


BEYOND THE WORKSHOP

Use Case Hunting — The Wrong Way to Start an AI Programme

A Practitioner's Perspective on Why Top-Down AI Opportunity Identification Produces Programmes That Never Deliver

Giovanni Leonardi

June 2022



Use Case Hunting — The Wrong Way to Start an AI Programme

Programme · No. 212

Originally written June 2022 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

"The use case is not the unit of AI value — the capability is, and capabilities cannot be discovered in a two-day workshop by people who have never built one."

1. The Ritual of Use Case Identification

There is a ritual that plays out in organisations launching AI programmes, and it is so consistent that it deserves to be named. A senior leader decides the organisation should "do something with AI." A programme team is formed. The team's first major activity is a series of workshops — sometimes facilitated by a consultancy, sometimes by the internal data science team — in which business stakeholders brainstorm potential applications of artificial intelligence across the organisation. The outputs are captured on sticky notes, transferred to spreadsheets, scored against criteria such as feasibility, impact, and data readiness, and ranked into a prioritised portfolio of use cases.

The programme then selects the top five or ten use cases, assigns teams to build proofs of concept, and declares itself underway.

This approach is now so widespread that it has become the default methodology for AI programme initiation. Vendor playbooks recommend it. Consultancy frameworks are built around it. Conference presentations celebrate it. And yet, across the organisations I have observed over the past several years, it consistently produces the same outcome: a portfolio of technically interesting pilots, most of which never reach production, and a programme that struggles to demonstrate value beyond the initial proof-of-concept phase.

The problem is not execution. The problem is the approach itself.

2. Why the Use Case Model Fails

The use-case-first approach rests on an assumption that sounds reasonable but is, on examination, deeply flawed: that the right way to start an AI programme is to identify specific applications and then build them. This assumption fails for several interconnected reasons.

Use cases identified in workshops are disconnected from operational reality. A brainstorming workshop is, by design, an exercise in imagination. Participants are asked

to envision how AI could improve their part of the business, and they respond with ideas that are shaped by their understanding of what AI can do — an understanding that is, in most organisations, shallow and heavily influenced by vendor marketing and media coverage. The resulting use cases tend to be either too vague to implement ("use AI to improve customer experience") or too narrow to justify the investment required ("automate the formatting of monthly reports").

More fundamentally, workshops do not surface the operational complexity that determines whether a use case is viable. The claims processing manager who suggests an AI-assisted fraud detection model may not know — and the workshop format does not require her to know — that the data needed for that model sits across three legacy systems with no common identifier, that the regulatory framework requires human review of all fraud determinations, and that the current process is so poorly documented that no one can articulate the decision logic the model would need to replicate.

The scoring frameworks create a false precision. Ranking use cases on a matrix of feasibility and impact is a comforting exercise that produces confident-looking outputs. It is also, in most cases, an exercise in guesswork. The feasibility score depends on data readiness, technical complexity, and integration requirements — none of which can be reliably assessed without detailed investigation that the workshop format does not allow. The impact score depends on assumptions about adoption, behaviour change, and process redesign that are rarely tested. The resulting rankings feel rigorous but are built on sand.

The approach optimises for pilots, not for production. The use-case model naturally produces a portfolio of independent pilots, each scoped as a standalone proof of concept. This is efficient for demonstrating technical possibility — a data science team can build a working model for a well-scoped use case in a matter of weeks — but it systematically neglects everything required to move from pilot to production: data pipelines, model monitoring, integration with operational systems, change management, governance, and ongoing maintenance.

The result is the pattern now visible across sectors: organisations with impressive proof-of-concept portfolios and very little operational AI capability. The pilots work. The programme does not.

3. The Deeper Problem: Confusing Use Cases with Capabilities

Beneath the tactical failures of the use-case approach lies a more fundamental conceptual error. Organisations treat the use case as the unit of AI value — the thing to be identified, prioritised, built, and measured. But value from AI does not come from individual use cases. It comes from capabilities — the organisational infrastructure, skills, data assets, and operational practices that enable AI to be deployed, maintained, and trusted at scale.

The distinction matters because capabilities are shared across use cases, while use cases are, by definition, specific. An organisation that builds a robust data pipeline for one use case has created something that can support many others. An organisation that develops the skills to monitor model performance in production has a capability that transfers. An

organisation that establishes a governance framework for AI-assisted decisions has laid groundwork that every subsequent use case can build on.

The use-case-first approach inverts this logic. It treats each application as an independent problem to be solved, and it funds each pilot as a standalone investment. The result is duplication, fragmentation, and a chronic inability to build the shared foundations that would make subsequent use cases faster, cheaper, and more likely to succeed.

The use case is not the unit of AI value — the capability is, and capabilities cannot be discovered in a two-day workshop by people who have never built one.

4. What a Capability-First Approach Looks Like

If use-case hunting is the wrong starting point, what should replace it? The organisations that have made the most progress with AI — and they exist, though they are a minority — have started not by asking "where can we apply AI?" but by asking "what do we need to be able to do?"

This capability-first approach begins with an honest assessment of the organisation's current position. Not an assessment of which AI tools are available or which vendors offer the best platform, but an assessment of the foundational capabilities that any AI programme requires:

- **Data infrastructure.** Not whether the organisation has data — every organisation has data — but whether that data is accessible, governed, documented, and of sufficient quality to support machine learning. In most organisations, this assessment alone reveals years of underinvestment that no use-case portfolio can overcome.
- **Technical foundations.** Does the organisation have the engineering capability to build, deploy, and maintain AI models in production? Not in a sandbox, not as a proof of concept, but as an operational system with monitoring, versioning, and incident response? For most organisations outside the technology sector, the answer is no.
- **Organisational readiness.** Does the organisation have the change management capability to redesign workflows around AI? Does it have the governance structures to make decisions about AI-assisted processes? Do the people whose work will change understand what is coming and have the skills to adapt? Again, for most organisations, the honest answer is no.
- **Leadership understanding.** Do the leaders sponsoring the AI programme understand what they are sponsoring? Not the technology in detail, but the nature of the investment, the timeline for returns, and the organisational changes required? The pattern I observe most often is leaders who understand that AI is important but do not understand what it takes to make it work, and who therefore approve programmes that are structurally unable to deliver what they promise.

A capability-first approach takes these findings and uses them to design a programme that builds the foundations first, rather than launching use cases on top of foundations that do not exist. This is slower. It is less exciting. It produces fewer impressive

demonstrations in the early months. But it creates the conditions for AI to deliver sustained value rather than a sequence of isolated experiments.

5. Why Organisations Resist This Approach

If the capability-first approach is more likely to succeed, why do so few organisations adopt it? The reasons are predictable but worth naming.

Speed pressure. Leaders want visible results quickly, and use-case pilots deliver visible results — a working demonstration, a proof of concept, a compelling narrative for the board. Capability-building is invisible. Nobody gets promoted for building a data pipeline or establishing a model governance framework. The incentive structures inside organisations systematically favour the approach that looks impressive over the approach that works.

Vendor influence. The organisations that sell AI capability — platform vendors, consultancies, specialist firms — are structured around the use-case model. Their engagements are scoped as use-case implementations, their pricing is based on per-model or per-application fees, and their success metrics are defined by pilot completion rather than operational value. Shifting to a capability-first model would require vendors to sell differently, price differently, and measure differently. Few have an incentive to make that shift.

The difficulty of honest assessment. A capability-first approach requires the organisation to confront uncomfortable truths about its data, its skills, its processes, and its leadership. The use-case approach avoids this confrontation. It assumes the foundations are adequate and moves directly to the applications. By the time the foundational gaps become visible — during the transition from pilot to production — significant investment has already been committed, and the sunk cost makes it difficult to change direction.

6. The Programme That Nobody Wants to Run

The irony of the current moment is that most AI programme leaders know, at some level, that the use-case-first approach is flawed. They have seen the pilots stall. They have watched the proof-of-concept portfolio grow while operational AI capability remains static. They understand that the foundations are weak.

But the programme that would address these weaknesses — a deliberate, capability-building programme that prioritises data infrastructure, technical engineering, organisational readiness, and leadership development over flashy demonstrations — is a programme that is extraordinarily difficult to sponsor, fund, and sustain in most organisations. It is slow. It is unglamorous. Its early milestones are technical and operational rather than strategic. And it requires leaders to tell their boards that the path to AI value is longer and harder than the vendor presentations suggested.

That is the programme that works. It is also the programme that almost nobody wants to run. And until that changes, the ritual of use-case hunting will continue — producing workshops, spreadsheets, pilots, and very little of lasting value.