


BEYOND THE HYPE CYCLE

ChatGPT Changed the Question — Why AI Strategy Cannot Start with the Technology

A Practitioner's Perspective on Strategic Misdirection in the Age of Generative AI

Giovanni Leonardi

August 2023



Contents

1. Executive Summary

2. The Rush to Strategy

3. Why Technology–First Fails

3.1 The Use Case Trap

3.2 The Governance Vacuum

3.3 The Capability Illusion

4. The Strategy That Starts Elsewhere

5. The Data Question That Nobody Wants to Answer

6. Organisational Readiness Is Not a Technology Problem

6.1 Skills and Understanding

6.2 Change Capacity

6.3 Governance Maturity

7. The Leadership Challenge

8. What a Credible AI Strategy Looks Like

9. The Road Ahead

3

3

4

4

4

5

6

6

6

7

7

8

8

ChatGPT Changed the Question — Why AI Strategy Cannot Start with the Technology

Leadership · No. 83

Originally written August 2023 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“The organisations that will succeed with AI are not those that moved fastest to adopt ChatGPT, but those that paused long enough to ask what problem they were actually trying to solve.”

1. Executive Summary

Since the public release of ChatGPT in late 2022, a pattern has emerged across sectors that is as predictable as it is damaging: organisations rushing to produce AI strategies that begin — and too often end — with the technology itself. Board papers are written around what large language models can do. Pilot programmes are launched to explore what generative AI might automate. Innovation teams are tasked with finding use cases for a solution that has not yet been matched to a problem.

This paper argues that this approach is structurally flawed. Not because generative AI lacks transformative potential — it plainly does not — but because strategy that begins with a technology inevitably produces a procurement exercise dressed in strategic language. The result is a portfolio of disconnected pilots, a growing gap between executive expectation and operational reality, and an organisation no closer to the capability shifts that would actually move the needle.

The alternative is not to ignore generative AI. It is to start in a fundamentally different place: with the business capability gaps, decision-making failures, and operational constraints that the organisation already knows it needs to address — and then to ask, rigorously, whether and how AI might close those gaps. This is not a semantic distinction. It produces a different strategy, a different investment portfolio, and a different set of organisational conversations.

2. The Rush to Strategy

The speed at which generative AI entered the corporate consciousness has no real precedent in enterprise technology. Cloud computing took years to move from technical curiosity to boardroom priority. Mobile strategy evolved over a decade. Even the internet itself — for all its revolutionary impact — gave organisations several years of gradual adoption before strategic urgency set in.

ChatGPT compressed that timeline to weeks. By early 2023, boards that had never discussed artificial intelligence were demanding AI strategies. Chief executives who could not define a neural network were announcing AI-first transformations. The pressure was not driven by a careful assessment of organisational need; it was driven by a combination of media intensity, competitive anxiety, and the genuinely impressive experience of interacting with a system that could write, summarise, and reason in natural language.

The pattern I have observed across multiple sectors is remarkably consistent. A senior leader — often the CEO, sometimes a board member — returns from a conference, a dinner, or simply an evening experimenting with ChatGPT, and issues a mandate: we need an AI strategy. The mandate flows downward. A team is assembled, typically drawing from IT, innovation, and sometimes strategy. A deadline is set — usually aggressive, because the implicit message is that competitors are already moving.

What follows is not strategy in any meaningful sense. It is a technology assessment wrapped in strategic framing.

3. Why Technology-First Fails

The fundamental problem with beginning an AI strategy from the technology is that it reverses the logic of strategic planning. Sound strategy starts with an honest assessment of where the organisation is, where it needs to be, and what is preventing it from getting there. Technology enters the conversation as a potential enabler — one of several — once the gaps are understood.

When technology leads, three predictable failures follow.

3.1 The Use Case Trap

The first failure is what might be called the use case trap. Tasked with finding applications for generative AI, teams naturally gravitate toward the obvious: document summarisation, customer service chatbots, content generation, code assistance. These are not wrong — they are real capabilities of the technology — but they are generic. They are the same use cases that every organisation in every sector identifies, because they are inherent to the technology rather than specific to the organisation's strategic position.

The result is a portfolio of pilots that looks impressive in a board presentation but has no strategic coherence. Each pilot exists because the technology can do it, not because the organisation needs it done. The cumulative investment may be significant, but the cumulative strategic value is negligible.

3.2 The Governance Vacuum

The second failure is a governance vacuum that technology-first strategies inevitably create. When an organisation begins with the question what can this technology do?, it defers the harder questions — what data will it need, what decisions will it influence, what risks does it introduce, who is accountable for its outputs — until implementation forces them to the surface. By that point, pilots are already running, vendor relationships are already forming, and the political cost of pausing to address governance is higher than

the perceived risk of continuing without it.

In my experience, this is not carelessness. It is structural. Technology-first strategies create momentum before they create the frameworks needed to manage that momentum. The governance conversation arrives late, underfunded, and competing for attention with the very initiatives it is supposed to govern.

3.3 The Capability Illusion

The third failure is subtler but arguably the most damaging. Technology-first strategies create what I would describe as a capability illusion: the belief that deploying a tool is equivalent to building a capability. An organisation that installs a large language model into its customer service operation has not built an AI capability. It has installed software. The capability — the ability to use AI systematically, to learn from its outputs, to improve its application, to govern its risks, to adapt as the technology evolves — requires organisational infrastructure that technology deployment alone does not create.

This illusion is particularly dangerous because it satisfies the board's demand for visible progress. Pilots are running. Demonstrations are impressive. The organisation appears to be moving. But the distance between a working pilot and an embedded organisational capability is vast, and technology-first strategies rarely acknowledge it, let alone plan for it.

4. The Strategy That Starts Elsewhere

The alternative approach begins not with the technology but with a disciplined assessment of organisational need. This is not a novel concept — it is the foundation of any sound strategic process — but it has been abandoned with remarkable speed in the face of generative AI's arrival.

The starting questions are deliberately technology-agnostic:

1. Where does this organisation make its worst decisions, and why?
2. Which operational processes consume disproportionate resource for the value they deliver?
3. What capability gaps have persisted despite previous investment, and what structural factors sustain them?
4. Where does institutional knowledge reside in fragile forms — in the heads of individuals, in unstructured documents, in processes that depend on experience rather than system?
5. What are the organisation's binding constraints on change — not the stated ones, but the real ones?

These questions produce a fundamentally different map of opportunity. Instead of a list of things generative AI can do, they produce a list of things the organisation needs to do better — and a clear-eyed view of why it has not done them already.

Only then does the technology enter the conversation. And when it does, it enters as a candidate solution to a defined problem, subject to the same rigorous evaluation that any

strategic investment should receive: Does it address the root cause or merely the symptom? Does the organisation have the data, governance, and operational maturity to deploy it effectively? What is the realistic timeline from pilot to embedded capability? What does failure look like, and can the organisation tolerate it?

5. The Data Question That Nobody Wants to Answer

No discussion of AI strategy is complete without confronting the data question — and confronting it honestly, not in the sanitised form it typically takes in strategy documents.

The pattern is consistent: organisations that begin with the technology assume their data is ready, or at least close to ready. This assumption is almost always wrong, and it is wrong in ways that are expensive to discover late.

The uncomfortable truth is that most organisations' data estates are not AI-ready. They are not even analytics-ready in the way that would be required for reliable AI deployment. Years of underinvestment in data quality, metadata management, and data governance have created environments where the raw material that AI needs to function is fragmented, inconsistent, poorly documented, and often wrong.

A technology-first strategy treats this as an implementation detail — something to be resolved during the pilot phase. A capability-first strategy treats it as a strategic constraint — something that must be understood, costed, and planned for before any AI investment can be credibly evaluated.

The difference is not academic. Organisations that discover their data limitations during pilot implementation face a choice between two bad options: invest significantly in data remediation (delaying the AI initiative and consuming budget that was not planned for) or proceed with poor data (producing AI outputs that are unreliable and potentially dangerous). Neither outcome serves the strategic intent.

6. Organisational Readiness Is Not a Technology Problem

Beyond data, the question of organisational readiness extends into territory that technology strategies rarely address: the human and structural factors that determine whether any technology deployment succeeds or fails at scale.

6.1 Skills and Understanding

Generative AI requires a level of organisational understanding that does not yet exist in most enterprises. This is not about technical skills — though those matter — but about the broader workforce's ability to work effectively alongside AI systems. When is an AI output trustworthy? When does it need human verification? How should workflows change to incorporate AI assistance without creating new dependencies or new risks?

These questions cannot be answered by a technology deployment. They require investment in education, in new ways of working, and in the gradual development of organisational judgment about AI's strengths and limitations.

6.2 Change Capacity

Every organisation has a finite capacity for change. AI does not exist in isolation; it competes for change capacity with every other initiative in the portfolio. Technology-first strategies rarely account for this constraint, because they frame AI as an addition to the portfolio rather than a demand on the same pool of organisational energy, attention, and tolerance for disruption that every other programme draws from.

The organisations I have seen navigate this most effectively are those that explicitly position AI within their broader change portfolio — not as a special category that operates outside normal portfolio disciplines, but as a set of initiatives that must justify their claim on organisational capacity alongside everything else.

6.3 Governance Maturity

AI governance is not a bolt-on. It requires decision-making frameworks, accountability structures, risk assessment capabilities, and monitoring mechanisms that most organisations have not yet built. A capability-first strategy identifies these requirements early and builds them into the investment case. A technology-first strategy discovers them late and treats them as obstacles to be navigated.

7. The Leadership Challenge

At its core, the question of where AI strategy begins is a leadership challenge. The technology-first approach persists not because it is analytically sound but because it is politically convenient. It responds to the board's demand for action. It produces visible progress quickly. It defers the harder conversations — about data, about capability, about organisational readiness — to a point where someone else may be responsible for them.

The leadership challenge is to resist the seductive simplicity of the technology-first approach and to insist on the harder, slower, but ultimately more productive discipline of starting with organisational need.

This requires a particular kind of courage: the willingness to tell a board that is excited about ChatGPT that the organisation's AI strategy should not begin with ChatGPT. That it should begin with a rigorous assessment of where AI could make a material difference to performance, and an honest evaluation of whether the organisation is ready to realise that difference.

It also requires a particular kind of patience: the willingness to invest in foundations — data quality, governance frameworks, skills development, change capacity — that do not produce impressive demonstrations but without which impressive demonstrations remain just that: demonstrations.

8. What a Credible AI Strategy Looks Like

A credible AI strategy, built from organisational need rather than technological capability, has several distinguishing characteristics:

- It begins with a capability gap analysis, not a technology assessment. The questions it asks are about the organisation, not about AI.
- It treats data readiness as a strategic investment, not an implementation detail. The data estate is assessed honestly, and the cost and timeline for remediation are built into the investment case.
- It includes a governance framework that is designed alongside the strategy, not bolted on after deployment begins.
- It accounts for organisational change capacity explicitly, positioning AI initiatives within the broader change portfolio rather than outside it.
- It distinguishes between pilots that test a hypothesis and pilots that merely demonstrate a capability. Every pilot has a clear decision gate: what must be true for this to proceed to scale, and how will we know?
- It sets realistic timelines that reflect the gap between a working prototype and an embedded organisational capability — typically measured in years, not quarters.
- It is honest about what the organisation does not yet know and builds in structured learning rather than assuming that knowledge will accumulate through doing.

9. The Road Ahead

Generative AI will reshape enterprise operations. The organisations that benefit most will not be those that moved fastest, but those that moved most deliberately — that took the time to understand what they actually needed before deciding what to buy.

The pattern that recurs across complex technology adoption is this: the early movers who invest heavily in the technology without investing equivalently in the organisational capability to use it effectively are overtaken by later movers who get the sequence right. Cloud computing taught this lesson. Digital transformation taught it again. There is every indication that AI will teach it a third time.

The question for today's leaders is whether they are willing to learn it in advance, or whether they will insist on learning it the expensive way.

“The organisations that will succeed with AI are not those that moved fastest to adopt ChatGPT, but those that paused long enough to ask what problem they were actually trying to solve.”
