

ARTICLE · AI ADOPTION

Your AI strategy is a document. Your operating model is a *mystery*.

Why so many AI programmes stall after the board says yes, and what turns a good strategy into real movement.

Michiel R. de Boer · Founder and Chief Architect, KAIVANT® · October 2026

IN SHORT

- Most organisations have an AI strategy. Very few have the operating model underneath it: who decides, what comes first, which rules apply, how results are measured and what happens when something fails.
- Without one, organisations either wait for permission or wander in every direction. Both look like progress on a slide.
- The fix is not another document. It is a short set of decisions, made early, and a clear baseline to measure against.

01

The strategy everyone liked

Six months after the board approved it, the AI strategy was still the most opened file on the shared drive. The head of strategy had sent it to every department head. The CEO had quoted it at the town hall. It was clear, well designed and backed by solid market research. It described a convincing picture of an AI-enabled organisation within three years.

Then I spoke to the team that had to make it happen. Since approval, three things had happened. One pilot had been extended. One vendor had been evaluated. And a working group had been set up to find more use cases.

The strategy was excellent. The operating model did not exist.

I have seen this story in many organisations, in many countries. The details change. The pattern does not.



Approved, admired, and still waiting for someone to act on it. Image: AI-generated.

02

Two different questions

An AI strategy answers one question: **what do we want to become?** It describes the ambition, such as higher productivity, new capabilities, lower costs or a stronger market position, and it explains why now.

An AI operating model answers a different question: **how will we actually get there?** It says who owns AI decisions and at what level, how use cases are chosen, which rules apply to which kind of AI, how results are measured, and what happens when something fails or falls short.

AI STRATEGY

What we want to become

The ambition

The expected gains

Why now, and why this way

Lives in a document

AI OPERATING MODEL

How we will get there

1 Who decides what

2 What comes first

3 Which rules apply

4 How we measure results

5 When we stop, fix or scale

Lives in weekly decisions

A strategy describes the destination. An operating model moves the organisation.

Most organisations have invested heavily in the first. Very few have the second. This is not a criticism of strategy documents. A document simply cannot close this gap on its own.

A strategy nobody can carry out is not a strategy. It is a wish. The operating model is what turns the wish into movement.

03

What the numbers say

The research tells the same story from different angles. A widely quoted, and widely debated, MIT study found that 95% of organisations were getting no measurable return from their generative AI investments.¹ Gartner expects over 40% of agentic AI projects to be cancelled by the end of 2027, because of rising costs, unclear business value or weak risk controls.² And McKinsey's latest global survey found that only 37% of organisations can link any profit impact to their use of AI.³

THE GAP BETWEEN AI ACTIVITY AND AI VALUE

95%

of organisations saw no measurable return from generative AI

MIT Project NANDA, 2025

40%+

of agentic AI projects expected to be cancelled by the end of 2027

Gartner, 2025

37%

of organisations can link any profit (EBIT) impact to AI


McKinsey, 2026

Three studies, three methods, one pattern. See the sources at the end.

Look closely and none of these is really a technology failure. In most cases the tools worked well enough. What failed was everything around them: who decides, what gets priority, how value is measured. That is an operating model failure.

The pattern behind these numbers is easy to recognise. The board approves a strategy. A small team gets the job on top of their normal work. They pick a use case, often one that was already being discussed before the strategy existed. A vendor is engaged. A pilot runs. The results go to leadership. Leadership is pleased. The pilot continues. Nothing scales.

HOW IT STARTS



The pilot loop. Every step looks like progress, so nobody stops to ask why nothing scales.

Meanwhile, the strategy document is updated to mention the pilot, and the next board report shows progress. On paper, the organisation has delivered on its AI commitment. In reality, it has done what it would have done anyway: one small, contained experiment that is not connected to how the organisation really works.

In my experience, many organisations only see this pattern after a year or two, and after a large part of the budget has gone.

How organisations set up their AI work matters too. In the same MIT study, AI tools built with outside partners reached full deployment about twice as often as tools built alone, roughly 67% against 33%.¹ Forrester made a similar point when it predicted that three out of four firms trying to build advanced AI agent systems on their own would fail.⁴

04

What an operating model contains

An AI operating model is not a thick document. The organisations that make steady progress with AI usually have one that is short, clear and used every week, rather than filed away. It has five parts.

1 Ownership and decision rights

Who has real authority over AI decisions? Not a sponsor in name, but someone with actual decision rights. Which decisions can a team make, which need a domain owner, and which need the executive team? In most organisations this is unclear, so every decision either escalates or stalls.

2 Clear priorities

How are AI initiatives chosen? What separates a valuable use case from a distraction, and who decides? Without clear criteria, organisations end up with the use cases that are easiest to build or pushed by the loudest voice. Neither reliably creates value.

3 Rules that match the risk

Different kinds of AI carry different risks. A tool that summarises internal meeting notes needs far less oversight than a system that makes decisions about customers. The operating model sets out which rules apply to which level of risk, and who makes sure they are followed.

4 Measurement against business results

Before a use case goes live, the operating model says how its success will be measured in business terms. Not usage figures, not the number of prompts, not time saved on its own, but revenue, cost, quality, speed and risk. The measurement is set up before launch, not pieced together afterwards. McKinsey's survey points the same way: the small group of AI high performers are the ones with defined processes to measure the impact of their AI work.³

5 A way to stop, fix or scale


What happens when an AI initiative underperforms? Who can decide to change it, pause it or stop it? Without a formal review, failing initiatives keep going simply because nobody has the mandate to end them.

The organisations making real progress usually have a few owners with a clear mandate, a simple decision framework that people actually use, and measurement that links AI activity to business results. They are still the minority.

05

Waiting or *wandering*

The missing operating model does not look the same everywhere. Across the countries I have worked in, it shows up in two very different ways.



Two cultures, one missing operating model. Many large organisations have both: head office waits while the edges wander.

In hierarchical organisations, the gap shows up as waiting. This is common in much of Asia, the Middle East, Latin America and parts of Europe, and in government-linked companies, banks and large corporates almost everywhere. Culture researcher Geert Hofstede called this high power distance: people accept and expect that decisions come from the top.⁵ Here, a missing operating model does not feel like a gap. It feels like an instruction to wait. Middle managers who could make AI decisions in their own areas wait for clear permission. Senior leaders who could give that permission wait for the strategy to be fully worked out first.

The result is a strange kind of standstill. Everyone has endorsed the strategy, yet nobody is carrying it out. This is not resistance. It is a sensible response to an unclear mandate.

In flatter organisations, the gap shows up as wandering. This is more common in places like the Netherlands, the Nordic countries, Australia and much of the tech sector. People do not wait for permission. They try tools on their own, often with personal accounts. The MIT study found that while only 40% of companies had bought an official AI subscription, workers at more than 90% of them were using personal AI tools for work.¹ There is plenty of activity, but no shared view of what works, what is safe or what creates value.

The cause is the same in both cases. That is why good governance is not about restriction. In a hierarchical culture, a clear mandate is the permission to move. In a flatter culture, it is the frame that turns scattered experiments into something the whole organisation can use.

Governance is not about restriction. It is about a clear mandate: who can decide what, and at what level.

The most common mistake: rules after the build

The operating model gap is made worse by a sequencing mistake that almost every organisation makes. The rules are designed after the AI is deployed, not before.

A team builds a pilot. The pilot works. Leadership approves a wider rollout. Then someone asks about the data rules. Someone else asks who is accountable if the system behaves in unexpected ways. Someone asks how success will be measured. These questions should have been answered before the pilot was designed. Now they are answered under time pressure, for a system people already use in their daily work.

The better approach is simple: design the rules together with the use case. The governance questions are not an extra step after the technical build. They come first. Organisations that get this order right, even imperfectly, are more likely to scale their AI work, because they planned for scale from the start.

Five first moves

- 1 Start with ownership, not policy.** The first document an AI operating model needs is a one-page ownership map. It shows which executive owns the AI strategy, which operational owners can make AI decisions in their areas, and which decisions need sign-off across functions. One senior meeting can produce it, and it removes the confusion that stalls most AI programmes.
- 2 Agree how you choose, before the next use case.** Before the next AI initiative is approved, agree the questions every proposal must answer. What business problem does it solve? How well does the current approach perform today? What does success look like in numbers? Who is accountable for delivering it? Used consistently, this simple check is worth more than any single use case it assesses.
- 3 Set risk levels before you need them.** Sort your likely AI use cases into three levels, and agree the criteria for each level before a specific use case forces the question.

LEVEL	WHO DECIDES	TYPICAL EXAMPLE
Low	The team, with standard controls	Summarising internal meeting notes
Medium	Review by the domain owner	Drafting customer emails that staff check before sending
High	Executive sign-off and ongoing oversight	AI that makes or shapes decisions about customers

Three risk levels. Agree them before a use case forces the question.

- 4 Measure business results from day one.** For every AI initiative already running, ask two questions. Which business result was it meant to improve? And how do we know whether it has? If there is no answer, set up the measurement now, even if building the baseline takes effort. Measurement added six months after launch is better than none. But by then, the baseline is gone.
- 5 Review every quarter, not every year.** The operating model is a living document. Tools, capabilities, risks and regulations change too fast for an annual review. A fixed quarterly review with the owners keeps the model current, and shows everyone that the organisation is serious.



An ownership map takes one senior meeting. It removes the confusion that stalls most AI programmes. Image: AI-generated.

08

Start with a baseline

Every one of these moves depends on something most organisations skip: knowing where they stand today. You cannot prove what AI has changed if you never measured the starting point. And you cannot set sensible priorities if you do not know how well your organisation already uses AI, or what it is doing to your people.

This is where KAIVANT® helps. It establishes your organisation's AI baseline, shows how effectively AI is being used and what it is doing to your people's capability, and gives you a solid starting point for an AI adoption strategy you can actually measure. It is not a ranking or a compliance score. It is a developmental map that shows where you stand and what to build next.

09

The honest assessment

The question is not whether your organisation has an AI strategy. It almost certainly does. The question is whether there is an operating model underneath it: the decision rights, the priorities, the risk levels, the measurement and the feedback loops that turn intention into movement.

If the answer is no, your strategy is still a document. That is not a failure. It is where most organisations are starting from today.

The operating model is what you build next.

SOURCES

1. MIT Project NANDA (A. Challapally, C. Pease, R. Raskar, P. Chari), "The GenAI Divide: State of AI in Business 2025", July 2025. Based on a review of 300+ public AI initiatives, 52 interviews and 153 survey responses. Its methods and its definition of success have been debated.
2. Gartner, "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027", press release, 25 June 2025. gartner.com/en/newsroom
3. McKinsey & Company, "The state of AI in 2026: On the road to ROI", August 2026 (1,719 respondents in 97 countries). mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai
4. Forrester, "Predictions 2025: An AI Reality Check Paves The Path For Long-Term Success", October 2024. forrester.com/blogs/predictions-2025-artificial-intelligence
5. Geert Hofstede, *Culture's Consequences*, 2nd edition, Sage, 2001.

YOUR NEXT STEP

Where does your organisation *really* stand?

Before you build the operating model, measure your starting point. Both first steps are free and need no account.

FOR LEADERS · ABOUT 25 MINUTES

What is AI doing to your organisation? Take the Leadership Pulse

kaivantscore.com/pulse

FOR INDIVIDUALS · UNDER 10 MINUTES

Is AI making you more valuable? Take the free Kaivant-I self-test

kaivantscore.com/KI-test

GO DEEPER

Why AI pilots do not scale, and what to check before you roll out

kaivantscore.com/decisions/pilot-purgatory

ABOUT THE AUTHOR

Michiel R. de Boer is the founder and Chief Architect of KAIVANT®, and co-owner and managing director of MindMagine. For more than 30 years, he has helped organisations in Asia Pacific and beyond with the best practices of programme, project and process management, governance and leadership. Based in Kuala Lumpur, he now focuses on helping organisations everywhere become truly AI-native, without losing the human capability that makes them work.



kaivant.org

© 2026 KAIVANT®. KAIVANT is a registered trademark in Malaysia.