

C E O / B O A R D B R I E F I N G • M A R C H
2 0 2 6 E D I T I O N

Enterprise AI Adoption Decisioning

“The most expensive AI decision in 2026 is not in your technology budget. It is in your hiring plan.”

A decision guide for boards and executives on where to adopt AI, how to govern it, when to hire versus invest in compute, and what control infrastructure must exist before you scale.

Tunde Kehinde • Cyril Simone

tk@5054holdings.com • Cyril.simone@dimensioneleven.com

What's Inside

01	How to read this paper	p. 3
02	One-Page Decision Framework	p. 4
03	Executive Summary	p. 5
04	What Changed Since the February 2026 Edition	p. 7
05	Why AI Pays: The Business Case That Actually Holds	p. 9
06	Ownership: Where Most AI Programs Fail Quietly	p. 9
07	Executive Support: The Air Cover That Makes or Breaks AI Programs	p. 10
08	What You Are Actually Funding When You Fund AI	p. 10
09	Should You Hire More People or Increase Token Budget?	p. 10
10	From Copilot to Agent: The Decision That Changes Your Risk Profile	p. 11
11	Model Context Protocol: The New Infrastructure Layer	p. 11
12	Vendor or Strategic Partner: A Different Decision Entirely	p. 12
13	The Control Stack: What Must Be in Place Before You Scale	p. 14
14	What Cannot Be Optional: Your Compliance and Control Baseline	p. 14
15	The Real Cost of AI: What Boards Are Getting Wrong	p. 16
16	Your 18-Month Playbook: Fund the Gate, Not the Hype	p. 17
17	Private Equity and AI Labs: A New Enterprise Adoption Channel	p. 18
18	Why Most AI Programs Stall: The Organizational Debt Few Boards Are Addressing	p. 18
19	Board Toolkit: Frameworks You Can Use Tomorrow	p. 19

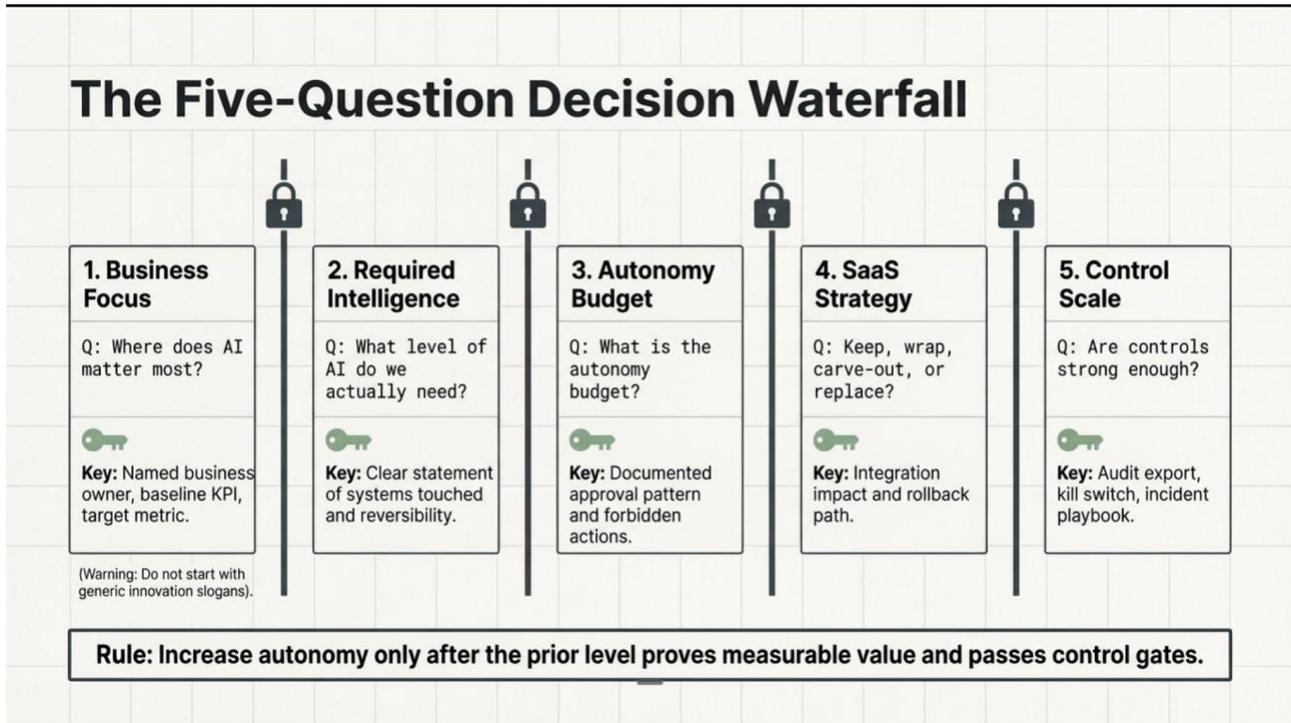
20	About the Authors	p. 20
21	Bibliography	p. 21

How to read this paper

- This is a decision tool, not a technical manual. It is designed to be read by a CEO and actioned in a board meeting.
- Two decisions drive everything here. First: where does AI actually move the needle for your business? Second: what must be in place before you give AI more autonomy? Answer both before committing budget.
- Default to bounded AI. Agents are earned through demonstrated governance, not assumed because the demos looked good.

One-Page Decision Framework

Run these five questions in order. Skipping one creates a decision you will regret funding. Business focus first, autonomy second, controls before scale. The sequence is not optional.



Board question	Decision choices	Minimum evidence required	What not to do
1. Where does AI matter most?	Select 2-4 value pools: productivity, cycle time, quality, capacity, or revenue support.	Named business owner, baseline KPI, target metric, and process scope.	Do not start with a generic innovation slogan or broad license rollout.
2. What level of AI do we actually need?	Assistive AI, bounded workflow AI, limited agents, or no-go for now.	Clear statement of systems touched, decisions influenced, and reversibility.	Do not jump to agents because demos look impressive.
3. What is the autonomy budget?	Read-only, recommend, execute with approval, or execute within guardrails.	Documented approval pattern, escalation path, and forbidden actions.	Do not allow money movement, legal commitments, pricing changes, or financial posting changes without multi-step approval.
4. Keep, wrap, carve-out, or replace SaaS?	Retain legacy SaaS, wrap with AI, carve out a sub-process, or replace over time.	Business case by process, integration impact, control implications, and rollback path.	Do not make one ideology-driven enterprise decision for every process.
5. Are controls strong enough to scale?	Go, pilot only, or hold.	Audit export, logging, replay tests, tool allowlist,	Do not fund scale if governance is still aspirational.

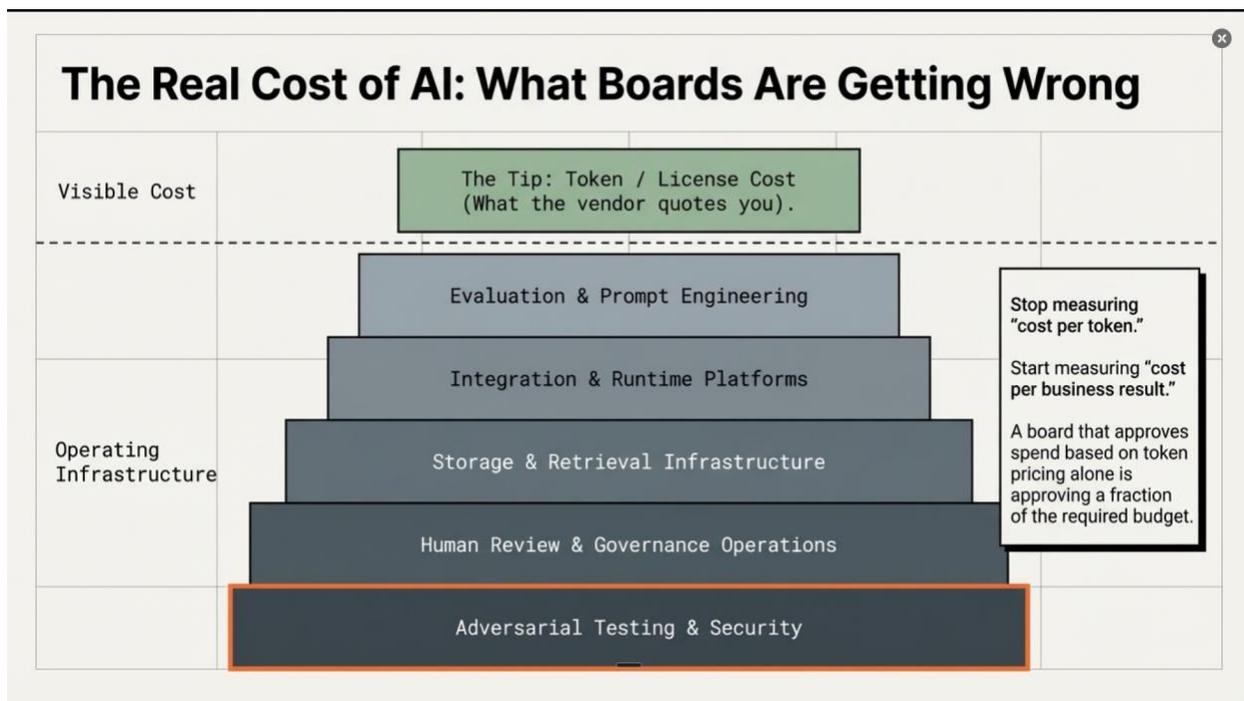
kill switch, owner, incident
playbook.

Board rule: increase autonomy only after the prior level proves measurable value and passes control gates.

Executive Summary

Most boards are asking the wrong question. It is not whether to adopt AI. It is whether your organization is built to govern what AI will do once you deploy it at scale.

The bottleneck is never the model. It is always the enterprise around it. Data discipline, clear ownership, tool controls, auditability, and approval structure are what separate organizations that capture value from those that multiply their exposure.



Since February, the stakes got higher and the decisions got more specific. Standards bodies moved. Regulators named obligations. Platforms shifted from selling model access to managing agent execution. Every one of those changes lands on the board agenda.

Buying licenses is not adopting AI. A company that adds AI tools without redesigning work, accountability, measurement, and controls has spent money without changing anything. That is the gap this paper is designed to close.

Board message: choose the business focus first. Do not let AI popularity become the strategy.

What Changed Since the February 2026 Edition

None of the changes since February made the front page. All of them affect decisions you need to make in the next 90 days.

1. **The governance bar just moved up.** The National Institute of Standards and Technology (NIST) launched the AI Agent Standards Initiative in February 2026, reinforcing that secure interoperability, trust, and standard agent behavior are moving toward formal enterprise requirements. This directly raises the bar for enterprise agent deployments: governance infrastructure will increasingly be assessed against these standards. See [The Control Stack](#) for the minimum requirements at each maturity level.
2. **Picking a vendor is now also picking an architecture.** Enterprise AI platforms are shifting from model access to managed agent execution. That increases the importance of permissions, shared context, state handling, and audit export. Selecting an AI vendor is now also a decision about shared context architecture and permissioning infrastructure. See [Vendor or Strategic Partner](#) for the full evaluation framework.
3. **Better models create better-looking failures.** Frontier model capability improved again in March 2026. Better performance raises value potential, but it also increases blast radius because business users are more likely to trust and connect these systems to production workflows. Better outputs also make failures harder to detect because they look more polished. See [What Cannot Be Optional](#) for the specific risk controls this capability improvement requires.
4. **The EU AI Act is no longer a future problem.** The European Union (EU) AI Act timeline is now operationally near. Boards with EU exposure should treat AI inventory, classification, logging, oversight, and documentation as 2026 work, not a future legal footnote. High-risk system obligations are now live, not theoretical. See [What Cannot Be Optional](#) for a board-level compliance path that covers 90-day, six-month, and 12-month targets.
5. **Speed without governance is not winning.** Enterprise surveys continue to show the same pattern: adoption is accelerating faster than formal oversight. That means governance maturity is becoming a competitive differentiator. The organizations capturing value are not the fastest movers. They are the ones pairing adoption speed with governance rigor. This paper is designed to help your board do both.
6. **The plumbing just got standardized.** Model Context Protocol (MCP) is emerging as a foundational standard for how AI agents connect to enterprise tools, databases, and services. Developed originally by Anthropic and now being adopted broadly, MCP defines a structured way for AI agents to discover and call external capabilities. For boards, MCP matters because it shifts interoperability from a proprietary integration problem to a standards-based one, reducing lock-in, simplifying audit of what tools agents can access, and enabling the tool allowlisting and permissioning controls that governance requires. Enterprises evaluating agent platforms in 2026 should ask explicitly whether their vendors support MCP and how tool permissions are managed within it.

Implication: governance, portability, and control evidence now belong in the same board conversation as productivity and cost.

Why AI Pays: The Business Case That Actually Holds

The value pools have not changed. Productivity, speed, quality, optionality. AI earns its keep fastest in information-dense, repetitive work: drafting, triage, knowledge retrieval, research synthesis, case assembly. If your team does a lot of that, the business case is immediate. If not, start somewhere else.

Stop leading with headcount elimination. It is the weakest argument in the room and the one most likely to create resistance. The stronger case is leverage: the same team moving faster, covering more ground, producing more consistent output. That is a growth story, not a reduction story.

A business case without a baseline metric is a pitch deck. Name the process. Name the owner. Set the before-and-after measurement. If you cannot do that, you are not ready to fund it.

Ownership: Where Most AI Programs Fail Quietly

AI programs do not fail loudly. They fail quietly, when nobody actually owns the outcome. The vendor does not. Technology cannot drive adoption without business workflow change. The business cannot own results without technical controls. You need a shared model with hard accountability lines, not a coordination committee.

Hard lines. Business owns outcomes, process redesign, and adoption. Technology owns architecture, integration, security, and observability. Risk and legal own policy and minimum control requirements. A single AI governance lead owns the inventory, review cadence, and the escalation path when something breaks. No shared ownership means no excuse for inaction.

When functions disagree on AI, and they will, the CEO or COO calls it. Without that forcing function, AI turns into a collection of disconnected experiments that collectively justify nothing and deliver less.

A board should insist that every scaled use case has one named business owner and one named technology owner.

Executive Support: The Air Cover That Makes or Breaks AI Programs

Your teams want the credit without the constraint. They want AI tools without being measured on what those tools actually produce. They want the upside of innovation without the discipline that makes innovation mean anything. Executive sponsorship exists to close that gap, by force if necessary.

Set the priorities. Approve the autonomy budget. Make workflow redesign non-negotiable. Then remove the friction that prevents it. The executive team's job is not to celebrate AI momentum from a distance. It is to create the conditions where momentum is actually possible.

What You Are Actually Funding When You Fund AI

The license is the small part. What organizations consistently underfund is everything that makes the license worth anything: evaluation, integration, adversarial testing, runtime engineering, governance operations. That is the gap between a working demo and a working production system.

Budget by maturity level, not by ambition. Assistive tools and workflow AI require product ownership, data engineering, applied AI capability, security review, and measurement. Agents add tool-gateway engineering, approval patterns, logging, and incident response. Anything more autonomous than that requires replay testing, rollback discipline, and governance operations as a permanent function, not a project.

- If the team measuring value is not the same team running the system, the measurement will eventually be fiction.
- Adversarial testing is not optional once agents touch systems of record. It is the price of deployment.
- A vendor can build the system. Only your people understand what risk it creates in your specific context.

Should You Hire More People or Increase Token Budget?

This is the defining operational question of 2026 for most enterprises that have moved past the pilot stage. It is also a question most boards are not yet framing explicitly, which means the decision is being made by default rather than by design.

The old model was simple: more work requires more people. The new model is more complex because AI extends the output capacity of existing workers before it eliminates the need for any of them. The relevant question is not “will AI replace my workforce” but “given a defined piece of work that needs to get done, what is the right combination of human capacity and AI capacity to execute it reliably and cost-effectively.”

There are three scenarios a board should think through explicitly. In the first scenario, the organization has significant unmet demand for a capability (faster research synthesis, more thorough compliance review, broader market monitoring) that currently cannot be resourced. Here the AI-first answer is almost always to extend token budget before adding headcount, because AI can close that gap faster and at a fraction of the cost, provided the task is structured enough for AI to handle with quality controls in place.

In the second scenario, the organization needs net new judgment, relationship management, or adaptive problem-solving that AI does not yet perform reliably. Here the hire is clearly justified, and the question becomes whether AI can be deployed to support that new hire so they function at higher leverage than a traditional hire would.

In the third scenario, the organization is considering a backfill after attrition. This is where the token budget question becomes most acute. Before posting the requisition, the board and executive team should ask: can we redesign the workflow so that AI handles the structured portions of this role, and either eliminate the backfill entirely or hire for a narrower, higher-judgment scope than the departing person occupied?

The financial discipline this requires is straightforward: model both options. For the hire: fully loaded cost per FTE, management overhead, time to productivity (typically three to six months), and fixed cost structure risk. For the AI investment: token cost at projected usage, integration and governance overhead, expected productivity multiplier based on task type, and the variable cost structure that allows scaling up and down without severance expense. The board that develops this model for one function owns a decision template that applies across the entire operating model.

One important caution: the “hire fewer people” argument is not the same as “invest less in talent.” The organizations that are extracting the most value from AI in 2026 are investing heavily in the talent that can direct AI effectively: prompt engineers, AI product managers, data curators, evaluation specialists, and governance leads. The workforce transformation is real, but it is a reshaping of what skills you need, not simply a reduction in headcount.

From Copilot to Agent: The Decision That Changes Your Risk Profile

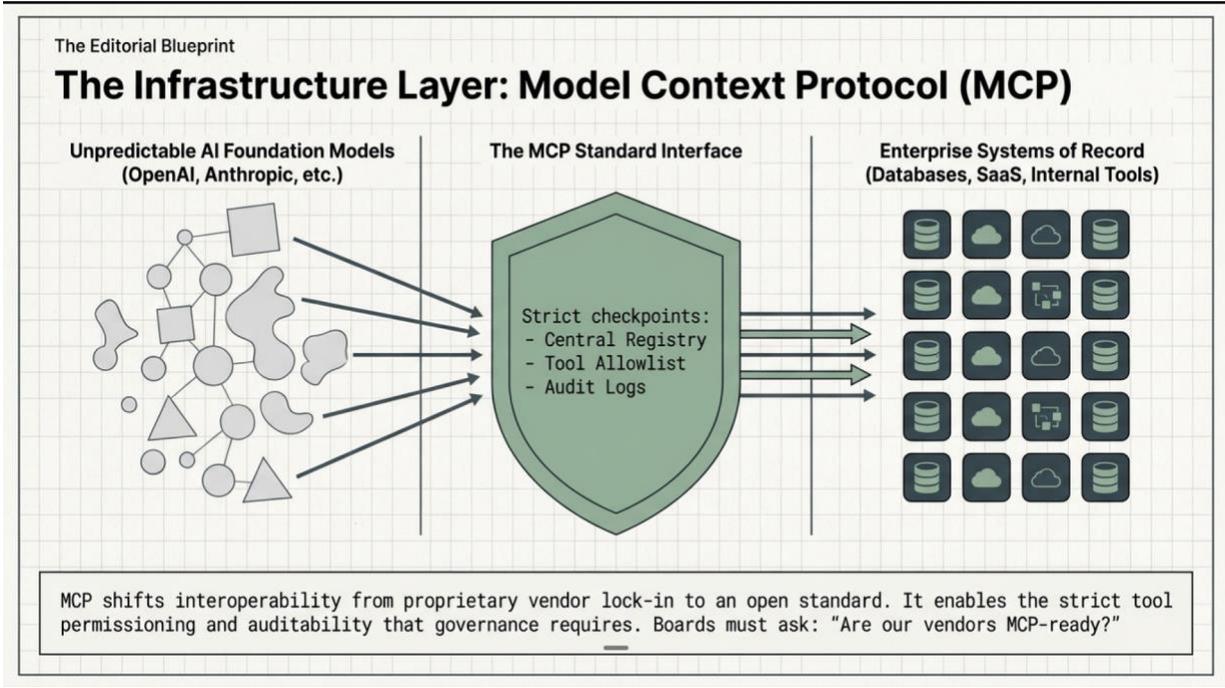
The move from copilot to agent is a risk decision, not a capability one. Make it only when the task is defined enough to bound the tools, make actions reversible, and set a clear approval checkpoint. Good starting candidates: case assembly, service-desk runbook execution, intake triage, procurement routing, data remediation proposals, internal operational routing.

The no-fly list has not changed. Anything that moves money, changes pricing, approves credit, touches financial postings, makes legal commitments, or talks to the outside world on material matters without a human checkpoint is not an agent use case yet. Not because the model cannot do it. Because your governance is not ready.

A model becoming more capable does not automatically make the use case safer.

Model Context Protocol: The New Infrastructure Layer

One of the most significant technical developments of early 2026 is the rapid adoption of the Model Context Protocol (MCP) as a standard interface for connecting AI agents to enterprise tools and data sources. MCP is an open protocol that allows AI models to discover available tools, query data systems, and take bounded actions through a structured and auditable connection layer. What makes MCP strategically important is not the protocol itself, but what it enables: a governed, extensible surface area for agent capabilities that can be inventoried, permissioned, and audited.



For board and executive audiences, MCP changes several decisions that were previously ambiguous. First, tool allowlisting becomes operationally tractable. Rather than building custom integrations for every agent capability, enterprises can maintain a central registry of approved MCP-compatible tools and restrict agent access to that registry. Second, vendor portability improves. An agent built against MCP-standard tool definitions is less coupled to a single model provider or orchestration platform. Third, audit and traceability improve because every tool call made through MCP is structured, logged, and inspectable by design.

The practical board question is not whether to adopt MCP, but whether enterprise AI vendors and internal platforms are building toward MCP compatibility. Boards should require their technology leads to report on MCP readiness as part of the quarterly AI ecosystem review. Enterprises that ignore MCP now risk building agent infrastructure that requires costly rework as the standard matures.

Vendor or Strategic Partner: A Different Decision Entirely

For most enterprises over the next 12-18 months, the AI decision is still largely a stack decision. The company is usually not building foundation models. It is selecting a combination of model provider, orchestration or runtime layer, retrieval layer, gateway pattern, and application controls.

Vendor evaluation now has to move beyond quality and price. Boards should ask about shared context, permissioning, agent state, audit export, model portability, evaluation tooling, and how much rework is required if the company needs to switch model providers or runtime platforms.

The distinction between a vendor and a strategic AI partner is not semantic. It reflects a fundamental difference in what the enterprise actually needs and the level of commitment required on both sides. A vendor sells a product or completes a bounded engagement, typically three to six months, with a defined deliverable and a clear exit. That model works for point solutions: a transcription tool, a summarization API, a compliance screening module. It does not work for AI transformation.

AI transformation is not a project. It is a continuous operating model shift that touches talent, data, workflow, governance, and culture simultaneously. The decision criteria for selecting a strategic AI partner should look more like a joint venture assessment than a procurement exercise. A strategic partner needs to understand your business model deeply, take

accountability for outcomes not just outputs, build institutional knowledge over time, and evolve the engagement as the technology and your organization change.

The board-level distinction looks like this: a vendor charges for hours or licenses and is done when the contract ends. A partner is measured against business outcomes, has skin in the long-term success of the program, and becomes harder to replace over time because the relationship carries embedded knowledge. Boards should be asking whether their AI engagements are structured as vendor relationships or partner relationships, and whether the incentive structures actually match the complexity of what they are trying to build.

Board-level vendor scorecard

Non-negotiables should include data usage and retention transparency, exportable audit logs, permissions and system boundary controls, regression and replay testing support, cost transparency under realistic workloads, and a credible portability story.

Criterion	What the board should ask
Audit export	Can we export prompts, responses, tool calls, approvals, and agent actions in a usable format?
Permissioning	How are access limits, system boundaries, and shared context governed?
Evaluation	Can we run replay tests, adversarial tests, and controlled regressions before rollout?
Portability	How much application logic breaks if we change model or platform vendors?
Cost	What is the all-in cost per business outcome, not just token or seat price?

The Control Stack: What Must Be in Place Before You Scale

Governance infrastructure is what determines whether your AI investment scales or collapses under its own weight. For assistive tools and workflow AI, the floor is: data ownership, access control, vendor review, logging, and basic evaluation. For agents, the floor goes up sharply: tool allowlisting, approval gates, schema validation, full audit logs, kill switches, incident response, and end-to-end behavior testing. Below the floor, do not launch.

As autonomy increases, the infrastructure compounds. Highly autonomous systems need an agent registry, versioned policy artifacts, replay testing, drift detection, rollback discipline, and board-level sign-off for any material change in agent behavior. The NIST AI Agent Standards Initiative provides the specific framework your technology team should be benchmarking against.

- Know what is running. Every scaled use case gets named, owned, and logged before anything else.
- Autonomy does not increase by default. Someone specific signs off every time it does.
- Read-only and write-access are not the same risk. Treat them differently from day one.
- No rollback path, no launch. That is not a suggestion.

What Cannot Be Optional: Your Compliance and Control Baseline

Five principles, no exceptions: least privilege, human oversight, full traceability, test before deploy, rapid rollback. They apply to copilots, workflow AI, and agents alike. Complexity changes how you implement them. It does not change whether they apply.

The threats are not theoretical. Prompt injection, tool abuse, data exfiltration, hallucinated justification, uncontrolled agency have all appeared in enterprise deployments. Better models do not eliminate these risks. They make failures harder to catch because the output looks more convincing.

If you have EU exposure, 2026 is not a planning year. It is an execution year. AI inventory, classification, oversight, logging, and documentation are this year's work. The board should be receiving readiness updates now, not at year end.

Understanding the Regulatory Landscape: Who Sets the Rules and Why It Matters Now

Two regulatory frameworks dominate the enterprise AI compliance discussion in 2026. The first is the European Union AI Act, which entered into force in August 2024 and is now rolling into active obligations by risk tier. The second is the standards infrastructure being built by the National Institute of Standards and Technology (NIST) in the United States, which, while not mandatory, is rapidly becoming the baseline that regulators, insurers, and large enterprise buyers reference when assessing AI program maturity.

The EU AI Act classifies AI systems into four risk categories: unacceptable risk (banned), high risk (heavily regulated), limited risk (transparency requirements), and minimal risk (largely unregulated). High-risk systems include AI used in credit scoring, employment decisions, access to essential services, critical infrastructure, and law enforcement. If your enterprise operates in the EU and uses AI in any of these domains, you are in active compliance territory now. The obligations for high-risk systems include mandatory risk management documentation, data governance requirements, human oversight mechanisms, accuracy and robustness standards, and registration in the EU AI database. Non-compliance carries fines of up to 3% of global annual turnover for violations of obligations, and up to 6% for prohibited AI use cases.

NIST published its AI Risk Management Framework (AI RMF 1.0) in January 2023 and has continued releasing profiles including the Generative AI Profile and, most recently in February 2026, the AI Agent Standards Initiative. NIST frameworks are voluntary in the US context, but they function as the de facto governance baseline. The AI RMF is organized around four functions: GOVERN, MAP, MEASURE, and MANAGE. Boards should treat these not as a compliance checklist but as

a structural diagnostic: does your organization have governance structures, can you map where AI operates, can you measure its performance and risks, and do you have processes to manage identified issues?

For CEOs, boards, and executive teams, the practical compliance question is not whether you have read the frameworks. It is whether you have operationalized them. The gap between stated intent and operational reality is exactly where regulators and auditors will look first. A compliance-ready AI program requires an AI inventory (knowing what is deployed and where), a risk classification for each system, documented ownership and oversight for high-risk applications, logging and audit exports that are exportable and reviewable, documented human oversight checkpoints for consequential decisions, and a tested incident response process.

Getting into compliance is a phased exercise. In the near term (next 90 days), the board should insist on a complete AI inventory and risk classification. In the medium term (next six months), organizations should close the largest documentation gaps on high-risk systems and confirm that logging and audit export infrastructure is in place. Over the following 12 months, the goal is a standing governance process rather than a one-time compliance effort, with quarterly reviews of the AI use case portfolio, escalation processes for new deployments, and board-level reporting on compliance posture. The regulation will evolve, but the governance infrastructure you build now will be reusable as requirements mature.

Minimum board control package for production agents: registered use case, named owner, tool allowlist, approval gate, audit export, replay tests, incident playbook, kill switch, and rollback path.

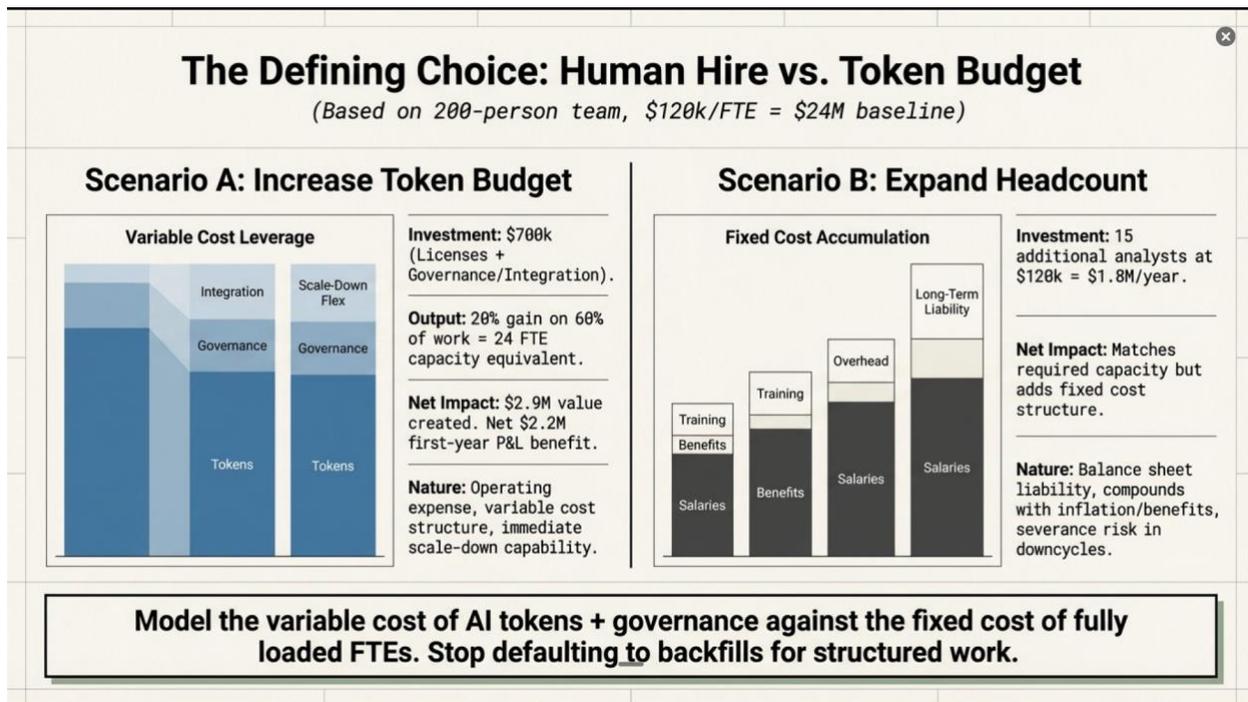
The Real Cost of AI: What Boards Are Getting Wrong

Token cost is what your vendor quotes you. It is also the smallest part of your actual investment. Runtime platforms, retrieval, storage, integration, evaluation, monitoring, human review, security engineering, and organizational change are the rest of the bill. A board that approves AI spend based on token pricing is approving based on the wrong number.

In March 2026, better models reduced some costs while new platform features added others. The net is not automatically positive. Stop measuring cost per token. Start measuring cost per decision, per case, or per business result. That is the number that tells you whether the investment is working.

Illustrative Financial Example: Making the Cost Decision Concrete

The following is an illustrative scenario to frame the board-level cost conversation. Assume a financial services firm with 200 knowledge workers in its research and compliance function, each earning an average fully loaded cost of \$120,000 per year. Total annual cost for that team: \$24 million.



Scenario A: The firm deploys an enterprise AI platform at \$2,000 per user per year (\$400,000 total in licenses), plus \$300,000 in integration, evaluation, and governance infrastructure, for a total first-year AI investment of \$700,000. Independent productivity research on similar deployments (Brynjolfsson et al., Noy and Zhang) suggests 20 to 40 percent efficiency improvements in drafting, research synthesis, and compliance review tasks. At a conservative 20 percent productivity gain applied to 60 percent of the team’s work, the effective capacity increase is equivalent to approximately 24 additional FTEs, or \$2.9 million in annual labor value created. Net first-year P&L impact: approximately \$2.2 million benefit after investment cost. The AI spend does not appear as headcount on the balance sheet. It is an operating expense that creates recurring capacity leverage.

Scenario B: The firm instead hires 15 additional analysts at \$120,000 fully loaded to address the same capacity gap. Total additional headcount cost: \$1.8 million per year, recurring and compounding with inflation and benefits. The headcount shows up on the balance sheet, creates fixed cost structure, and requires management overhead. There is no flexibility to scale down in a downcycle without severance expense.

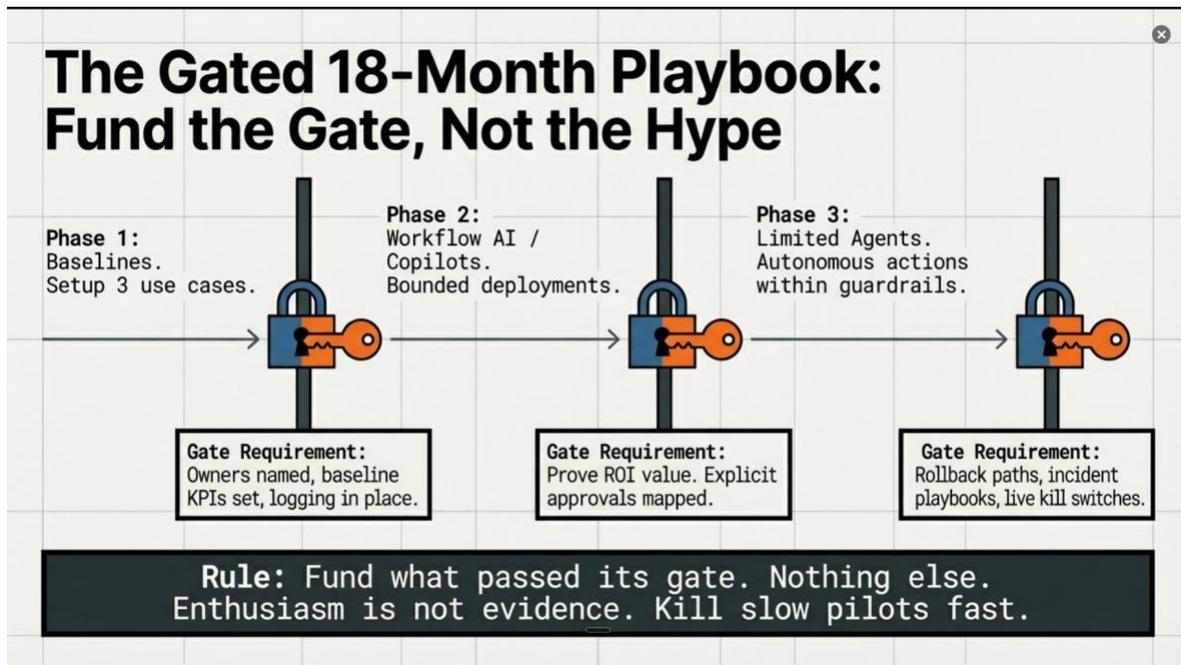
The comparison is not that AI always wins. It is that the decision needs to be modeled explicitly at the board level using real numbers. The key variables to model are: what is the fully loaded cost per FTE in the affected function, what is the

realistic productivity improvement percentage for the specific task type, what percentage of total work hours is affected, and what is the total cost of the AI program including governance infrastructure, not just licenses. Boards that approve AI spend without this model are operating on intuition rather than decision support. Boards that build this model for one function have a template they can apply across the portfolio.

A further consideration for the balance sheet and cash flow: AI infrastructure cost is largely variable (token consumption scales with usage) rather than fixed, which changes the risk profile of the investment compared to headcount. In a downcycle, AI spend can be dialed back without severance cost. In a growth cycle, AI capacity can be increased immediately without a six-month hiring cycle. That optionality has real financial value that rarely appears in the initial cost comparison.

Your 18-Month Playbook: Fund the Gate, Not the Hype

Phase 1: three use cases, owners named, baselines set, logging and evaluation in place before a single workflow goes live. Phase 2: bounded copilots or workflow AI with explicit ownership and approvals. Prove value before moving. Phase 3: limited agents only in environments where tools, approvals, rollback, and monitoring are already operational. Adaptive memory and learning loops come after that, and only in narrow, well-governed domains.



Fund what passed its gate. Nothing else. Enthusiasm is not evidence. Each move up the autonomy ladder requires proof that the prior level delivered measurable value and cleared its control checks. No proof, no funding.

- Put governance on its own line in the budget. Buried inside platform spend, it disappears.
- Your vendor will always say you are ready for the next level. Your own evidence makes that call.
- Kill slow pilots fast. A weak pilot running for twelve months teaches nothing and burns credibility.

Private Equity and AI Labs: A New Enterprise Adoption Channel

A development that is reshaping enterprise AI adoption in early 2026 deserves direct board attention: the emergence of formal partnerships between frontier AI laboratories and private equity firms, specifically designed to accelerate AI deployment across portfolio companies. OpenAI, Anthropic, and others have begun structured relationships with major

PE sponsors to offer enterprise AI programs at scale to portfolio companies, often bundled with adoption support, governance frameworks, and implementation resources.

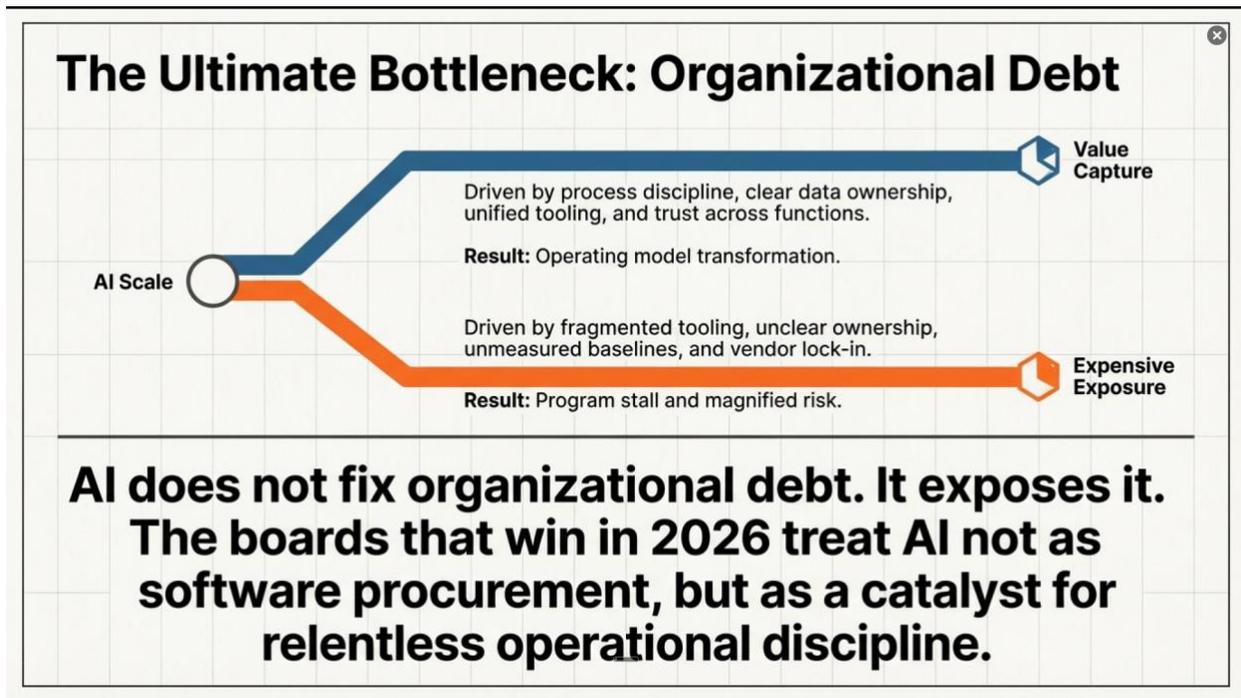
The strategic logic from the AI lab perspective is straightforward. Enterprise sales cycles are long, and the bottleneck is not model capability but organizational readiness. PE-backed companies offer a concentrated channel: a single sponsor relationship can unlock dozens of portfolio companies, each with a motivated operating partner and board pushing for AI adoption as a value creation lever. For the PE sponsor, the partnership offers a proprietary advantage in portfolio value creation and a differentiated narrative for fundraising.

For boards and executive teams at PE-backed companies, this development has several practical implications. First, if your sponsor has an AI lab partnership, the terms of that relationship may create economic incentives to favor specific platforms over others, regardless of which is the best fit for your specific workflows. Boards should ask sponsors to disclose the structure of any AI lab arrangements and ensure that vendor selection remains a technology-fit decision, not a channel economics decision. Second, the availability of structured adoption programs through PE channels can genuinely accelerate implementation timelines and reduce governance setup costs, particularly for companies that lack internal AI capability. Third, this channel does not eliminate the need for independent judgment on what AI should do in your specific business. No sponsor-level agreement substitutes for the board-level decisions on autonomy, controls, and ownership described in this paper.

Independent of the PE channel, the broader pattern matters for any enterprise board: AI adoption is now being structured as a value creation program, not just a technology deployment. That framing is correct. The board that treats AI as a software procurement will underinvest in the organizational change required to capture value. The board that treats AI as an operating model transformation, managed with the same discipline as any other major capital allocation decision, is far better positioned to separate signal from noise in what is, by any measure, one of the most consequential technology transitions in a generation.

Why Most AI Programs Stall: The Organizational Debt Few Boards Are Addressing

The model is not the problem. It has never been the problem. What stops AI programs is the organizational debt that was there before you started: unclear data ownership, undefined processes, fragmented tooling, no measurement discipline, and functions that do not trust each other. AI does not fix organizational debt. It exposes it.



The gap is wider now because the technology moved faster than the organizations running it. Process discipline and governance maturity, not model quality, determine whether this ends in value creation or expensive exposure. That is not a technology problem. It is a leadership problem, and it belongs on the board agenda.

Board Toolkit: Frameworks You Can Use Tomorrow

Use these tools to make better decisions faster. Each is designed to be completed by your team, not by a consultant, and to force the specificity that separates a real AI program from a funded experiment.

- Decision one-pager: four decisions, one page, one meeting. Business focus, autonomy budget, SaaS displacement dial, and gating controls. If you cannot complete it, the program is not ready.
- Use-case intake sheet: six questions every use case must answer before it gets funded. No answers, no budget.
- Vendor scorecard: six non-negotiables. A vendor that cannot answer all six is not ready to be your partner.
- Agent charter: the contract between your governance team and every agent you deploy. No charter, no launch.
- Board reporting pack: five numbers that tell you whether your AI program is real or theater.

About the Authors

Tunde Kehinde is the Managing Partner of 5054 Holdings, where he invests in and builds businesses across technology and infrastructure. He co-founded Jumia, helping scale the business through its formative growth years and contributing to over \$1 billion in enterprise value by designing go-to-market strategy, operating cadence, and trust mechanisms in complex, low-trust markets prior to its public listing on the NYSE. He subsequently led growth and investment strategy at Lidya, a cross-border fintech lender, where he originated and deployed over \$150 million in private credit across the U.S., Europe, and Africa, achieving a 99% repayment rate through rigorous underwriting, dynamic pricing, and disciplined borrower acquisition. Across his career, Tunde has led over 1,500 direct and indirect reports across multiple continents, installing the operating rhythm, ownership structures, and performance standards required to sustain growth at scale. He is a Guest Lecturer at Howard University School of Business and serves as an advisor to enterprises navigating AI adoption, operating model transformation, and capital allocation strategy. Tunde completed his MBA at Harvard Business School.

Cyril Simone is a technology executive and researcher with over 20 years of experience spanning AI, machine learning, and enterprise transformation. He has held senior technology and data science leadership roles across media, consulting, and enterprise transformation, and currently leads work in advanced AI and computing architectures while founding iLiner Studios and Mercado AI.

Contact and Additional Frameworks

The authors maintain a library of additional frameworks, diagnostic tools, and decision templates that support the material in this paper, including a full AI readiness assessment, a vendor-versus-partner evaluation rubric, a headcount-versus-token-budget modeling worksheet, and a board reporting pack template. These are available on request.

To reach Tunde Kehinde directly: tunde@athenai.com

To reach Cyril Simone directly: cyril@mercadoai.com

If this paper has been useful to your board or executive team, the authors welcome feedback on which frameworks would be most valuable to publish in future editions.

Bibliography

References verified for this March 2026 edition. These sources support the standards, governance, regulation, security, and enterprise adoption updates incorporated into the body of the paper.

Governance, standards, and regulation

- NIST. Artificial Intelligence Risk Management Framework (AI RMF 1.0). <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>
- NIST. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- ISO/IEC 42001:2023. Artificial intelligence management systems. <https://www.iso.org/standard/42001.html>
- European Commission. AI Act. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- NIST. Announcing the AI Agent Standards Initiative for Interoperable and Secure Innovation. <https://www.nist.gov/news-events/news/2026/02/announcing-ai-agent-standards-initiative-interoperable-and-secure>

Enterprise platforms and model developments

- OpenAI. Introducing OpenAI Frontier. <https://openai.com/index/introducing-openai-frontier/>
- OpenAI. Introducing GPT-5.4. <https://openai.com/index/introducing-gpt-5-4/>
- OpenAI. GPT-5.4 Thinking System Card. <https://openai.com/index/gpt-5-4-thinking-system-card/>
- OpenAI. Improving instruction hierarchy in frontier LLMs. <https://openai.com/index/instruction-hierarchy-challenge/>

Enterprise adoption, value, and security

- Brynjolfsson, Li, and Raymond (2023). Generative AI at Work. NBER Working Paper 31161. <https://www.nber.org/papers/w31161>
- Noy and Zhang (2023). Experimental evidence on the productivity effects of generative AI. Science. <https://www.science.org/doi/10.1126/science.adh2586>
- Peng et al. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot. <https://arxiv.org/abs/2302.06590>
- Deloitte. 2026 State of AI in the Enterprise / press release. <https://www.deloitte.com/us/en/about/press-room/state-of-ai-report-2026.html>
- EY US. Autonomous AI adoption surges as oversight falls behind. https://www.ey.com/en_us/newsroom/2026/03/ey-survey-autonomous-ai-adoption-surges-at-tech-companies-as-oversight-falls-behind
- OWASP. Top 10 for Large Language Model Applications. <https://owasp.org/www-project-top-10-for-large-language-model-applications/>
- OWASP Cheat Sheet Series. AI Agent Security Cheat Sheet. https://cheatsheetseries.owasp.org/cheatsheets/AI_Agent_Security_Cheat_Sheet.html
- Greshake et al. (2023). More than you've asked for: indirect prompt injection. <https://arxiv.org/abs/2302.12173>