SINGULARICS

White Paper

# The Intent Gap

Why AI accelerates organizations in the wrong direction, and what it takes to close the gap.

Alexander S. Petty
SINGULARICS | singularics.ai



First published April 2026.

singularics.ai

# Contents

## Abstract

*Every large organization now has access to AI. The models work. The tools are mature. And yet most of the value remains uncaptured. This paper argues that the bottleneck is not technology, talent, or budget. It is intent: the distance between what an organization means to do and what it actually does, now widened and accelerated by the most powerful amplifier the enterprise has ever encountered. We examine the structural roots of the intent gap, the emerging shadow layer of invisible AI usage, the collapse of tolerance for vague requirements, and the structural changes required to close the gap before AI widens it further.*

## The Most Expensive Assumption in Business

Here is a scenario that plays out in thousands of organizations every quarter.

A leadership team articulates a strategy. It is clear, compelling, and well-funded. It is communicated broadly. Town halls are held. Slide decks are distributed. Everyone nods.

Six months later, the work that reaches the ground looks nothing like what was intended. Not because anyone disobeyed. Not because the strategy was wrong. But because somewhere between the boardroom and the teams doing the actual work, the signal degraded. Context was lost. Priorities were reinterpreted. Assumptions went unspoken.

If you have spent any time inside a large organization, you have seen this. You may have caused it. You have certainly been on the receiving end of it.

Henry Mintzberg documented this pattern decades ago, distinguishing between deliberate strategy (what leaders intend) and emergent strategy (what actually happens) [1]. Kaplan and Norton built an entire framework around the gap between strategy and execution [2]. The problem is well understood. It has never been solved.

And it has never mattered as much as it does right now.

Because the gap between intent and execution is about to get a rocket engine strapped to it. That rocket engine is AI. And unlike previous technologies, it does not wait for the organization to catch up. It accelerates whatever it finds, coherent or not, aligned or not, intentional or not.

This is the most expensive assumption in business: that intent survives translation.

It almost never does. And now it bites.

## Enter the Accelerant

There is a useful way to think about AI that strips away the hype.

AI is an accelerant.

Not a strategy. Not a solution. Not a replacement for human judgment. An accelerant. It takes whatever it is pointed at and makes it happen faster, at greater scale, with fewer people in the loop.

This is the source of its extraordinary value. It is also the source of its extraordinary risk.

An accelerant does not have opinions about direction. Gasoline does not care whether the building should be on fire. A jet engine does not ask whether the runway is pointed at the right city. AI does not ask whether the outcome it is optimizing for is the one the organization actually intended.

> *Pour AI on clarity and you get speed. Pour it on confusion and you get expensive confusion. At scale. The accelerant is neutral. The direction is everything.*

Consider what this means in practice. A marketing team with clear positioning and well-defined customer segments can use AI to produce months of content in weeks. The same team with vague positioning and undefined segments will use AI to produce months of wrong content in weeks. Both teams will be busy. Both teams will feel productive. Only one will have created value.

The McKinsey Global Institute estimates that generative AI could add $2.6 to $4.4 trillion annually in value across industries [3]. That number is real. But it is a ceiling, not a floor. The distance between the ceiling and what organizations actually capture will be determined by one thing: how clear their intent is before they accelerate.

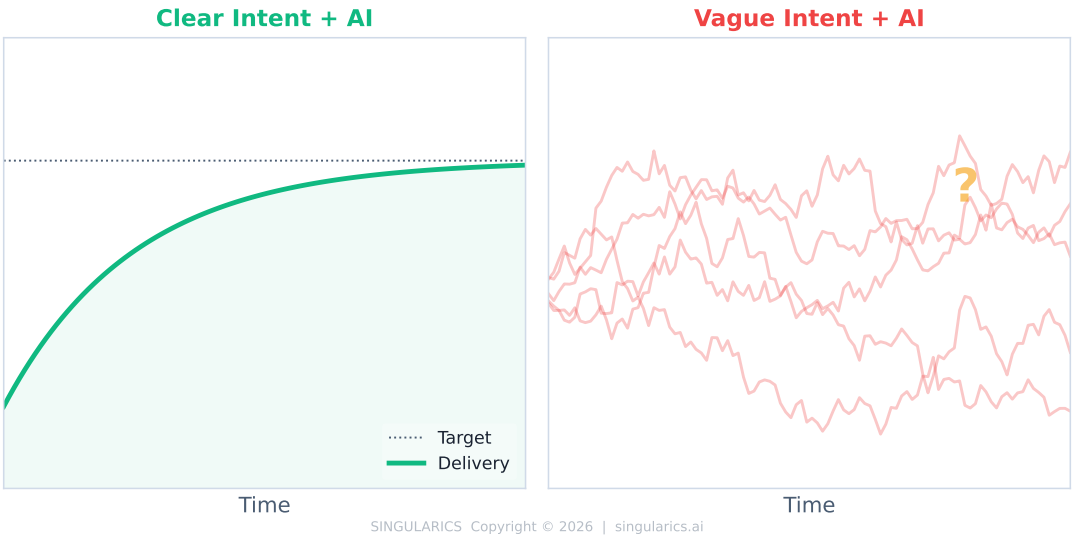


**Figure 1:** Clear intent produces convergent delivery. Vague intent produces divergent paths that AI accelerates simultaneously. The gap between these two outcomes is the intent gap, and it widens with every cycle.

## The Story Organizations Tell Themselves

Walk into the boardroom of any large organization and ask about AI. You will hear some version of the following:

"We have a comprehensive AI strategy. We have invested in the tools. We have trained our people. We are seeing early results."

This is what the organization believes is happening. It is built on metrics that feel reassuring: licenses purchased, pilots launched, training sessions delivered, innovation labs staffed. The dashboards are green. The quarterly report looks good.

Now walk downstairs. Talk to the people doing the work. You will hear something very different.

"I use ChatGPT on my phone because the company tool is locked down." "I have no idea what the AI strategy means for my job." "We did the training. Nobody uses it." "My manager does not know what AI can do, so they cannot tell me what to do with it."

The gap between these two conversations is the intent gap in action.

**80%** of AI projects do not deliver their intended business value. The technology works. The organizational absorption does not. [4]

This number is striking, but it should not be surprising. It mirrors the historical pattern of every major enterprise technology adoption. ERP implementations that met every technical milestone and were abandoned by users within a year. Agile transformations that changed process but not cognition. Digital transformation programs that produced new systems running old workflows.

The technology was never the hard part. The hard part has always been the same: changing how people actually work. And AI has not changed this fact. It has made it more urgent.

## The Adoption Funnel Nobody Wants to See

There is a funnel that every organization deploying AI should understand. Most have never drawn it.

It starts at the top with the number that feels good: 100% of licenses deployed. Every seat activated. Full coverage.

Then it narrows.

72% of those users have logged in at least once. Good. That means 28% never bothered to open it.

38% use it weekly. More than half of the people who tried it have stopped.

14% have actually changed how they work. Not just used the tool occasionally, but altered their workflow, their process, their output.

6% have experienced a measurable behavior shift. Their work looks different. Their output has changed shape. Their role has evolved.

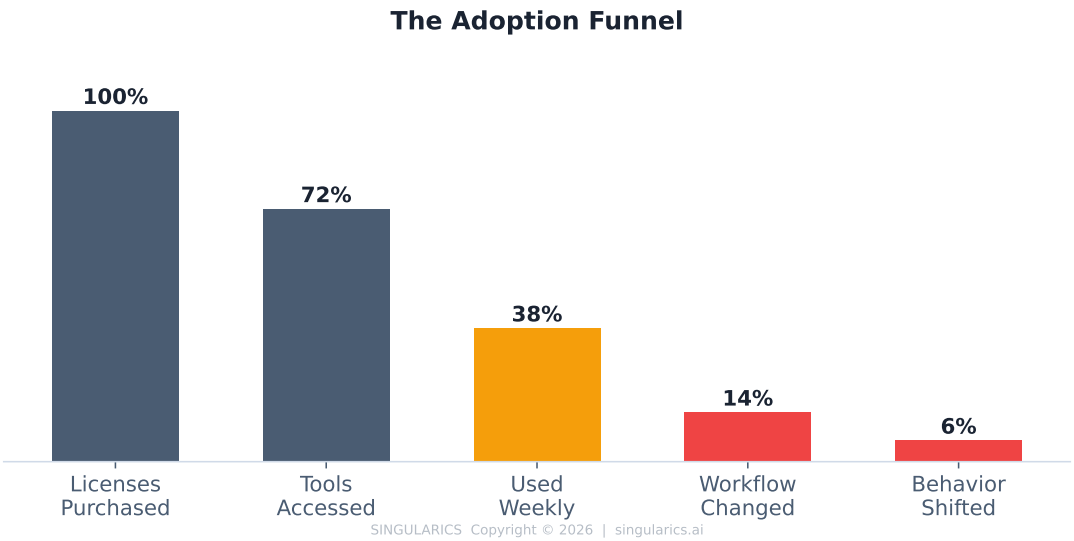


**Figure 2:** The adoption funnel reveals the decay between deployment and behavior change. Most organizations track the top of the funnel and declare success. The value lives at the bottom.

This funnel is the difference between an AI expense and an AI capability. An organization with 10,000 licenses and 600 people whose work has actually changed does not have 10,000 AI-enabled employees. It has 600. The other 9,400 have access to a tool they are not using, or using so sporadically that no value compounds.

Most organizations report the top of the funnel. "We deployed AI to 10,000 employees." This is a procurement metric. The real metric is at the bottom: "600 people work differently because of AI, and we can prove it."

Nobody wants to draw that funnel. It is too honest.

## The Shadow Layer

While leadership tracks deployment metrics and celebrates pilot launches, something else is happening inside the organization. Something no dashboard can see. Something that will reshape the organization whether leadership acknowledges it or not.

Individuals are adopting AI on their own.

### Three stories that are already happening

**The analyst.** A senior regulatory analyst at a financial services firm has been using Claude to draft compliance summaries for three months. Her output has doubled. Her error rate has dropped. The quality of her analysis has measurably improved. She has told no one, because her organization has not yet decided whether external AI tools are permitted for regulatory work. She uses it anyway, because not using it feels like choosing to be worse at her job.

**The product manager.** A PM at a technology company runs every customer feedback report through GPT-4 before his weekly leadership review. The insights he presents have become sharper. His recommendations are better supported. His ability to synthesize large volumes of qualitative data has gone from adequate to exceptional. His colleagues assume he is simply working harder. He is not working harder. He is working differently.

**The engineering lead.** A team lead uses AI to pre-review pull requests before the team's code review meetings. Obvious issues, style violations, and potential bugs are caught before they consume meeting time. Code reviews that took 90 minutes now take 30. The team has not discussed this change. It simply happened. Nobody has asked why the reviews got shorter, because the outcome is better and everyone is relieved.

### The implications

These are not edge cases. They are the norm. And the implications for organizations are profound.

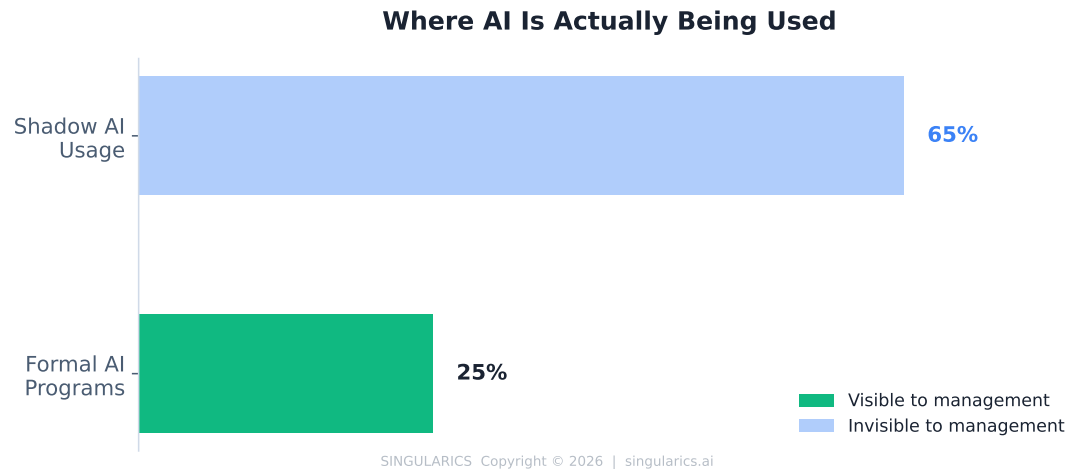


**Figure 3:** Formal AI programs represent a fraction of actual AI usage. The majority of adoption is invisible to management, unstandardized, and impossible to scale.

First, some employees are now operating at two to ten times the effectiveness of their peers. Their manager does not know. Their teammates may not know. The organization certainly cannot measure it, replicate it, or scale it. The productivity gains are real but invisible.

Second, this creates a new kind of inequality within the organization. Not one based on title, tenure, or department, but on who has quietly figured out how to augment their work and who has not. Two people with the same role and the same salary may now have fundamentally different capabilities. The organization cannot see this gap, which means it cannot address it.

Third, the organization cannot standardize what it cannot see. Best practices are emerging in isolation, one person at a time, with no mechanism to capture, share, or scale them. Team A figured out that AI transforms how they do documentation. Team B, doing similar work two floors up, has no idea.

> *The question is not whether people will adopt AI. They already have. The question is whether the organization will catch up to what individuals are already doing.*

And fourth, for some, the shadow layer produces the opposite response. Individuals who perceive AI as a threat to their expertise, their identity, or their job security do not adopt it. They resist it. Sometimes by ignoring it. Sometimes by actively undermining adoption efforts. Research on technology resistance has documented this pattern for decades [5]: when individuals feel their professional identity is threatened, rational self-preservation drives resistance. This is not

sabotage in a moral sense. It is a predictable human response to ambiguity about the future.

The shadow layer is not a bug in the adoption process. It is the adoption process, happening without organizational permission, support, or awareness.

## The Collapse of Tolerance for Ambiguity

Here is something that has always been true about organizations, but has never mattered this much:

Humans compensate for vague intent. AI does not.

When a product owner writes a requirements document that says "improve the customer experience," a human team absorbs the ambiguity. They ask questions in standup. They read between the lines. They bring judgment, context, and the accumulated knowledge of months or years of working together. They fill in the gaps that the document left empty.

This compensation is invisible. It happens in hallway conversations, in Slack threads, in the instinctive decisions that experienced professionals make a hundred times a day. It is so embedded in how organizations work that most leaders do not realize it exists.

AI cannot do any of this.

Hand AI a vague prompt and it will produce a vague output with complete confidence. It will not raise its hand and say, "I am not sure what you mean by this." It will not push back on unclear acceptance criteria. It will not notice that the stated goal conflicts with what the customer actually needs. It will not sense that the team has been quietly interpreting the requirement differently than the product owner intended.

It will execute. Literally, quickly, and at scale.

> *AI punishes ambiguity at machine speed. What used to be a communication problem absorbed by human teams is now a production problem amplified by machines.*

This represents a fundamental shift in what organizations must demand from the people who define intent. For decades, the role of the product owner, business analyst, or requirements author has been important but forgiving. The safety net of human interpretation covered for imprecision.

That safety net is gone.

The role of the person who defines intent, whatever an organization calls them, has gone from "nice if they are good at it" to "the entire system degrades if they are not." Every unclear outcome, every fuzzy acceptance criterion, every unstated assumption now multiplies through the system at machine speed.

This is the most underappreciated consequence of AI adoption. It has not just changed what technology can do. It has changed what clarity is worth. Clarity is now load-bearing.

## The Structural Root

The intent gap is not a communication problem that better meetings can solve. It is not a training problem that more workshops can fix. It is not a leadership problem that a more charismatic executive can overcome.

It is structural.

Large organizations are built on layers.

Strategy is articulated at the top. Execution happens at the bottom. In between, intent is translated through a series of handoffs: from executive vision to departmental objectives, from departmental objectives to team goals, from team goals to individual tasks.

This is how organizations have worked for a century. It is how they must work at scale. The problem is not the layers. The problem is what happens to intent as it passes through them.

At each handoff, something is lost. Context is simplified. Priorities are reinterpreted. Constraints are assumed rather than stated. The executive who said "we need to reduce customer churn by 20%" envisioned a comprehensive retention strategy. By the time that intent reaches the engineering team, it has become "add a cancel flow survey," which is one small piece of one possible approach, stripped of the context that gave it meaning.

Shannon's foundational work on communication channels [6] established a principle that applies directly here: every transmission introduces noise. The more links in the chain, the more noise accumulates. Organizations are extraordinarily noisy channels. Intent is the signal.

This has always been true. But it used to be slow enough that the damage was manageable. A misaligned team working without AI might lose a sprint. They would course-correct in a retrospective. They would ask questions. They would gradually find their way back toward the intent, even if the path was inefficient.

A misaligned team working with AI does not get that luxury. They can produce months of misaligned output in days. By the time anyone realizes the work has drifted from intent, the cost of the drift has already been incurred at machine speed.

The organizational structure that tolerates loose translation between layers is no longer viable when AI is in the loop. The tolerance for imprecision has collapsed. Not because organizations chose to raise their standards, but because AI has made the cost of imprecision visible, immediate, and expensive.

## The Change Management Gap

There is a second gap that compounds the first, and it is hiding in the budget.

Every AI initiative has a technology cost. Licenses, infrastructure, integration, security review. These are visible, quantifiable, and funded. They appear on spreadsheets. They are approved by procurement. They have owners, timelines, and success criteria.

Almost none have an equivalent investment in the organizational change required to absorb the technology.

The work of redefining roles, updating expectations, redesigning workflows, building governance, training people in context rather than in classrooms, and fundamentally shifting behavior is either unfunded, underfunded, or assumed to happen naturally.

It does not happen naturally. It has never happened naturally. The history of enterprise technology adoption is definitive on this point.

**$7.2M** average sunk cost per abandoned AI initiative. Not per failed technology. Per initiative where the technology worked and the organization could not absorb it. [7]

This number should reframe how organizations think about AI investment. The technology risk is manageable. The models work. The APIs are reliable. The tools are mature. The risk that deserves executive attention is the organizational absorption risk: the probability that the technology will be deployed successfully and adopted poorly.

AI transformation without dedicated change management is not transformation. It is tool deployment with a bigger budget. And the evidence suggests that tool deployment, by itself, does not produce the outcomes that justify the investment.

> *Every AI budget has a line for tools. Almost none have a line for the organizational change required to absorb them. The technology is the smaller problem. The change is the larger one.*

The organizations that capture the value of AI will be the ones that fund the change, not just the technology. That means investing in the people who define intent, the systems that translate it, the governance that sustains it, and the measurement that proves it.

This is not a secondary consideration. It is the primary one.

## Closing the Gap

Closing the intent gap is not a single initiative. It is not a project with a start and end date. It is a structural change in how the organization connects what it intends with what it delivers.

It requires three things.

### Precision at the source

Intent must be defined with enough clarity for AI to act on it productively. This means explicit outcomes, measurable success conditions, and unambiguous scope.

This sounds obvious. It is not. Most organizations operate on implicit intent: shared assumptions, unspoken expectations, and the ambient context that experienced teams build over time. This was adequate when human teams could compensate for ambiguity. It is not adequate when AI is in the loop.

Practically, this means investing in the capability of the people who define intent. Product owners, business analysts, program managers, and executive sponsors must learn to articulate outcomes with the precision that AI requires. Not the precision that satisfies a project management framework. The precision that a machine can act on without human interpretation.

This is a skill. It can be taught. It must be prioritized. And it must be valued by the organization as the critical capability it has become.

## A brief, worked

The artifact at the source is small enough to show whole. An insurance claims division declares an outcome, and the brief that reaches triage reads like this. The outcome, first notice of loss to coverage decision falls from fourteen days to two. Why now, two carriers in the footprint already decide in under three days and retention is moving. The sponsoring officer, the claims VP, who has approved the release. The candidate owner, a regional claims director, four of four on the qualification screens. Known constraints, state notification rules and the fraud review threshold, both already encoded in the network. Success, a coverage decision recorded within two business days for 90% of standard claims, measured from the intake timestamp.

Seven sentences, and triage can price it, the owner can write the first spec from it, and the harness will eventually check its success line. Notice what is absent. No solution, no architecture, no timeline, no resource plan. Those emerge from the daily cycle, and every attempt to front-load them into the brief is the project charter trying to come back.

The test that keeps the source honest is a one-minute walk. Pick any daily spec, read its trace sentence, follow it to the Effort, the Effort to the outcome, the outcome to its officer. If any hop needs a meeting to explain, the trace is broken at that hop, and it gets named in the open.

## Reinforcing translation across layers

Intent must survive the journey from strategy to execution without degrading. This requires deliberate, reinforcing translation between organizational layers.

The test is simple and unforgiving: can someone three levels below the executive who declared the strategy explain how their daily work connects to it? Not in the abstract. Not with corporate language. Can they draw a line from what they are doing today to the outcome the organization declared?

If the answer is no, the translation is broken. And AI is now accelerating the broken version.

Building this translation requires infrastructure, not just communication. Alignment must be observable: visible to everyone in the system, not just the people at the top. It must be measurable: quantified in terms that reveal drift before it compounds. And it must be continuously governed: maintained as an ongoing system property, not produced in an annual planning cycle and left to decay.

### Continuous measurement of alignment

Intent must be measured, not just declared.

A strategy offsite produces declarations. A quarterly business review produces updates on those declarations. Neither produces alignment. Alignment is a system property that must be continuously observed, in the same way that a manufacturing process monitors quality, a financial system monitors risk, or a software system monitors performance.

The organizations that stay aligned will be the ones that treat coherence as infrastructure. Instrumented. Measured. Maintained. Governed. Not as a meeting outcome that degrades the moment everyone returns to their desks.

This is a higher standard than most organizations have ever held themselves to. It is also the standard that AI makes necessary. Because when the system moves at machine speed, the feedback loop that detects drift must move at least as fast.

## The Opportunity

The intent gap is not inevitable. It is structural, and structures can be redesigned.

The organizations that close the gap before they accelerate with AI will capture value that others cannot reach. Not because they have better technology. Not because they have better talent. Not because they have a bigger budget. Because they know what they intend, and they have built the organizational infrastructure to translate that intent into execution without degradation.

> **74%** of global leaders say AI is a top investment priority, even in a recession. The investment is coming. The question is whether it will be directed by clear intent or diffused by organizational ambiguity. [8]

The technology is ready. The models work. The tools are available. These have been true for some time, and they will only become more true.

What remains unready, in most organizations, is the organizational capacity to absorb what the technology makes possible. The capacity to define intent clearly. To translate it faithfully across layers. To measure alignment continuously. To change how people work, not just what tools they use.

Agile taught organizations to change process. AI requires them to change cognition and role definition. The organizations that recognize this distinction, and invest in the change it demands, will define the next era of enterprise performance.

> *The question has never been whether your organization can use AI. The question is whether your organization knows what it intends. And whether its structure can hold that intention long enough for AI to deliver on it.*

The ones that do not will spend their AI budgets accelerating in circles. They will have the tools. They will have the talent. They will have the investment.

They will not have the intent.

And in the age of AI, intent is the only thing that determines whether acceleration creates value or destroys it.

**Is your intent clear enough for AI to act on? Let's find out.**

30 minutes. No pitch. Just an honest look at the gap.

Book a call | singularics.ai

**SINGULARICS**

We help large organizations close the intent gap between strategy and execution so that when they accelerate with AI, they accelerate in the right direction.

singularics.ai | Book a conversation

## References


[1] Mintzberg, H. (1994). *The Rise and Fall of Strategic Planning*. Free Press.

[2] Kaplan, R.S. & Norton, D.P. (2008). *The Execution Premium: Linking Strategy to Operations for Competitive Advantage*. Harvard Business Press.

[3] McKinsey Global Institute (2023). *The Economic Potential of Generative AI: The Next Productivity Frontier*. McKinsey & Company.

[4] RAND Corporation (2025). *AI Project Outcomes in Large Organizations: Patterns of Value Delivery and Failure*. RAND Research Report.

[5] Lapointe, L. & Rivard, S. (2005). A Multilevel Model of Resistance to Information Technology Implementation. *MIS Quarterly*, 29(3), 461–491.

[6] Shannon, C.E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27(3), 379–423.

[7] Deloitte (2025). *State of AI in the Enterprise*, 8th Edition. Deloitte Insights.

[8] KPMG (2026). *Global AI Pulse Survey: Q1 2026*. KPMG International.