

WHITE PAPER

The Four Gaps: Where AI Investment Stalls

Enterprises have committed the capital. Most are not yet seeing the return. A field guide to the four organizational gaps behind AI underperformance: what the evidence shows, and what it takes to close each one.

**DR. IAN
MCCULLOH**

Founding Partner

**MARIA
CHALOUX**

Founding Partner

OCTANTADVISORY.COM

EXECUTIVE SUMMARY

Four gaps stand between AI investment and AI performance

The enterprise has committed to AI. The performance it was promised has largely not arrived, and the evidence now shows why. The technology mostly works. The failure sits in the organization around it.

The numbers are stark. MIT's 2025 study of enterprise AI found that roughly 95% of generative-AI pilots deliver no measurable impact on the bottom line. McKinsey reports that more than 80% of organizations see no tangible enterprise-level earnings impact from their AI use. BCG finds that only 26% of companies have built the capability to move AI beyond proofs of concept into value. The most-cited diagnosis is consistent across all three: the models are good enough. The constraint is organizational integration, leadership, and governance.

Across our diagnostic work and conversations with more than 100 C-suite leaders, AI executives, and board directors, that organizational failure resolves into four recurring gaps: the mandate gap, the leadership gap, the execution gap, and the performance gap. They are not mutually exclusive (an organization can carry more than one), but each is distinct, each has a recognizable signature in the data and in the room, and each closes through a specific move.

Two of the four are gates. Until the mandate is funded and the leadership is in place, closing the others produces local wins that do not hold. The order in which the gaps are addressed determines whether the investment compounds, and the stakes are not symmetrical. BCG finds that AI leaders, the organizations that close these gaps, have achieved 1.5 times the revenue growth and 1.6 times the shareholder returns of their peers. This paper takes each gap in turn, and then sets the sequence.

THE PARADOX

Investment is surging. Returns are not.

Ninety-two percent of companies plan to increase their AI investment over the next three years. Only 1% describe themselves as AI-mature. That distance between the capital committed and the capability built is the most consequential and least examined risk on the enterprise AI agenda, and it is widening rather than closing.

95%

of enterprise generative-AI pilots deliver no measurable impact on the P&L.

It is easy to call this a technology problem and reach for better models or more data. It isn't one. MIT traces the failure to a "learning gap" in how organizations absorb AI, not to the models. Market analysts land in the same place: what moves earnings is redesigning how the work actually gets done, and having a CEO who owns AI governance. The technology already works in most companies. What they have not built is an organization that can turn it into results.

This is the premise of the four gaps. Each one names a specific place where AI investment leaks out of the system before it becomes enterprise value, and each is visible in both the external evidence and the conversations we have with leadership teams. The sections that follow take them one at a time.

THE FOUR GAPS

One pattern, four signatures

When we sit with a leadership team, the same four gaps keep surfacing. Each one shows up as a phrase you can hear around the table, and the outside data backs it up.

THE GAP	WHAT'S MISSING	THE SIGNAL YOU'LL HEAR
The Mandate Gap	A funded, board-owned ambition tied to named outcomes	"Everyone agrees AI matters, but ask us separately and you'll hear different priorities."
The Leadership Gap	Enterprise leadership capacity to drive change at scale	"Our AI lead is brilliant technically and completely underwater."
The Execution Gap	An enterprise framework to scale and govern the portfolio	"The pilots work. We just can't get them past the pilot."
The Performance Gap	Evidence that the investment is producing value	"Spend is up. I couldn't tell you the return."

The mandate and leadership gaps are gating: until they close, work on the other two does not compound. The diagnostic discipline is to find which gap is binding, and the most revealing signal is often not the average view of the organization, but the variance between executives who are describing the same company differently.

GAP 01

The Mandate Gap

Board-level strategy that never became operating authority.

What it is. An AI strategy exists at the board level, but it has never translated into authority, governance, or operating-model alignment. The ambition is real and is usually written down, yet the funded mandate behind it is absent: the budget and the incentives still sit where they were, not behind two or three named enterprise outcomes.

28%

of organizations have CEO-level ownership of AI governance, the factor most correlated with bottom-line impact.

MCKINSEY, THE STATE OF AI, 2025

HOW IT SHOWS UP

- Asked separately, executives describe the AI priorities differently.
- Capital is allocated initiative by initiative rather than against a portfolio.
- Use cases originate in enthusiasm rather than from mapped, high-value processes.
- Each business unit interprets the mandate against its own agenda.

IN THE ROOM

"Honestly? We signed off on the AI strategy last year and everyone nodded. But I sat down with my four BU heads last month, one at a time, and asked what we're actually going after. I got four different answers. And I'm sitting there thinking, okay, so what did we actually approve?"

Why it persists. The strategy was approved but never operationalized. No single forum owns the ambition, no outcome is named, and no baseline exists against which progress could be judged, so the mandate stays a statement of intent that every part of the organization is free to read its own way.

What it costs. Investment compounds slowly, or not at all. Capital flows to convenient, low-value use cases while the high-value opportunities that would justify the program go unfunded. This is the gap that gates the others: until the mandate is funded, money spent on operating model, leadership, or measurement cannot accumulate into performance.

How to close it. A one-page mandate the CEO can communicate inside and outside the company: specific outcomes, a defined horizon, named accountability, and a documented process baseline against which AI ROI can later be attributed. The signature that matters is the board's; without it, the strategy keeps being reinterpreted at the business-unit level.

GAP 02

The Leadership Gap

Leaders chosen for technical credibility, asked to carry enterprise change.

What it is. The roles meant to drive AI transformation are filled by technically capable people who do not have the organizational authority or change-leadership capability to drive it at scale. The individuals are genuinely capable; the role they have been placed in demands authority and change-leadership their selection never tested for.

3×

more likely: at AI high performers, senior leaders actively own AI initiatives, not just approve the budget. Leadership, not employee readiness, is identified as the biggest barrier to AI success.

MCKINSEY, SUPERAGENCY IN THE WORKPLACE, 2025

HOW IT SHOWS UP

- Selection prioritized technical credibility over enterprise leadership.
- The AI leader is isolated and depends on CEO escalation to move decisions.
- A small number of individuals carry the program; the succession layer is not in place.
- The function is one departure away from a meaningful capability gap.

IN THE ROOM

"Look, he's the smartest ML person in the building, no question. But I've basically asked him to get six business units to change how they work, and he doesn't run any of them. He has to come to me every time something needs to move. He's underwater. And if he leaves... I don't even want to think about where that leaves us."

Why it persists. AI ambition has outrun AI leadership. The organization can define a strategy but cannot yet execute it at scale, because execution demands authority, stakeholder reach, and change-leadership judgment that a technically weighted selection never tested for. Research finds C-suite leaders are more than twice as likely to blame employee readiness for stalled adoption as their own leadership, a misattribution that keeps the real gap open.

What it costs. Transformation pace is set by individual bandwidth rather than portfolio strategy, and the departure of one or two people becomes a single-point failure for the entire program. Talent decisions made under this pressure, usually adding headcount, tend to compound the miscast rather than correct it.

How to close it. An evidence-based read of who is in role and what each leader can actually carry, scored against the traits that predict transformation success, followed by targeted development for the leaders worth retaining and a structured search for the seats that must be re-filled. The assessment comes before the next hire, not after it.

GAP 03

The Execution Gap

Pilots that succeed in isolation and stall at the seams.

What it is. The pilots succeed; they simply do not scale. Cross-functional coordination breaks down, business units move in their own directions, and there is no enterprise framework to govern the portfolio as a whole.

26%

of companies have built the capability to move AI beyond proofs of concept into value. Gartner expects at least 30% of gen-AI projects to be abandoned after the pilot.

BCG, WHERE'S THE VALUE IN AI?, 2024 · GARTNER, 2024

HOW IT SHOWS UP

- Promising initiatives stall at the seams between functions.
- Underperforming pilots are never retired, because no executive owns the outcome.
- There are no defined pathways to move a pilot into production.
- Authority to approve, halt, or reassign use cases sits with no single forum; risk reviews convene after the architecture is already set.

IN THE ROOM

"We've got something like forty pilots going. Three of them are actually in production. Three. And they all demo great, that's the maddening part. But the second one of them needs another team to change how they do something, it just... dies. Nobody can make that call, so it sits."

Why it persists. The organization is not structured to develop, deploy, and scale AI, and its governance authority exists on paper but cannot be exercised. The finding that workflow redesign is the single strongest driver of earnings impact points to the root: scaling AI is an operating-model change, and the seams between functions are where that change has not been made.

What it costs. Capital keeps flowing to in-flight initiatives no one owns, pilots accumulate without ever being retired, and the program generates motion without crossing a single organizational boundary.

How to close it. Stand up the decision rights, escalation paths, and operating model that let governance actually be exercised; define the pathways that carry a pilot into production and the discipline to retire the ones that do not earn it. Where the root cause is executive misalignment, a facilitated session that surfaces the variance and produces a shared diagnosis comes first: alignment before architecture.

WHAT WE HEAR MOST

Three failures that feed each other

1. **Misaligned leadership and implementation.** The business never defined the vision, set the priorities, resourced the work, or built the incentives to deliver it. The technology gets built with no mandate behind it.
2. **Flashy use-cases over pragmatic ones.** The goal of AI was never perfection; it was beating the manual process that runs today. That takes knowing your own processes, pain points, and bottlenecks, and holding an honest baseline of how well the work gets done now. Most organizations never measured that baseline, so they cannot prove the return later.
3. **The wrong data.** You don't need more data; you need the right data, and you need to see all of it. High-quality data trapped in silos is invisible. Integrated data of poor quality is a swamp. And where two sources disagree, surfacing the contradiction is worth as much as clean consensus. That is a leadership call on standards and visibility, not an IT one.

The cascade. Misalignment leads to chasing flashy use-cases with no baseline, which leaves the data ungoverned, and the effort never shows its value. Each fix unlocks the next.

GAP 04

The Performance Gap

Rising investment the organization cannot yet measure.

What it is. Investment is growing; measurable performance is not. The ROI is anecdotal. The board can't see much, and money gets allocated on faith instead of evidence.

80%+

of organizations report no tangible enterprise-level earnings impact from their use of generative AI.

HOW IT SHOWS UP

- Activity is tracked, but outcomes are not attributed to it.
- There are no documented baselines against which ROI could be measured.
- The organization cannot reliably separate the initiatives that work from the ones that merely look busy.
- The next funding round rests on conviction rather than evidence.

IN THE ROOM

"Last quarter the board asked me, point blank, what we got for the forty million we've put into AI. And I had... examples. Anecdotes. Stuff that's going well. What I didn't have was a number I could stand behind. I'm still thinking about that one, frankly."

Why it persists. There is no value-measurement architecture: no process baselines, no attribution discipline, and no integration of AI performance into strategic planning and capital allocation. What gets reported is effort, because effort is what the organization is set up to see. Unclear business value is, in Gartner's analysis, among the leading reasons AI projects are abandoned after the pilot.

What it costs. The board cannot exercise real oversight, capital is allocated on faith, and a program that may well be working has no way to prove it, leaving it one budget cycle away from losing support it might deserve. Only 17% of organizations report that their board oversees AI governance at all.

How to close it. Build the performance architecture: establish process baselines, attribute outcomes to AI, and integrate AI performance into planning and capital allocation. Make it durable with a standing quarterly AI performance review at the board, structured around evidence and reporting failures as openly as successes. The cadence is the discipline.

A BOARD-LEVEL SELF-CHECK

Which gap is binding in your organization?

The four gaps look alike from the inside; each presents as "our AI isn't delivering." The questions below are a first separation. Put them to your leadership team individually, then compare the answers. Where the team cannot answer cleanly, or answers differently from one another, that gap is likely live, and the disagreement itself is a finding.

GAP

ASK YOUR LEADERSHIP TEAM

Mandate

GAP	ASK YOUR LEADERSHIP TEAM
	Can each executive name the two or three enterprise outcomes our AI investment is meant to produce, and would they all name the same ones? Is the mandate funded, or only stated?
Leadership	Does the person leading AI have the authority, not just the expertise, to change how other functions work? If they left tomorrow, would the program survive?
Execution	How many pilots have we run, and how many are in production? Is there a defined pathway from one to the other, and a forum with the authority to retire what isn't working?
Performance	If the board asked what our AI investment returned last quarter, could we answer with evidence rather than anecdote? Do we have baselines to attribute the outcome to?

Aggregate answers tell you where the organization stands. The variance between your executives tells you whether leadership is aligned. That is the more important reading.

SEQUENCE

The gaps compound. Close them in order.

The four gaps are connected. The mandate gap sits upstream of the others: until a funded ambition is in place, closing the leadership, execution, or performance gap produces local improvements that do not hold. The leadership gap caps execution, because a portfolio cannot be scaled by leaders who lack the authority to scale it. The execution and performance gaps sit downstream, where structure and evidence either compound the earlier work or expose its absence.

The implication is a fixed order. Align the executive team, fund the mandate, then resolve leadership, then build the structure and the measurement that let the program scale and prove itself. Each phase produces an instrument leaders can use: a mandate, a variance map, an operating rhythm, a re-assessment, and earlier phases unlock later ones.

HORIZON	WHAT IT CLOSES	THE MOVE
Days 1-90	Executive misalignment	A facilitated alignment session that surfaces where the team sees the organization differently and produces a shared diagnosis and a draft mandate.
Months 4-6	The Mandate Gap	Fund the one-page mandate against named outcomes; establish the process baseline that makes ROI attributable.
Months 7-12	The Leadership Gap	Assess the named AI leadership layer; develop the leaders to retain and run a structured search for the seats that must be re-filled.

HORIZON	WHAT IT CLOSES	THE MOVE
Months 13-18	Execution & Performance	Govern the portfolio through real decision rights and scale pathways; institutionalize value measurement and a quarterly board review.

HOW OCTANT DIAGNOSES THIS

From symptom to binding constraint

Telling the four gaps apart takes an outside read, not a self-assessment. What a company says about itself and what its structure actually allows are usually two different things. The research bears that out, and so does our own work.

100+

C-suite leaders, AI executives, and board directors interviewed in Octant's primary research underpin the four-gap framework.

OCTANT ADVISORY PRIMARY RESEARCH

Octant locates the binding gap through the AIR-O assessment, the organizational lens of the Octant AIR Index™. It scores eight organizational conditions on a common maturity scale, assigned by advisors from executive interviews, document review, and governance analysis, and it reads the variance between executives as a finding in its own right. The result is a precise answer to which gap is binding, and in what order to close them. Each gap maps to one or more of those conditions.

THE GAP	UNDERLYING ORGANIZATIONAL CONDITION
The Mandate Gap	Strategic Mandate, a gating condition. Below threshold, it caps the organization's maturity regardless of every other strength.
The Leadership Gap	Leadership Capability, a gating condition. Below threshold, it triggers a focused assessment of the named AI leadership layer before any further talent decision.
The Execution Gap	Operating Model, Governance & Decision Rights, Executive Alignment, the conditions that let a portfolio be coordinated, governed, and scaled.
The Performance Gap	Value Measurement & Improvement, whether outcomes are baselined, attributed, and fed back into capital allocation.

The diagnosis sets the agenda

Most organizations have already committed to AI; the question is no longer whether to invest but whether the investment will perform. Naming which of the four gaps is open, and closing them in the order that lets each one hold, is what turns committed capital into compounding return. The upside is documented: BCG finds that the organizations that close these gaps have achieved 1.5 times the revenue growth, 1.6 times the shareholder returns, and 1.4 times the return on invested capital of their peers. For a board, the highest-return decision available is usually an honest read of the organization that shows what to fix first, well ahead of the next model or platform.

ABOUT OCTANT ADVISORY

Octant Advisory is an AI transformation and executive search firm. We help boards and executive teams convert AI investment into enterprise performance, diagnosing the organizational conditions behind the four gaps and placing the leaders who can close them. The AIR Index™, spanning organizational, technology, and leadership readiness, is our proprietary framework for that diagnosis. Dr. Ian McCulloh and Maria Chaloux are the founding partners of Octant Advisory. This paper draws on Octant's primary research, interviews and working conversations with more than 100 C-suite leaders, AI executives, and board directors across our diagnostic engagements, alongside the published research cited below. To request a briefing or discuss an assessment, contact info@octantadvisory.com.

SOURCES

1. MIT NANDA, "The GenAI Divide: State of AI in Business 2025." Finding: ~95% of enterprise generative-AI pilots show no measurable P&L impact; failure attributed to organizational integration, not model quality. 2. McKinsey & Company, "The State of AI: How Organizations Are Rewiring to Capture Value," 2025. Findings: 80%+ report no tangible enterprise EBIT impact; workflow redesign is the strongest correlate of impact; 28% have CEO oversight of AI governance, 17% board oversight. 3. McKinsey & Company, "Superagency in the Workplace," 2025. Findings: leadership is the biggest barrier to AI success; only 1% of companies are AI-mature while 92% plan to increase investment; ~46-47% cite talent/skill gaps; high performers 3× more likely to have senior leaders owning AI. 4. Boston Consulting Group, "Where's the Value in AI?," Oct. 2024 (survey of 1,000 CxOs across 20+ sectors, 59 countries). Findings: only 26% have moved beyond proofs of concept to value; AI leaders achieved 1.5× revenue growth, 1.6× shareholder returns, 1.4× ROIC. 5. Gartner, "Gartner Predicts 30% of Generative AI Projects Will Be Abandoned After Proof of Concept by End of 2025," 2024. Drivers cited: poor data quality, inadequate risk controls, escalating costs, unclear business value. 6. Deloitte Global Boardroom Program, "Governance of AI: A Critical Imperative for Today's Boards," 2025 (695 board members and executives, 56 countries). Findings: 31% report AI is not on the board agenda; only 3% consider their organization very ready to deploy AI.