timing, funding costs, margin and contractual risks. The displayed threshold cannot be used as a commercial offer or an estimate of the price available in a particular market.

The benchmark should prompt an evidence request. Ask whether gross cash benefit is supported, whether the platform and evidence budgets are complete and whether the review workload has been measured. If required oversight makes the proposal unaffordable, reconsider the workflow, scope or procurement decision. Reducing essential review merely to obtain a positive investment result would invalidate the approved operating assumptions.

16 Value the bounded investment under explicit scenarios

Discount the fictional annual net cash amounts over five years and subtract USD 600,000 initial expenditure. The results are approximately USD 120,249, negative USD 69,290 and negative USD 258,829 for the three workload cases. These figures describe the bounded workflow investment under the illustration's assumptions. They establish no enterprise valuation, forecast or probability-weighted return.

Figure 3. Fictional five year investment sensitivity and proposed commercialisation roadmap

Fictional cash cases; proposed roadmap has no promised timing or conversion probability.

Figure 3 places the cash sensitivity alongside a proposed commercial sequence. Define the scope, cost the review plan, test the permission boundary, reconstruct the evidence and review expansion. The roadmap has no promised duration. Each decision should have its own evidence requirements and expenditure authority. A positive fictional cash case cannot demonstrate that technical acceptance or contractual conditions have been met.

Keep the financial comparison consistent when considering alternatives. A conventional workflow, an existing rules-based service or a manual process may carry different costs and dependencies. Obtain actual evidence before making a procurement comparison. The paper does not assume that agentic AI is the preferred solution. It sets out the commitments that should be priced if an agentic service is proposed.

An investment committee can approve a limited evaluation with a maximum spend and defined exit conditions. It can also withhold further commitment when benefit evidence or required rights remain unresolved. The value of that decision depends on the expenditure and information actually obtained. This paper assigns no invented option value to a pilot or future expansion.

17 Allocate responsibility around verified events

Contractual responsibility should be examined using events that the parties can identify and investigate. A proposed record might distinguish incorrect source information, an unauthorised tool invocation, a changed request after approval and a destination-system failure. Each event involves potentially different actors and evidence. Legal advisers should determine the effect of the contract and applicable law on the actual facts.

Preserve separate questions about contractual breach, causation and recoverable loss. Evidence that a supplier failed an agreed control does not by itself establish every claimed loss. Evidence that an employee approved a request does not establish that all other participants fulfilled their obligations. The investigation should reconstruct the sequence and identify the relevant contractual duties without assuming a predetermined allocation.

The author proposes an incident responsibility schedule describing the party responsible for each component and the records needed to assess performance. This is a drafting aid for qualified review. It expresses no universal legal rule and contains no opinion on enforceability in a particular jurisdiction. The schedule should be revised when the service boundary or supplier chain changes.

Avoid assigning all consequences to an unnamed human approver. Identify what that person was authorised and equipped to review. A reviewer cannot meaningfully inspect information withheld by the interface or hidden downstream actions. Procurement should address those design obligations separately. Actual responsibility must be determined against the evidence, agreed scope and governing legal position.

18 Make supplier obligations specific

Identify the model provider, agent application operator, integration contractor and destination-system provider where they are different parties. Describe each component and its contractual owner. A utility contracting with one supplier should understand which obligations that supplier accepts for subcontracted services. The scope of any pass-through rights or exclusions needs explicit review.

Ask for obligations tied to the service purchased. These can include maintaining the approved tool boundary, notifying material changes, preserving records and cooperating with incident investigation. Specify evidence delivery and response arrangements. Qualify the obligations for the actual technology and contract rather than adopting broad claims about trustworthy AI. Technical reviewers should confirm that promised controls can be tested.

Allocate implementation and acceptance costs separately from continuing obligations. Permission updates, interface changes and evidence exports may involve several suppliers. State who coordinates the work and who pays when one party's change requires another party's retesting. An apparently modest annual software price can leave these costs outside the commercial comparison.

Review exit arrangements before commitment. The utility needs a practical way to retrieve its records, revoke access and continue the approved manual process. Identify formats, dependencies and the party supporting transition. These terms should remain consistent with data protection and confidentiality requirements established by advisers. The framework offers proposed commercial questions and no assurance that a supplier will agree to the requested obligations.

19 Review caps insurance and recovery separately

Assess contractual liability caps against the obligations and loss categories actually negotiated. A cap should be reviewed with its exclusions, aggregation rules, claim conditions and relevant law. This paper provides no recommended universal cap multiple. A percentage of subscription fees carries little analytical meaning without a description of the exposure and enforceable recovery arrangements.

If insurance is proposed, obtain the actual policy wording and qualified advice. Examine insured parties, covered activities, exclusions, notification requirements and the treatment of subcontracted services. A certificate or product description cannot establish that a particular agent-related event is covered. The economic illustration includes no insurance recovery because no policy or event evidence has been supplied.

Consider counterparty capacity independently. A contractual right can depend on the supplier's ability to perform or pay. Review available financial evidence and any security proposed for material obligations. Identify the conditions affecting that protection. The paper makes no claim about the credit quality of any supplier and assigns no assumed recovery rate.

Keep incident response expenditure visible even where recovery is disputed. The utility may need to investigate, preserve evidence or restore administrative records before a claim is resolved. Assign authority for those actions under the actual operating arrangements. Budget decisions and ultimate liability determinations can occur at different times. Qualified legal and insurance review should inform the proposed service contract before production acceptance.

20 Assess geography and data movements

The framework can organise reviews in different jurisdictions, including the UK, United States and Gulf markets. Its records do not establish compliance in any of them. Determine the locations of the utility, service operators, model processing, stored evidence and subcontractors. Legal advisers should identify applicable obligations from those facts and the actual service proposed.

Map information movements at the level of the approved workflow. Record what leaves the utility environment, the receiving party, purpose and retention arrangement. Sanitised information still needs assessment against the underlying dataset and the applicable classification rules. A supplier's regional marketing description may differ from the actual routing of model requests or support access.

Connect geographic review to commercial commitments. A restriction on processing location can require a different service design or additional expenditure. A requirement to retain evidence locally can change storage

and retrieval arrangements. Record those consequences before using the fictional cost model as a comparison. The illustration's dollar amounts contain no country-specific tax, wage or infrastructure assumption.

Review the contract's governing law, dispute process and rights to obtain evidence with qualified counsel. Identify the effect of cross-border subcontracting on practical enforcement and investigation. The same action-rights record may be useful across locations, while its legal and operating conditions differ. Expansion into another jurisdiction should therefore have its own approval record and supported budget.

21 Build a staged commercialisation plan

Start commercialisation with a service whose scope and evidence can be tested. The proposed sequence begins with a bounded evaluation of administrative draft preparation. Price configuration, permission tests, reviewer training and evidence retrieval. Define the resulting decision and maximum authorised spend. An evaluation agreement should state the work included and the conditions for termination or further procurement.

Production pricing should include ongoing oversight responsibilities. Separate platform fees from review labour, evidence services and change administration where they are performed by different parties. Identify fixed and workload-dependent costs. Agree how measured volumes will be reported and what happens when the approved capacity is exceeded. The fictional cases demonstrate why volume matters; actual prices require supplier offers and utility evidence.

Expansion requires a revised service boundary and financial comparison. Reusing a tested interface may reduce some implementation work, but that reduction needs evidence. New source data, users or destinations can create additional reviews. Avoid multiplying an initial pilot's assumed savings across a portfolio without establishing the conditions at each site. The plan should preserve the distinction between demonstrated scope and proposed scope.

Commercial advisory work can help assemble the investment record, coordinate procurement questions and compare funding structures. Technical acceptance remains with qualified reviewers and authorised utility decision makers. Legal allocation remains subject to counsel. This paper establishes no demand estimate, mandate pipeline or consulting fee forecast. Its commercial contribution is a structured decision process that an actual procurement team can evaluate.

22 Present an investment committee memorandum

The memorandum should begin with the requested expenditure and permission decision. Describe the workflow and exclusions in plain language. State the proposed implementation version, qualified review owner and accepted source information. Attach the action-rights record and unresolved conditions. The committee should be able to identify exactly what its approval would authorise.

Present the financial case with a bridge from gross benefit to cash receipts or avoided expenditure. Distinguish observed evaluation results from management assumptions. Include platform, review, evidence and change costs. Explain the timing of initial expenditure and the service term available under the contract. Show relevant sensitivities without attaching unsupported probabilities to them.

Summarise acceptance evidence and contractual questions. State which tests have passed, which have failed and which remain unperformed. Identify any dependence on future supplier cooperation or additional legal work. Required evidence should remain traceable to its actual record. A polished dashboard cannot replace the underlying observations or executed terms.

Recommend a defined approval, conditional approval or deferral. Each condition needs an owner and a record demonstrating fulfilment. Expansion should require a fresh memorandum when permissions or costs materially change. The committee's decision should preserve a version of the proposal and conditions reviewed. This supports later comparison between the service funded and the service actually operated without claiming that the memorandum resolves every technical or legal risk.

23 Limitations and further research

The framework has not been validated through a deployed utility agent service. Its records are author proposals and its numerical examples are fictional. The paper supplies no measured error rate, incident frequency or causal estimate of the effect of human review. Readers should apply the framework only after checking its relevance to their own environment and authorised use.

The cost model uses constant annual benefit and simplified timing. It omits peak staffing, tax, debt, working capital, delays and residual value. It assigns no cash amount to an operational incident. Those omissions limit conclusions about affordability and valuation. An actual investment should incorporate supported assumptions relevant to the service, including the cost of meeting required oversight conditions.

Further research should measure review workload under defined, low-risk conditions. Preserve the request characteristics, reviewer qualifications, exceptions and actual time spent. Examine whether independent records support reconstruction of approved and rejected requests. Report sample boundaries and missing information. Such evidence could inform budgets without transferring findings automatically to plant-control applications.

Contract research could examine how specific service obligations, evidence rights and supplier chains are treated in executed transactions. Any analysis should identify jurisdiction, agreement scope and access limitations. The evolving standards initiatives cited here require monitoring as their deliverables change. This paper reflects sources accessed on 4 October 2026 and should be reviewed against later publications before use in a new procurement decision.

24 Conclusion

The proposed oversight decision joins a bounded service, qualified human review, reconstructable evidence and responsibilities defined for legal assessment. Fund those commitments before granting production permissions. Begin with a low-risk information workflow whose actual connections and rights can be tested. Maintain separate authority for maintenance execution, plant operation and safety decisions.

The fictional arithmetic shows how review workload can change a bounded investment result under fixed benefit assumptions. Increasing daily requests from 80 to 240 raises the assumed annual human review cost from USD 70,000 to USD 170,000. The five-year cash comparison moves from positive to negative under the stated costs and discount rate. That result is a sensitivity illustration, with no prediction about a utility's performance.

The next practical decision is to define the proposed workflow and finance the evidence needed to assess it. Complete the action-rights record, identify authorised reviewers and agree the acceptance tests. Obtain an actual cost proposal and a supported benefit estimate. Review the contract's change, evidence and incident obligations with the relevant specialists.

Approval should relate to a recorded service version and its permitted environment. Subsequent expansion needs new evidence about the additional work and commitments. A disciplined record of scope, authority and outcome gives investment and procurement teams a common basis for reviewing the service. Its usefulness depends on the quality of the evidence supplied and the competence of the people applying it.

Frequently Asked Questions

What utility workflow does this paper consider?

The example prepares administrative maintenance drafts using an approved, non-sensitive information extract in an IT environment. A human reviews the proposed record under existing utility procedures. The example excludes plant controls, safety decisions and automated maintenance authorisation. The utility must verify that the actual systems and information satisfy the boundary before considering deployment.

Does a human approval button establish safe operation?

A button records an interaction. Effective oversight also requires an authorised reviewer, sufficient information, a defined decision and a connection between the approved request and executed outcome. Engineering and safety acceptability need independent assessment for the actual use. The paper makes no safety certification claim.

How should human review be included in the investment case?

Estimate request volume, review effort, exception work, training and change administration using relevant evaluation evidence. Include the resulting expenditure in recurring costs. Assess peak coverage and reviewer qualifications separately from annual hours. The fictional model illustrates the arithmetic and supplies no market staffing estimate or minimum roster.

Which evidence should be preserved for a completed request?

Preserve the initiating request, source snapshot, tool operation, permission result, human decision and destination outcome with connected identifiers and versions. Test reconstruction before acceptance. Establish access, retention and export rights with security and legal advisers. Records held by a supplier require practical retrieval arrangements.

Who carries liability when an agent makes an incorrect request?

Responsibility depends on the verified sequence, contractual obligations and applicable law. Examine source information, permissions, interface design, approval scope and execution separately. Qualified legal advisers should assess breach, causation and recoverable loss. The paper offers a review structure and no universal allocation or enforceability opinion.

Can the fictional annual thresholds be used as supplier prices?

They are cost-equivalent calculations based on fictional expenditure, term and discount assumptions. They include no quoted financing cost or supplier margin. Actual procurement requires supported costs, service obligations and payment terms. The thresholds help identify the assumptions affecting affordability.

What changes require a fresh review?

Material changes to data sources, tools, destinations, model versions, credentials, delegation or request volumes should be assessed against the approved boundary. Identify the tests, training, evidence and expenditure affected. The utility should define its change procedures and who has authority to approve each revised service version.

How can the framework be used outside the UK or United States?

Its action, review and evidence records can organise a transaction-specific assessment in another jurisdiction. Local engineering, security, privacy and legal requirements must be established independently. Map actual information movements and service locations. The paper contains no jurisdiction-specific compliance conclusion or forecast of commercial demand.

References

1. NIST. Guide to Operational Technology Security, SP 800-82 Revision 3, final 2023.
<https://www.nist.gov/publications/guide-operational-technology-ot-security> Accessed 4 October 2026.
2. Australian Signals Directorate and partner agencies. Principles for the secure integration of Artificial Intelligence in Operational Technology, 4 December 2025. <https://www.cyber.gov.au/business-government/secure-design/operational-technology-environments/principles-for-the-secure-integration-of-artificial-intelligence-in-operational-technology> Accessed 4 October 2026.
3. Australian Signals Directorate and partner agencies. Careful adoption of agentic AI services, 1 May 2026.
<https://www.cyber.gov.au/business-government/secure-design/artificial-intelligence/careful-adoption-of-agentic-ai-services> Accessed 4 October 2026.
4. NIST. AI Risk Management Framework, voluntary framework and current programme materials.
<https://www.nist.gov/itl/ai-risk-management-framework> Accessed 4 October 2026.
5. NIST. New Concept Paper on Identity and Authority for Software Agents, 5 February 2026.
<https://www.nist.gov/news-events/news/2026/02/new-concept-paper-identity-and-authority-software-agents> Accessed 4 October 2026.
6. NIST. Comments on Software and Agentic AI Identity Concept Paper, 29 September 2026.
<https://www.nist.gov/news-events/news/2026/09/comments-software-and-agentic-ai-identity-concept-paper> Accessed 4 October 2026.
7. UK Department for Science, Innovation and Technology. Code of Practice for the Cyber Security of AI, 31 January 2025.
<https://www.gov.uk/government/publications/ai-cyber-security-code-of-practice/code-of-practice-for-the-cyber-security-of-ai> Accessed 4 October 2026.
8. NIST. Announcing the AI Agent Standards Initiative for Interoperable and Secure Innovation, 17 February 2026.
<https://www.nist.gov/news-events/news/2026/02/announcing-ai-agent-standards-initiative-interoperable-and-secure> Accessed 4 October 2026.

About the Author

Authored by Chennakeshav Adya
Independent Researcher

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.