


THE AI INSTITUTE / OPERATING INTELLIGENCE

FLAGSHIP EVIDENCE PAPER · 04 / 2026



Draw the AI Accountability Boundary Before You Scale

The board's decision on which outcomes, evidence, authority and exit rights stay inside the enterprise—and what happens when they do not.

Boards, C-suite and senior management

Research period: January 2023 – September 2026

First published: 1 September 2026

Current to 1 September 2026

PUBLICATION RECORD

An Institute thesis and decision framework

An AI supplier can provide capability and some controls, but enterprise accountability does not move with the contract. Boards should treat retained control capacity as operating capital and gate scale on evidence that every consequential outcome, input, acceptance decision, authority, incident response and exit route still has an empowered owner inside the enterprise.

WHAT LEADERS SHOULD TAKE FROM THIS PAPER

- 01

Approve an accountability boundary before approving a model, platform or transformation partner.
- 02

Fund retained control capacity as part of the investment case, not as compliance overhead added after deployment.
- 03

Require independent evidence for business outcomes, rights, performance, incidents and exit—not only supplier assurance.
- 04

Match the boundary to local law and institutional capacity while preserving one global enterprise principle: consequential accountability remains owned.
- 05

Sequence or stop scale where no internal owner can challenge the system, contain harm or change supplier without losing the operation.

Research method

This flagship evidence paper synthesises Institute daily research briefs published from 6 August to 1 September 2026 and refreshes their strongest recurring signal against primary government and multilateral guidance, official regional datasets, a US public-sector audit, two current NBER working papers and the OECD's 2026 cross-study productivity review. The topic was scored against executive decision relevance, evidence strength, novelty, global applicability, Institute fit and multi-channel potential. Claims retain the source population, geography, period and study type. Guidance is used as normative evidence, not adoption proof. Supplier integration is treated as a market-structure observation rather than independent outcome evidence.

First published 1 September 2026 · Updated 1 September 2026 · Research period January 2023 – September 2026 · Version 1.0 · Suggested citation: The AI Institute, *Draw the AI Accountability Boundary Before You Scale* (2026).

CONTENTS

Inside this paper

01	Decide what the enterprise must still be able to know, judge and stop The strategic decision
02	Adoption is widespread; enterprise impact is not yet automatic The evidence gap
03	The contract does not supply the missing management system What buyers are learning
04	Fund six obligations before funding scale The retained control portfolio
05	Retained control capacity is operating capital The Institute thesis
06	The obligation travels; the control design does not Global principle, local implementation
07	A stronger boundary can include more outsourcing, not less Contrary evidence and exceptions
08	Put one owner, six decisions and three capital gates in front of the board The 90-day executive agenda
M	Method and limitations
R	References

How to use this brief

Start with the decision and use the framework to challenge the current plan. Figures retain their original populations and definitions. Institute analysis is labelled, and limitations are recorded before the references.

Decide what the enterprise must still be able to know, judge and stop

AI procurement is moving from a discrete software choice to a stack decision. Models, data services, evaluation, infrastructure, workflow tools and implementation support increasingly arrive as one managed path. That can compress delivery time and give smaller organisations access to capabilities they could not build alone. It can also make the supplier's interface look like the organisation's operating authority.

The board-level question is not whether every component should be internal. It is whether the enterprise can still define the outcome, establish lawful and reliable inputs, accept or reject performance, constrain consequential action, respond to incidents and change course. Those are the six retained obligations in this paper. Capability may be bought. Accountability for the consequence remains with the enterprise.

This distinction is easy to lose in a successful pilot. Supplier dashboards, evaluation scores and service commitments can all be valuable, but they describe the service through the provider's operating frame. The enterprise still needs a view of the people, money, rights, commitments and operational dependencies affected by that service.

The boundary test: if the supplier fails, changes terms or produces a contested outcome, can the enterprise still know what happened, decide what to do and keep the operation under control?

Adoption is widespread; enterprise impact is not yet automatic

The strongest current cross-country executive survey is a useful warning against both complacency and hype. NBER Working Paper 34836 surveyed nearly 6,000 senior executives at firms in the United States, United Kingdom, Germany and Australia. Sixty-nine per cent reported active AI use, yet nine in ten reported no impact on employment or productivity over the prior three years. The same respondents expected material gains over the next three years. The study is a working paper based on reported firm experience in four advanced economies; it is not a global census or a causal estimate.¹

The counter-evidence matters. A separate 2026 NBER working paper using nearly 750 US corporate executives found positive but heterogeneous labour-productivity gains, strongest in high-skill services and finance, and a gap between perceived and measured gains. OECD's 2026 productivity review reports emerging task-, firm- and industry-level evidence, including a short-run association of about 4% higher labour productivity among 12,000 firms across 27 EU countries, while stressing intangible-investment and measurement problems.^{2 3}

The implication is not to slow every programme. It is to replace adoption as the investment proof. Scale should depend on a visible chain from workflow change to measured outcome, with an accountable owner and a comparison that would still be credible if the supplier were not grading its own work.

EXHIBIT / USE, REPORTED IMPACT AND EMERGING MEASURED VALUE ARE DIFFERENT EVIDENCE

69% Active use Share of surveyed firms reporting active AI use across four advanced economies. NBER WP 34836	9 IN 10 No reported prior impact Executives reporting no employment or productivity impact over the prior three years. NBER WP 34836	~4% Short-run association Higher labour productivity associated with adoption among 12,000 firms in 27 EU countries. OECD 2026 synthesis
--	--	--

The contract does not supply the missing management system

The US Government Accountability Office reviewed 13 AI acquisitions across four federal agencies through fiscal year 2025. It found recurring challenges in specialist expertise, requirements and contract terms, data and intellectual-property protections, early testing, continuous evaluation, pricing and total cost. The agencies had not systematically captured lessons, including when an acquisition was retired. This is public-sector US evidence from a deliberately selected sample, not a prevalence estimate for global business. Its operational lessons are nevertheless concrete: write for portability and knowledge transfer, test before larger commitments, define data rights and record what failed.⁴

The United Kingdom's updated Data and AI Ethics Framework makes the buyer's retained role explicit when a third-party supplier is involved: responsibilities, risk allocation, operating information and legal accountability have to be agreed; where responsibility is unclear, the buyer should identify and assume it. NIST's voluntary AI Risk Management Framework similarly places executive responsibility, third-party risk, contingency for third-party failure, inventory and safe decommissioning inside organisational governance. These are guidance frameworks, not proof that firms implement them.^{5 6}

A supplier can deliver excellent assurance services. It cannot remove the buyer's need to decide whether those services are sufficient for the actual use, population, market and risk appetite. The management system around the contract remains an enterprise asset.

Institutional learning is part of that asset. Record rejected designs, failed tests, incidents, workarounds, realised cost and the assumptions that changed. Without that record, each renewal or new procurement begins from the supplier's current proposition rather than the enterprise's accumulated evidence.

Fund six obligations before funding scale

Treat the following controls as a portfolio of operating assets rather than a compliance checklist. Each needs an accountable executive, defined decision rights, evidence that can be inspected independently and a funded response when the evidence is weak.

First, outcome ownership: a business leader defines the result, baseline, affected population and stop condition. Second, authoritative inputs and rights: a data or domain owner establishes which sources can be used, where they came from and what permissions travel with them. Third, independent acceptance: an assurance owner tests performance, failure modes and business consequences against enterprise criteria rather than supplier claims.

Fourth, consequential authority: an operating executive decides what the system may recommend, commit, approve or execute and where human intervention is mandatory. Fifth, incident and recovery authority: a named leader can contain harm, preserve evidence, notify affected parties and restore a safe service. Sixth, exit and continuity: procurement and operations maintain practical routes to retrieve data, transfer workflows, preserve records, substitute components or operate safely through a supplier disruption.

Not every control has to be performed by an employee. An external evaluator, sector consortium, insurer, managed security provider or regulator may supply stronger capability. Retention means the enterprise chooses the evidence standard, can inspect the result and holds the decision right when evidence is contested.

EXHIBIT / SIX RETAINED OBLIGATIONS

OUTCOM E Business consequence Baseline, benefit, affected population and stop condition. Business owner	INPUTS Authority and rights Source, permission, quality and regional suitability. Data or domain owner	ACCEPT Independent judgment Enterprise tests, comparison, uncertainty and sign-off. Assurance owner	ACT Consequential authority Recommendation, commitment, approval and execution limits. Operating executive
RECOVE R Incident control Containment, evidence, notification and safe restoration. Risk and operations	EXIT Strategic choice Data, workflow, records, substitute path and continuity. Procurement and operations		

Retained control capacity is operating capital

Financial cases commonly place licence, integration and compute on the investment side while treating evaluation, data stewardship, workflow redesign, training, incident readiness and exit testing as overhead. That accounting choice makes a thin control boundary look artificially efficient. OECD's 2026 review notes that the intangible investments required for AI adoption—including training and changes to business processes and software—are imperfectly captured in firm accounts and national statistics.³

The Institute proposes a control-capital account beside the AI investment case. It does not capitalise expenditure for accounting purposes. It makes the economic dependency visible. For each retained obligation, record the internal capacity required, the supplier contribution, the independent evidence, the failure exposure, the substitute route and the cost to maintain them through the planned operating life.

This changes the capital gate. A low subscription price is not a low-cost operating model if the enterprise cannot measure the benefit, prove the authority behind inputs, contain an incident or leave without rebuilding the workflow. Conversely, a managed service can be the rational choice for a smaller firm when it supplies scarce capability and the consequential boundary remains narrow, inspectable and reversible.

The account should expose concentration as well as direct cost. When the same provider supplies the model, workflow, evaluation, incident evidence and migration path, each service may appear efficient in isolation while the combined dependency becomes expensive to challenge. The board needs the portfolio view before the dependency is operationally irreversible.

Control capital is not an accounting label. It is the funded capacity to remain the principal in an AI-enabled operating model.

The obligation travels; the control design does not

In the European Union, deployers of high-risk AI systems have legal duties that include using systems according to instructions, monitoring operation, acting on risks or serious incidents, assigning equipped human oversight and ensuring relevant input data where they control it. Application dates and classifications remain material, and the current framework is changing through implementation guidance.⁷

The United States offers influential but often voluntary enterprise frameworks alongside sectoral law and procurement rules. Singapore's 2026 Model AI Governance Framework for Agentic AI emphasises that humans remain ultimately accountable and combines technical and non-technical measures. Australia provides a public-sector assurance framework and policy architecture that can inform a common enterprise vocabulary, but it does not replace sector law, contractual duties or implementation capability.^{6 8 9}

The same control design cannot simply be exported to markets with different infrastructure, language, skills or institutions. The World Bank's 2025 global report shows large gaps in connectivity, compute, local context and competency: high-income economies held 77% of colocation data-centre capacity, while low-income economies held less than 0.1%; internet use and basic digital skills also diverged sharply. ECLAC's 2025 index finds materially different infrastructure, talent, adoption and governance maturity across 19 Latin American and Caribbean countries.^{10 11}

For Africa and the Middle East, country and sector conditions must replace regional shorthand. Infrastructure and language constraints may justify smaller or local models, managed platforms or shared assurance capacity. Sovereign-investment priorities may change data-location and procurement choices. None of those conditions transfers the enterprise's responsibility for the consequential use; they change how much control capacity can be internal, shared or independently supplied. Workforce exposure also varies with occupation, income and digitisation: the ILO's 2025 global task index estimates that one in four workers is in an occupation with some generative-AI exposure, while only 3.3% of global employment sits in its highest exposure category. Exposure is not observed displacement, and the study concludes that job transformation is the more likely aggregate effect.¹²

A stronger boundary can include more outsourcing, not less

This paper does not argue that vertical integration or managed AI services are inherently undesirable. Integrated suppliers can lower transaction cost, coordinate security improvements, provide specialist evaluation and make advanced capability accessible to smaller organisations. Positive productivity evidence in selected tasks, sectors and firms shows that value can arrive before a complete enterprise measurement system exists.^{2 3}

The practical exception is proportionality. A low-consequence drafting tool does not need the same control portfolio as a system that influences credit, employment, health, safety, public entitlement or material financial commitments. A small enterprise may reasonably rely on supplier controls when its alternative is no access at all. A globally standardised control may also be stronger than a weak local process.

But proportionality is a decision, not an absence. Leaders should document why a use is low consequence, what would change that classification, which supplier evidence they accept and how they would stop or substitute the service. Integration becomes strategically dangerous when the organisation can no longer see the dependency or exercise judgment at the point where the consequence is created.

Put one owner, six decisions and three capital gates in front of the board

Name one accountable executive—normally the COO for enterprise deployment, with the CIO or CTO owning platform integrity and the CRO or equivalent holding independent challenge. Give that owner authority to stop scale when a retained obligation lacks evidence. The board or an authorised committee sets risk appetite, approves the boundary and reviews exceptions; management owns implementation.

In days 1–30, select the three most consequential AI-enabled workflows and map the six retained obligations. Record the current supplier, internal owner, decision rights, evidence, regional variation and substitute path. In days 31–60, run one independent acceptance test and one tabletop exercise covering supplier failure, contested output and data or rights challenge. In days 61–90, put the control-capital account beside the benefit case and decide whether to scale, sequence, redesign or stop.

Release capital in three stages. Discovery capital requires an explicit outcome and authorised inputs. Deployment capital requires independent acceptance, bounded authority and incident readiness. Scale capital requires realised outcome evidence, tested continuity and a reviewed regional operating note. The board should receive the exceptions, not a catalogue of pilots.

The board pack: one accountable owner, six retained obligations, three staged capital gates, current evidence and every exception to the approved boundary.

EXHIBIT / THREE CAPITAL-RELEASE GATES

DISCOVER	DEPLOY	SCALE
Outcome and authority Baseline, affected population, authorised inputs and explicit stop condition. Release limited discovery capital	Acceptance and containment Independent test, bounded authority, incident response and operating owner. Release deployment capital	Realised value and continuity Outcome evidence, tested substitute route and reviewed regional operating note. Release scale capital

RESEARCH RECORD

Method and limitations

Method

This flagship evidence paper synthesises Institute daily research briefs published from 6 August to 1 September 2026 and refreshes their strongest recurring signal against primary government and multilateral guidance, official regional datasets, a US public-sector audit, two current NBER working papers and the OECD's 2026 cross-study productivity review. The topic was scored against executive decision relevance, evidence strength, novelty, global applicability, Institute fit and multi-channel potential. Claims retain the source population, geography, period and study type. Guidance is used as normative evidence, not adoption proof. Supplier integration is treated as a market-structure observation rather than independent outcome evidence.

Limitations

The cross-country firm survey covers four advanced economies and relies on executive reports. The second NBER study is a US working paper; neither is peer-reviewed journal evidence at the research cutoff. OECD findings combine studies with different definitions and methods and include associations that do not establish causality. GAO examined 13 deliberately selected US federal acquisitions, so its findings are operational lessons rather than global prevalence estimates. NIST, UK, EU, Singapore and Australian materials differ in legal force and scope. World Bank and ECLAC evidence describes national foundations and ecosystems, not individual enterprise controls. Sector law, contract enforceability, institutional capacity, infrastructure, language and market structure can materially change implementation.



Sources and limitations

Source facts link to the original publication. Institute analysis reconciles definitions and turns the material into a management proposition. Frameworks are decision aids, not claims of universal causality. Read every important statistic with its population, date, definition and caveat.

Research cutoff: 1 September 2026. Review cycle: six months, or earlier after a material change in the underlying facts, standards or policy.

REFERENCES

Evidence behind the thesis

-
- | | |
|----|--|
| 01 | NBER, Firm Data on AI |
| | Working Paper 34836, revised March 2026; nearly 6,000 senior executives in US, UK, German and Australian firms; self-reported use and effects. |
-
- | | |
|----|--|
| 02 | NBER, Artificial Intelligence, Productivity, and the Workforce |
| | Working Paper 34984, March 2026; survey of nearly 750 US corporate executives; positive but heterogeneous productivity findings and perceived-versus-measured gap. |
-
- | | |
|----|---|
| 03 | OECD, Compendium of Productivity Indicators 2026 |
| | Official June 2026 synthesis of task, firm and industry evidence; reports measurement limits and imperfect capture of intangible AI investment. |
-
- | | |
|----|--|
| 04 | US GAO, Artificial Intelligence Acquisitions |
| | April 2026 audit of 13 acquisitions across DOD, DHS, GSA and VA; selected public-sector cases, not a representative business survey. |
-
- | | |
|----|--|
| 05 | UK Government, Data and AI Ethics Framework |
| | Current UK public-sector guidance on third-party accountability, contracts, risk allocation and supplier information; normative rather than implementation evidence. |
-
- | | |
|----|---|
| 06 | NIST, AI Risk Management Framework Core |
| | Voluntary US framework, version 1.0 under revision at cutoff; executive responsibility, third-party controls, contingency and decommissioning outcomes. |
-
- | | |
|----|--|
| 07 | European Commission, Navigating the AI Act |
| | Official current FAQ on provider and deployer duties, including high-risk monitoring, incident action, human oversight and input-data responsibilities; application depends on classification and dates. |
-
- | | |
|----|---|
| 08 | Singapore IMDA, Model AI Governance Framework for Agentic AI |
| | Official January 2026 enterprise guidance; voluntary framework emphasising ultimate human accountability. |
-
- | | |
|----|--|
| 09 | Australian DTA, National framework for the assurance of AI in government |
| | Australian public-sector assurance architecture; comparative implementation lens, not evidence of private-sector adoption. |
-
- | | |
|----|--|
| 10 | World Bank, Digital Progress and Trends Report 2025: Strengthening AI Foundations |
| | Global multilateral report on connectivity, compute, context and competency; country-income evidence, not firm-level controls. |
-
- | | |
|----|---|
| 11 | ECLAC, Latin American Artificial Intelligence Index 2025 |
| | Official comparison of enabling factors, R&D and adoption, and governance across 19 Latin American and Caribbean countries. |
-

12 ILO, Generative AI and Jobs: A Refined Global Index of Occupational Exposure

May 2025 global task-exposure working paper; exposure is not observed displacement and transformation is the more likely aggregate effect.

THE AI INSTITUTE

Make the decision.
Then measure what changes.

Research · Readiness · Capability

theai.institute