


BEYOND THE ROADMAP

AI Strategy Cannot Start with the Technology

A Practitioner's Perspective on the Gap Between Boardroom Enthusiasm and Operational Readiness

Giovanni Leonardi

August 2023



Contents

1. Executive Summary

3

2. The Moment Everything Changed

3

3. The Strategy Rush

4

4. Why Technology-First Strategy Fails

4

4.1 The use-case trap

4

4.2 The vendor dependency problem

5

4.3 The capability gap is not where leaders think it is

5

5. What Strategy Should Start With

5

6. The Leadership Gap

6

7. The Path Forward

7

8. Conclusion

8

AI Strategy Cannot Start with the Technology

Leadership · No. 95

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

"The organisations that will extract lasting value from AI are not the ones that moved fastest — they are the ones that understood what they were moving towards."

1. Executive Summary

Nine months after ChatGPT's public launch in November 2022, the pattern across industries is now clear enough to name. Organisations of every size and sector have rushed to produce AI strategies, often within weeks of board-level pressure to "have a plan." The overwhelming majority of those strategies start with the technology: which large language models to adopt, which vendor platforms to evaluate, which use cases to pilot. This paper argues that this approach is structurally flawed — not because the technology is unimportant, but because starting with it guarantees that the organisation will optimise for procurement rather than for value.

The evidence from the first wave of enterprise responses to generative AI suggests that the organisations making genuine progress share a common characteristic: they started with a clear articulation of what problems they were trying to solve and what capabilities they needed to build, and only then asked which technologies could help. The organisations that started with the technology — the majority — have produced strategies that are, in effect, vendor selection exercises dressed in strategic language.

This paper examines why the technology-first pattern persists, what it costs, and what a more defensible approach looks like. It is written for senior leaders who suspect their AI strategy may be answering the wrong question.

2. The Moment Everything Changed

The launch of ChatGPT on 30 November 2022 was not, in any technical sense, a breakthrough. Large language models had existed for years. GPT-3 had been available through OpenAI's API since 2020. Researchers and specialist teams had been working with transformer architectures since 2017. What ChatGPT changed was not the technology — it was the audience. For the first time, a capability that had been confined to technical teams became viscerally accessible to anyone with a browser.

The consequences for enterprise strategy were immediate and profound. Within weeks, boards and executive committees that had treated AI as a medium-term consideration were asking urgent questions: What does this mean for us? What are our competitors doing? Where is our AI strategy? The speed of this shift was unprecedented. No enterprise technology in recent memory — not cloud, not mobile, not even the internet itself — moved from specialist curiosity to boardroom imperative in under three months.

By early 2023, the pressure to respond had produced a visible pattern. Organisations stood up AI working groups, appointed AI leads, commissioned strategy documents, and launched pilot programmes. The activity was intense, visible, and — in a striking number of cases — misdirected.

3. The Strategy Rush

The speed of the organisational response was understandable. The quality was not.

The pattern I have observed across financial services, professional services, healthcare, and the public sector follows a remarkably consistent sequence. First, a senior leader — often the CEO or a board member — experiences ChatGPT directly or reads about it and asks the organisation what it is doing. Second, the technology function is asked to produce an AI strategy, typically within four to eight weeks. Third, the resulting strategy focuses overwhelmingly on technology: platform selection, model evaluation, data infrastructure readiness, and a portfolio of potential use cases ranked by feasibility and impact.

This sequence produces a document that looks like a strategy but functions as a technology procurement plan. It answers the question "which AI tools should we buy and where should we deploy them?" rather than the question that actually matters: "what do we need AI to do for this organisation, and are we capable of absorbing it?"

The distinction is not semantic. It determines whether the organisation's AI investment creates lasting capability or produces a collection of disconnected pilots that never reach operational scale.

4. Why Technology-First Strategy Fails

There are specific, structural reasons why starting with the technology leads to poor outcomes. They are worth examining because they are not obvious to leaders who are, understandably, focused on keeping pace with a fast-moving landscape.

4.1 The use-case trap

The centrepiece of most technology-first AI strategies is a use-case portfolio: a ranked list of potential applications, typically generated through a combination of brainstorming workshops and vendor-supplied frameworks. The implicit logic is appealing — identify the highest-value opportunities, prove the technology works in controlled conditions, then scale the winners.

The problem is that use cases identified this way almost never scale. They are selected for technical feasibility, not for organisational readiness. A use case may be technically straightforward — a document summarisation tool, an automated first-pass review, a chatbot for internal queries — and still fail in practice because the process it touches is poorly defined, the data it needs is inaccessible, the people it affects have not been consulted, or the governance framework for AI-assisted decisions does not exist.

The result, visible across sectors by mid-2023, is a growing inventory of successful proofs of concept that cannot make the transition to production. The technology works. The organisation does not.

4.2 The vendor dependency problem

A strategy that starts with technology inevitably becomes shaped by vendor capabilities. When the first question is "which platform should we use?", the available answers are defined by what vendors sell. This creates a dependency that most organisations do not recognise until it is too late: the strategy becomes bounded by the vendor's roadmap, the vendor's pricing model, and the vendor's assumptions about how AI should be deployed.

The large cloud providers — and the growing ecosystem of AI-specialist vendors — are sophisticated organisations with their own commercial incentives. Their interest is in platform adoption, consumption-based revenue, and long-term lock-in. These are legitimate business objectives, but they are not the same as the organisation's objectives. A strategy that starts with the vendor's capabilities rather than the organisation's needs will inevitably optimise for the vendor's success.

4.3 The capability gap is not where leaders think it is

The most consequential failure of technology-first strategy is that it misidentifies the organisation's primary constraint. Most AI strategies treat the technology as the binding constraint — if only we had the right platform, the right models, the right data infrastructure, we could deliver AI at scale. In practice, the binding constraint is almost never the technology. It is the organisation's capacity to absorb AI into its operations: to redesign workflows, to reskill staff, to establish governance, to make decisions about how human and machine judgement should interact.

The technology is the easiest part of AI adoption. It is also the part that receives the overwhelming majority of investment, attention, and strategic focus. This inversion explains why so many AI programmes deliver technical capability without organisational value.

This capability gap is structural, not incidental. It reflects decades of underinvestment in the organisational capabilities that AI adoption actually requires: change management, process design, data governance, and — above all — leadership that understands technology well enough to make strategic decisions about it without being captured by it.

5. What Strategy Should Start With

If not the technology, then what? The organisations I have seen make genuine progress with AI share a common starting point, and it is not a platform evaluation.

They start with a clear, specific articulation of the problems they are trying to solve or the capabilities they are trying to build. Not "we need an AI strategy" but "we need to reduce the time from application to decision in our underwriting process from fourteen days to three, and we believe AI-assisted triage could be part of that." Not "we should be using large language models" but "our knowledge management is broken — critical expertise leaves when people leave, and our institutional memory is fragmented across systems that do not talk to each other."

This problem-first framing changes everything that follows:

1. It makes the technology decision subordinate to the business need. The question becomes "which technology best addresses this specific problem?" rather than "where can we apply this exciting technology?" This is a fundamentally different question, and it produces fundamentally different answers.
2. It forces early engagement with operational reality. A problem-first approach requires the organisation to understand the current process, the current pain points, the current data landscape, and the current people — all before any technology is selected. This understanding is precisely what technology-first strategies skip.
3. It creates natural success criteria. When the starting point is a specific problem, success is defined by whether the problem is solved, not by whether the technology is deployed. This distinction determines whether the organisation measures what matters or celebrates what is visible.
4. It reveals the real constraints. A problem-first approach surfaces the organisational barriers — data quality, process ambiguity, skills gaps, governance voids — that will determine whether any technology solution succeeds. Technology-first strategies discover these barriers late, usually after significant investment has been committed.

6. The Leadership Gap

The deeper question beneath the strategy question is a leadership question. Why do intelligent, experienced leaders allow their organisations to pursue technology-first AI strategies when the limitations of that approach are, on reflection, foreseeable?

Part of the answer is speed. The pressure to respond to the generative AI moment has been extraordinary, and a technology-first strategy can be produced quickly — vendors are eager to help, consultancies have frameworks ready, and the output looks substantive. A problem-first strategy takes longer because it requires the organisation to do the harder work of understanding its own operations, capabilities, and constraints.

Part of the answer is knowledge. Many senior leaders do not understand AI well enough to challenge the technology-first framing when it is presented to them. They know they do not understand it, and this knowledge gap makes them dependent on the people who do — who are, almost always, technologists or vendors with their own perspective and incentives.

“The organisations that will extract lasting value from AI are not the ones that moved fastest – they are the ones that understood what they were moving towards.”

And part of the answer is institutional. The structures that produce strategy in most organisations – the planning cycles, the governance forums, the investment committees – are designed for capital-intensive, technology-led programmes. They are good at evaluating platform costs, sizing implementation teams, and tracking deployment milestones. They are poor at evaluating organisational readiness, measuring capability-building, or governing the messy, iterative process of embedding a new technology into how an organisation actually works.

7. The Path Forward

This paper is not an argument against AI investment. The potential of large language models and the broader generative AI landscape to transform how organisations operate is real, and the organisations that ignore it will pay a competitive price. The argument is narrower and more specific: that the way most organisations are approaching AI strategy – starting with the technology and working outward – is producing strategies that will not deliver their intended value.

A more defensible approach would include the following elements:

Start with problems, not platforms. Identify the three to five operational challenges where AI could make a material difference, and build the strategy around solving those challenges. Let the technology choice follow the problem definition, not precede it.

Invest in organisational readiness alongside technical capability. For every pound spent on AI technology, allocate a proportionate investment in the capabilities the organisation needs to absorb it: process redesign, data governance, change management, and skills development. The current ratio – overwhelmingly weighted toward technology – is a predictable recipe for underperformance.

Build AI literacy at the leadership level. Leaders do not need to understand transformer architectures. They do need to understand what AI can and cannot do, how it differs from previous technologies, and what questions to ask when evaluating AI proposals. Without this literacy, leaders will continue to delegate strategic decisions to technologists and vendors who are not positioned to make them.

Accept that speed is not the primary virtue. The pressure to move fast is real but misleading. An organisation that takes six months to build a well-founded AI strategy and twelve months to deliver its first operational capability will outperform one that produces a strategy in six weeks and spends the next two years cycling through pilots that never scale.

Create governance that fits the technology. AI is not like previous enterprise technologies. It is probabilistic, it degrades, it can produce harmful or biased outputs, and it operates in ways that are difficult to explain. The governance frameworks designed for

deterministic systems — ERP, CRM, core banking — are not adequate. Organisations need governance that addresses AI's specific characteristics: model monitoring, output validation, bias detection, and clear accountability for AI-assisted decisions.

8. Conclusion

ChatGPT did not change what AI can do — it changed who knows about it. That shift in awareness created a wave of strategic activity across every sector, most of which has been directed at the wrong question. The question is not "how do we deploy AI?" It is "what do we need to become in order to use AI well?"

The first question produces a technology roadmap. The second produces a transformation strategy. The difference between them will determine which organisations create lasting value from the most significant technology shift since the commercial internet, and which simply add another layer of underused capability to an already complex technology estate.

The evidence from the first nine months is clear. It is not too late to change the question.