TIRA STRATEGIC ADVISORY · RESEARCH

The Conviction Engine

Closing the Strategic Latency Gap in Enterprise AI Transformation

Mahesh M. Thakur

TIRA Strategic Advisory

July 2026

Abstract

Enterprise AI investment has reached a scale without precedent, and the financial return on that investment remains largely undemonstrated. This paper argues that the shortfall originates in governance rather than in technology. It introduces the Strategic Latency Gap, the widening distance between the exponential advance of AI capability and the linear rate at which organizations absorb that capability into decisions, workflows, and P&L. Drawing on two decades of operating experience inside Microsoft, Amazon, Intuit, and GoDaddy, and on advisory engagements with Fortune 100 leadership teams, the paper presents the Conviction Engine, an operating discipline that aligns capital allocation, product strategy, and execution governance on a single quarterly cadence. The output of that alignment is fiduciary conviction: the capacity of a leader to state, in front of a board, what the capital bought and what it returned. Recent survey and market evidence is examined and its limitations noted.

Keywords: *enterprise AI transformation, AI ROI, fiduciary governance, capital allocation, agentic workflows, executive decision-making, organizational velocity, board oversight*

1. Introduction: the gap that does not show up on a server

I was advising the chief executive of a Fortune 100 company when I first saw the problem clearly. I had been invited to observe one of their senior leadership meetings, the room where the company's largest decisions are made. The technology was not the issue. The large language models worked. The budget was already committed. What I watched instead was a slower kind of failure: the distance between what the organization was capable of deciding and what it actually decided. I had no name for it that afternoon. The symptom was visible long before the label was, a lack of alignment, clarity, and momentum at the C-Suite and board level that no tool on the market could close.

I now call that distance the Strategic Latency Gap. It is the lag between how fast AI capability advances and how fast an organization absorbs that capability into decisions, workflows, and P&L. The capability curve is nearly vertical. The absorption curve is close to flat. The space between the two is where most enterprise AI budgets quietly disappear.

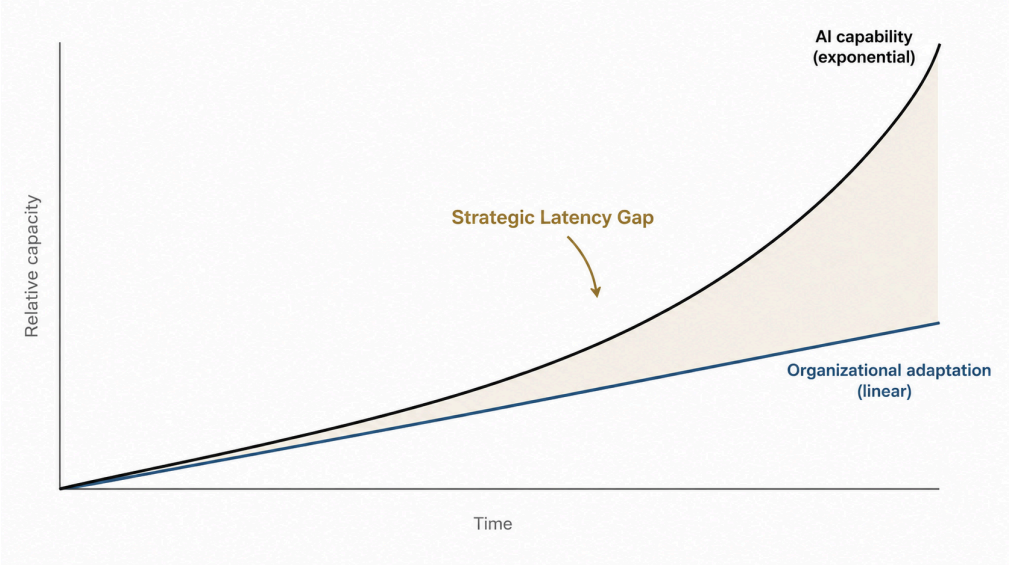

Figure 1. The Strategic Latency Gap. Capability compounds. Organizational absorption advances in a straight line. The shaded area is where committed capital produces no measurable return.

This paper argues that the gap is a governance problem wearing a technology costume, and that closing it requires a specific operating discipline rather than a better tool. I call that discipline the Conviction Engine.

2. The AI ROI Mirage

Boards are approving AI budgets at a pace no prior technology has matched. What most of them are not seeing is the return. The capital is committed, the pilots launch, the dashboards get built, and twelve months later the P&L reads the same as it did before. This is the AI ROI Mirage: the appearance of enterprise transformation without its financial substance.

The scale of the shortfall is now measurable. PwC's 29th Global CEO Survey, covering 4,454 chief executives across 95 countries, found that 56 percent report neither increased revenue nor reduced costs from AI over the preceding twelve months. Thirty percent report revenue gains, 26 percent report cost reductions, and only 12 percent report both. The pattern holds at the largest scale of spending. When Alphabet raised its 2026 capital expenditure guidance to between $195 billion and $205 billion in July 2026, its shares fell more than five percent. An analyst on the call asked the chief executive directly about “the size of the overall gen AI ROIC opportunity and the timing” of that return. A question once asked quietly inside boardrooms is now being asked aloud on earnings calls.

The mirage survives because most organizations measure AI by activity rather than by outcome. Number of pilots running, number of tools deployed, number of employees holding a license. None of those is a figure a board can defend to shareholders, because none of them connects to capital returned. A company can score well on every one of those measures and still have moved nothing on the only statement that matters.

The reason boards accept activity metrics is that AI arrived framed as an IT concern, and IT has always reported in the language of deployment rather than the language of return. A fiduciary body cannot govern what it cannot measure in its own terms. Until AI investment is expressed as capital allocated against capital returned, the board is approving spend it has no instrument to evaluate.

3. Why the failure is fiduciary rather than technical

MIT's Project NANDA, in its 2025 report on the state of AI in business, found that enterprise AI initiatives purchased from external vendors succeeded roughly twice as often as those built internally, about 67 percent against 33 percent. The report is not peer reviewed and its methodology has been contested, so it reads as a signal rather than a settled fact. The signal still points somewhere useful. What separated the two groups was the discipline of the organization deploying the capability. The most capable large language model did not rescue a company that lacked that discipline.

Alexander Karp, the chief executive of Palantir, put it more bluntly in July 2026, saying that AI models “have been completely irresponsibly oversold” and arguing that durable enterprise value sits in the combination of model, application layer, and compute. His framing supports a broader point. The organizations that capture return from AI are the ones that built the surrounding operating system to hold it. Capability without that system leaks value rather than compounding it.

I watched the internal version of this years earlier, inside Microsoft. I was on a team reviewing the integration of machine learning models across two organizations ahead of a proposed acquisition valued in the billions, one that would have ranked among the largest in the company's history. The models were strong on both sides. The question I raised with the team and the leadership went past whether the technology worked. I asked whether the combined capability created any durable competitive advantage, or whether we were about to pay a premium for something a competitor could assemble from the same public components within a year.

The room went quiet. I got some answers. Those answers filled the silence without addressing the question. Capability against durable advantage is the question most AI budgets never ask, and the silence that follows it is diagnostic.

4. The Conviction Engine

The Conviction Engine™ is the operating discipline that closes the Strategic Latency Gap. It aligns three systems that most organizations run separately, and it runs them on a single cadence. The three are Capital Allocation, Product Strategy, and Execution Governance. The output of the three working together is fiduciary conviction, the state in which a leader can stand in front of a board and defend an AI investment decision in the language the board actually speaks.

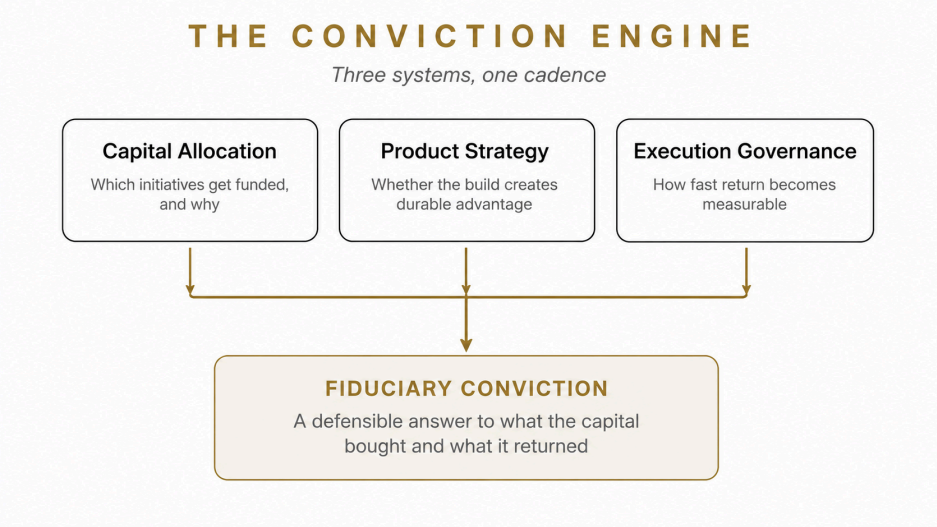

Figure 2. The Conviction Engine. Three systems run on one cadence, producing a single defensible output.

4.1 Capital Allocation

This pillar governs which initiatives get funded and why, rather than how large the budget is. The distinction matters, because most organizations treat the AI budget as a number to be set once a year and then defended. The discipline here concerns the allocation decision itself: which small number of initiatives can move EBITDA, and which experiments should stop so that those initiatives reach the scale where return becomes visible. Funding spread thin across a portfolio of pilots guarantees that none of them arrives anywhere a board can see.

The most common capital mistake I see in boardrooms is a contradiction the board does not notice it is making. A leader is handed a hundred million dollars of AI funding and, in the same conversation, asked to remove a quarter of the engineering organization that would have to build against it. The board has funded the ambition and defunded the capacity to deliver it in a single motion. This is the Strategic Latency Gap created deliberately, by the very body that later asks why the return never arrived.

This contradiction is now visible in public. In January 2026, Teradata's chief executive told roughly 5,100 employees that the company would fund its AI investment by “reallocating the budget from 2026 annual salary adjustments,” a decision reported publicly the following June. The strategic logic is defensible on its face. The organizational consequence is the part that rarely enters the capital discussion: an AI transformation depends on the same people whose compensation funded it choosing to change how they work.

4.2 Product Strategy

The second pillar governs what the capital actually builds. The move is from bolted-on AI features to an Agentic-First business model, where autonomous systems are designed into how the company operates rather than attached to the edges of existing products. A feature can be copied by any competitor with access to the same model. An operating model rebuilt around agentic workflows is far harder to replicate, because the advantage lives in the organization's design rather than in the tool.

The cost of getting this pillar wrong is concrete. Before working with me, one client had spent twelve months and tens of millions of dollars attempting to build an AI capability internally from scratch, and the effort

collapsed. The sunk cost was the smaller loss. The larger one was market position, because while they were building, their competitors moved ahead in the agentic race and did not wait. They treated a strategic capability as an engineering pilot project rather than a governed investment.

This is where the MIT signal earns its place in the argument. The organizations that failed at internal builds did not fail because building is inherently wrong. They failed because they lacked the operating discipline to convert capability into advantage, and an internal build exposes that absence faster than a vendor purchase does. The pillar is a requirement to know, before the capital moves, whether buying or building produces a durable advantage in the specific case at hand.

4.3 Execution Governance

The third pillar sets the cadence. The Conviction Engine replaces the twelve-month roadmap with 90-day Conviction Cycles, each one moving an initiative from concept toward measurable P&L impact inside a single quarter. The twelve-month roadmap is where Strategic Latency compounds, because it defers the moment of measurement past the point where a course correction is cheap. A ninety-day cycle forces the organization to meet the goal four times a year rather than once, and it holds technical velocity to a fiduciary standard rather than an engineering one.

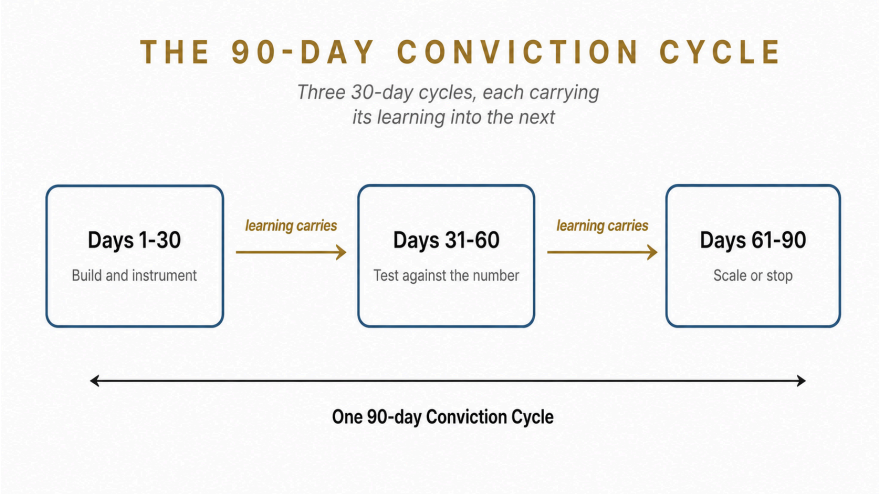

Figure 3. The 90-day Conviction Cycle, composed of three 30-day cycles. Learning from each cycle is carried forward rather than archived.

A ninety-day cycle only works if it is subdivided. In practice the quarter breaks into three 30-day cycles: one to build and instrument, one to test the build against the number it was funded to move, and one to scale it or stop it. The compression is deliberate. A thirty-day boundary is short enough that a team cannot defer the uncomfortable measurement to a later phase.

The mechanism that makes this compound is the transfer of learning between cycles rather than the speed of any single one. One company I work with runs weekly and bi-weekly sprints to build new products while testing AI tools that augment their engineers. What they learned in one sprint about where the tools helped, and where the tools quietly cost them time, is written into the design of the next. Across a quarter that produces three cycles of accumulated organizational knowledge rather than one delayed verdict, and the

knowledge travels beyond the team that generated it. Most organizations run the sprints and discard the learning, which is how a company can move quickly for a year and arrive nowhere.

5. The board-defensible number

The purpose of running the three pillars on one cadence is to produce a single output the previous approach never could: a number a fiduciary can defend. Capital Allocation establishes what was invested and why that initiative rather than another. Product Strategy establishes whether the investment bought a durable advantage or a copyable feature. Execution Governance establishes the return inside a quarter rather than a year. Together they convert AI from a cost center the board tolerates into a revenue driver the board can govern.

Fiduciary conviction is a governance capacity: the ability to answer, in front of the board, what the capital bought and what it returned, with evidence rather than activity. The Strategic Latency Gap closes at the point where a leader can answer that question without reaching for a deployment metric. That is the difference between an organization that spends on AI and one that compounds on it.

If you want to see where your own organization sits on the Strategic Latency Gap, the Fiduciary Audit maps your current AI investment against the three pillars above: maheshmthakur.typeform.com/FiduciaryAudit

---

References

[1] PwC. 29th Global CEO Survey: Leading Through Uncertainty in the Age of AI. Davos, January 2026. pwc.com

[2] Alphabet Inc. Second Quarter 2026 Earnings Call, 22 July 2026. Coverage and analyst exchange reported by CNBC. cnbc.com

[3] Challapally, Pease, Raskar and Chari. The GenAI Divide: State of AI in Business 2025. MIT Project NANDA, July 2025. Not peer reviewed; methodology contested. nanda.media.mit.edu

[4] Karp, Alexander. Interview, CNBC, 1 July 2026. Analysis in Forbes, 2 July 2026. forbes.com

[5] Teradata internal memo, January 2026, first reported by Business Insider and subsequently by Fortune, June 2026. fortune.com

Appendix A. Evidence summary

The table below consolidates the external evidence cited in this paper, with scope and limitations stated so that each figure can be weighed rather than assumed.

Source	Finding	Sample / scope	Date
PwC 29th Global CEO Survey	56% report neither revenue gain nor cost reduction from AI. 12% report both.	4,454 CEOs, 95 countries	Jan 2026
MIT Project NANDA	Externally purchased AI initiatives succeed roughly 67% of the time against 33% for internal builds.	Enterprise survey, not peer reviewed	Jul 2025

Alphabet Q2 2026 earnings	Capex guidance raised to $195B–$205B. Shares fell over 5%. Analysts questioned ROIC size and timing.	Public earnings call	Jul 2026
Teradata internal memo	2026 salary adjustments reallocated to fund AI investment across roughly 5,100 employees.	Single company	Jan 2026, reported Jun 2026

Note: the MIT Project NANDA figure is the least settled of the four. It has not been peer reviewed, and independent researchers have publicly requested release of its underlying data. It is presented here as directional evidence rather than as a confirmed measurement.

Appendix B. Fiduciary readiness: six diagnostic questions

A board or executive team can preview its own position on the Strategic Latency Gap by answering the following. These mirror the dimensions assessed in the full Fiduciary Audit.

Capital Allocation

- How would you describe your organization's current capital posture on AI: experimental, committed, or defensive?
- How often does your governance body actually review AI capital allocation against return, not deployment?

Product Strategy

- Does your AI investment build a moat a competitor with the same models could not replicate in twelve months, or a feature they could?

Execution Governance

- How many times per year does an AI initiative meet a financial number rather than a status update?
- Is the velocity of the people operating the system keeping pace with the velocity of the system itself?

These questions are a preview of the same dimensions assessed in the full Fiduciary Audit, which produces a scored alignment result and a specific next step for your organization:
maheshmthakur.typeform.com/FiduciaryAudit

About the Author

Mahesh M. Thakur coaches and advises senior leaders navigating AI-era transformation across Silicon Valley and Fortune 500 companies, including VPs, CIOs, and Executive Leadership Teams from organizations such as Meta, Google, NVIDIA, and Apple. A former leader at Microsoft, Amazon, Intuit, and GoDaddy, he is recognized as the world's #1 executive coach helping successful leaders remain impactful in the AI era. He is a member of the elite 100 Coaches organization and one of only 15 Master Certified Coaches globally in Stakeholder Centered Coaching.