

Manuka



Beyond the Band-Aid:
How the 5% Turn
AI Pilots into Profit

Gareth Martin

Chief Executive Officer

**In 2023, AI was everyone's favourite toy.
In 2024, it became a board obsession.**

In 2025, reality hit.

An MIT study suggests 95% of generative AI projects are failing to deliver measurable returns, with only 5% of organisations managing to get anything into production at scale.¹

S&P Global also reported that 42% of companies now abandon most of their AI initiatives before they reach production, with an average of 46% of projects scrapped between proof of concept and rollout.²

Many leaders are now sitting in a trough of disillusionment, wondering whether AI was just a short-lived and very expensive phase. If that's you, you're not alone.

But here's the uncomfortable truth:

AI isn't the problem. Your approach is.

The next few years will be brutal for organisations that treat AI as a string of experiments or a quick fix for isolated issues. We call this the Band-Aid problem.

The 5% who win take a different approach: they move beyond Band-Aids to rebuild how their business works. They're the ones who'll benefit from the \$15.7 trillion that AI is expected to contribute to the global economy by 2030.³



**Want to join the 5%?
This is how.**

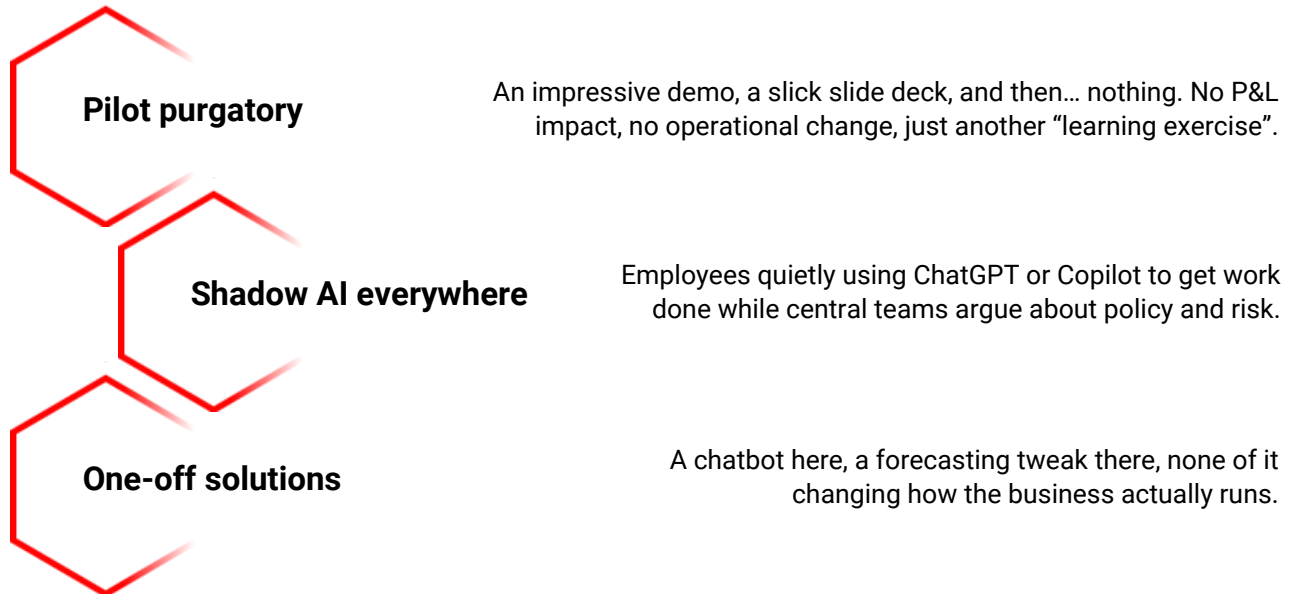
¹ https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf

² <https://www.ciodive.com/news/AI-project-fail-data-SPGlobal/742590/>

³ <https://holisticds.com/en/articulo/ai-economy/>

First, admit you've got a Band-Aid problem

Most enterprises we meet are stuck in some combination of:



This shows up as three symptoms:

1. **ROI not realised** – you’ve spent the money, you’ve hired the teams, but you can’t point to a single line in the P&L that moved.
2. **Risk delaying decision making** – fragmented data and unmanaged shadow AI make you nervous about scaling. Instead of enabling safe experimentation through governance and platforms, you’re trapped in ‘not data-ready’ loops and zero risk approval processes.
3. **Lack of trust from users** – you’ve invested 80% in tech and 20% in people. It should have been the other way round. Front-line teams don’t trust the tools. Middle managers don’t know how to manage with them. Adoption stalls.

If that sounds familiar, you’ve got a Band-Aid problem. You don’t need a new AI model. You need a new operating model.



Early warning sign: If “let’s use AI for something” is your starting point, or you choose a pilot in a small, low-risk back-office process “just to prove it works”, then stop. Immediately. You’re choosing the solution before you understand the problem, align on the goal / destination and evaluate the data.

Build a strategy engine, not a slide deck

Most “AI strategies” are 40 pages of ambition, four bullets of risk and no plan.


The 5% run a strategy engine – three parallel streams that work hand in hand to constantly connect vision, value and infrastructure:

Value realisation – follow the money



Keep refining your value chain so you’re always clear on where money is made, lost or at risk. Turn that into a living backlog of AI products and use a simple value vs complexity view to decide what to build now, what comes next and what never gets done.

Strategy & operating model – change how work gets done



Be honest about where you are today across people, process, data, tech and governance, then decide what changes to make to accelerate your ability to deploy AI strategically in the business. From there, define the roles you’ll need so someone is clearly accountable for making it real.

Architecture – build for scale, not just the next PoC



Start with a clear inventory of your current platforms, data stores and integration patterns, and design a target state that can handle agents as well as dashboards – interoperable, secure and scalable. Then decide what belongs on your core platform (in our case, Databricks) and what elements remain in other parts of the architecture. Design patterns should be agreed so you’re not slowed by architectural governance.

The output isn’t a vision statement. It’s a roadmap of capabilities: which data products you need, which AI products you’ll build, which processes you’ll re-engineer and in what order.

That’s the difference between “We’re doing AI” and “We’re building an intelligent enterprise”.

Now let’s put this into action.

Early warning sign: If your AI roadmap doesn’t obviously connect to the strategic objectives in your annual report, you don’t have a strategy. You have a wish list.





Value realisation – follow the money

The 5% don't start with "What can we do with AI?". They start with "Where does this business actually make and lose money?"

That's a value chain question. Answering it is the only way to identify how AI can truly transform your organisation. Try this 3-step verification process.

Step 1 – Map the value chain with brutal honesty

Bring together a small C-suite group (CEO, CFO, COO, maybe CCO/CIO) and a handful of people who run the processes – ops leads, finance, front-line managers. Sketch the end-to-end chain at a high level: from demand creation to revenue and retention, or from "citizen in need" to "case resolved", depending on your sector. **Then layer on where money and value move.**

Here is the method we use, feel free to use it and tweak it for your purposes or get in touch with our team and we can help you through this approach.

Where are the biggest costs, delays and failure rates?

Where are service metrics visibly broken?

Where are customers, citizens or partners clearly feeling pain?

Goals:

- Identify the highest ROI use case opportunity areas
- Down-select one high-priority use case
- Bring AI solution design to the line of business
- Produce a business case

Features:

- Duration: 2 weeks
- Activities: Meetings, workshops, prototype solution design
- Participants: Business area leads & SMEs

Outputs:

- Value chain map with several value "hotspots" identified
- One high-value use case high-level design & business case
- Easy next steps defined to move the product to prototype and MVP build



Stage 1

Value chain mapping

Stage 2

Use case identification & prioritisation, business case

	Current state & context gathering	Value chain mapping	Hotspot identification	Focus area deep dive	Use case design	Business case development
Objectives	Understand the current state of AI, data, technology, and strategic business priorities.	Map the business's value chain, considering whether to map the whole or a sub-section.	Pinpoint high-potential links in the value chain for process re-engineering or data & AI product development.	Assess a priority hotspot to understand current process, data, pain points, and AI applicability.	Develop a use case primer for the top use case, including high-level solution design.	Build the business case for the solution to ensure the ROI is understood and documented.
Audience	Key client stakeholders from IT, data, and business leadership.	Cross-functional relevant client stakeholders (e.g. operations, sales, HR).	Primary client stakeholders, business leadership, and data & AI strategy lead (intentionally small audience).	Process owners, product managers, and technical SMEs for selected area.	Same as previous step with additional validation by finance, business sponsors, and tech leadership.	Product owner, business sponsor, relevant finance team members, architecture review team.
Format	Several stakeholder interviews, review of documents, and summary of findings.	Collaborative workshop, co-development of an annotated value chain visualisation.	Small-format workshop, value v. complexity opportunity matrix, final shortlist of 3+ hotspots.	Deep-dive workshop, process mapping, high-level opportunity estimation, and solutioning.	Follow-up workshop, use case primer with expected benefits, costs, risks, approach, and success criteria.	One-to-one meetings and workshops to identify KPI and OKR impacts and technical complexities to estimate cost.
	Current state assessment playback		Pipeline of high-priority use cases playback			Business case playback

You're looking for hot paths through the business, places where a small improvement in speed, accuracy or consistency would have an outsized commercial or service impact. Those are the points where AI can explain, predict or optimise something that really matters.



Early warning sign: Mapping every detail of a complex organisation is a waste of time. Start at 30,000 feet. Identify the few functions that add the most value or cost for your business. Then zoom into those, and only those. Within a priority area such as new customer acquisition or claims management, you can create a more detailed map of how value is added, where it leaks, and where it touches other parts of the organisation.



Step 2 – Apply the four tests of a catalyst use case

From those hotspots, the 5% don't pick whatever sounds coolest. They run a simple, ruthless test:

- 1. Cross-team impact** – does this touch multiple teams, processes or customer journeys? If your AI only helps one person in one corner, it's fragile.
- 2. Leadership backing** – is there a senior owner who will defend this when priorities shift? No sponsor, no success.
- 3. Measurable outcome** – can you describe, in one sentence, what you'll improve and how you'll measure it? For example: "Reduce claims handling time by 30% while holding error rates flat." If you can't say it, don't start.
- 4. Foundations ready (enough)** – do you have the basic data, systems and governance to attempt this? If your source data is a spreadsheet on someone's desktop, fix that first. More on this later.

On one page, you should be able to answer these questions for each use case:

1. What problem are we solving, in plain business language?
2. What value is at stake – revenue, cost, risk, service – and what's the upside vs doing nothing?
3. What exactly will the AI do – explain, predict, optimise, generate, orchestrate?
4. What data do we need, from where, at what quality?
5. Who will use this, inside which workflow, through what interface?
6. How will we measure success?
7. What are the key risks, dependencies and milestones?

Only once a hotspot passes all four tests does it earn the right to become your catalyst use case – the one that proves AI can move the needle in your real business operation, not just in a pilot environment.




Early warning sign: Optimise with care. In value-chain mapping exercises, strategic opportunities often get lost in a long list of process niggles. It's tempting to chase the quick wins – the whack-a-mole fixes. Prioritise the golden nuggets: use cases that deliver dramatically higher value for the same level of effort.



Step 3 - Build, test and evolve your solutions

Once you've chosen your catalyst use case, this is how the 5% actually build and run the product. They don't treat AI like a one-off experiment, they treat it like a product with a lifecycle.


From identifying the right use case, to prototyping fast, to launching a Minimum Viable Product (MVP) with governance and a roadmap.



This isn't just about getting something live. It's about creating a repeatable path from idea to production. One that can be reused across teams and scaled across the business.



Early warning sign: Put it in front of real users and be blunt with the test: if you switched it off tomorrow, would anyone actually miss it? If the answer is no, kill it.



Strategy & operating model - change how work gets done

The most effective organisations design their operational structure for impact from day one, setting up the roles, rhythms and funding models that allow AI products to move fast and scale safely. Without this foundation, even the best use cases stall.

You need to be clear on the current state of your core pillars: people, operating model, processes, data, and governance, then decide what changes to make to accelerate your ability to deploy AI strategically in the business.

Let's focus on the first three - we'll come on to governance and data as the heart of your architecture next.



People & culture

Start with mindset. If your teams see AI as someone else's job, or innovation as a side hustle, you'll never get traction. The top performers build a culture of ownership and experimentation. Where leaders champion AI, and frontline teams feel empowered to solve problems with it. That means training, yes, but also storytelling, incentives and visible support from the top.



Operating model

Your AI capability needs a home. Whether it's centralised, hub-and-spoke or fully federated, the structure must reflect your ambition and appetite for scale. Mature organisations create cross-functional teams that can move from idea to production repeatedly, reusing patterns and components. Funding mechanisms should follow suit, enabling strategic initiatives to tap into shared AI resources without friction.



Processes

You need lightweight, well-defined processes that let teams build quickly without getting blocked. That includes clear governance pathways, reusable design patterns, and pre-approved architectural templates. Done right, these processes don't slow you down, they accelerate production by removing ambiguity and reducing risk.



So begin by assessing the maturity of your core pillars with these simple checklists:



People & culture

- ☐ Internal data & AI community
- ☐ Data-driven AI decision making
- ☐ Data & AI champions in the leadership team
- ☐ Accessible training & support



Operating model

- ☐ Formalised data & AI capability structure
- ☐ Effective funding mechanism process
- ☐ Centralised or federated best practice function



Processes

- ☐ Documented data & AI governance processes, followed by all
- ☐ Project or workload prioritisation mechanism
- ☐ Project / product prioritisation process
- ☐ Business engagement mechanism
- ☐ Standard methodologies for new projects, evolutions & day-to-day maintenance



Early warning sign: If every AI product feels like a one-off, or if teams are constantly reinventing the wheel, your operating model isn't working. You're not set up to scale, and your roadmap will cost a fortune to execute.



Architecture – build for scale, not just the next PoC

Now let's focus on the core of your architecture: data & governance.

Most AI roadmaps assume a single "AI platform" will emerge and run everything. That's unlikely.

AI will be found in workplace tools, content management systems, line-of-business apps, home-grown solutions and off-the-shelf products. It will appear in the ones you have, and you'll acquire new ones where it is core. And they'll talk to each other through cross-application, AI-powered processes.

An AI strategy that recognises this future state is essential to drive value and reduce risk. And it needs to be underpinned by an open, interoperable platform that can build home-grown solutions and orchestrate AI processes across multiple systems using shared data. In our world, that's Databricks.

Cross-application AI-powered processes



**AI-powered
workplace
tools**

**AI-powered
content
management
systems**

**AI-powered
business
applications**

**Home-grown
AI applications**

**Off-the-shelf
AI applications**

**AI vibe-coded
applications**




The first step to establishing an effective platform is accepting that no single vendor will own all your AI. Many vendors get greedy and will keep pushing you to “do it all in our stack”, but most enterprises don’t work that way.

Instead, you need a clear view of the ecosystem you’re building:

- The core platforms where data, intelligence and governance will live.
- The satellite systems – SaaS tools, operational applications, niche AI products – that will plug into that core.
- The cross-application processes that matter most and where AI agents will pass information, decisions and actions between systems.

Architecture is the strategy for that ecosystem: how everything coexists, shares data and supports end-to-end processes without creating chaos.

The intelligent enterprise emerges from repeatable capabilities you can apply across the value chain, with governance at the core and data as the foundation.



Early warning sign: If your architecture assumes AI will live in one place, or if every new use case requires a new build from scratch, you’re not ready to scale. You’re designing for the next PoC, not the next hundred.



Governance as a growth engine, not red tape

Good governance isn't the handbrake on AI, it's at the heart of AI that moves safely and at speed. It starts with the basics done properly. You need clear, unambiguous policies on what data can be used where, by whom and for which kinds of AI workloads. You need safe, sandboxed environments where teams can experiment with guardrails in place, not in the wild on unsanctioned tools. And you need auditability baked in: who accessed what, when, and which decisions were influenced by AI. If you can't answer those questions, you're not in control.

This is where platforms like Databricks, with Unity Catalog at the core, change the game. Instead of spraying sensitive data out to dozens of external APIs, Unity Catalog:

- Brings AI to the data inside a governed perimeter.
- Manages data, models and AI agents under a single set of access controls, lineage and audit trails.
- Enables security and compliance teams to see what's happening without slowing everything to a crawl.
- Supports multiple LLMs and model types behind one governance layer, allowing innovation without creating chaos.

The concerns about hallucination, security and model behaviour are legitimate. They are also solvable. The guardrails, evaluation techniques and monitoring tools exist. The real test is not whether something is possible; it's whether the use case is valuable enough to justify using them.

If you want AI to accelerate core processes and support meaningful decisions in your business, governance can't be an afterthought. It has to be built into how the platform works, not just written down in governance documents.

Stop using data readiness as an excuse.

Data readiness isn't optional – but it also isn't an excuse to stall before you've made any real architectural choices. In practice, some priority use cases will already have data available, even if it lives in operational systems rather than a pristine warehouse, and part of the job is to make that data available to the AI in the right way. For other use cases the data simply isn't there yet, which is your signal to change how and what you collect going forward, so future AI products don't hit the same wall.

AI-ready data means:

- Clear standards for what "good enough, structured enough, governed enough" looks like for your chosen use cases.
- Quality fixed as close to source as possible, not endlessly patched downstream.
- Shared visibility via a catalogue, so everyone knows which data is certified, experimental or off-limits.



So, how do you join **the 5%**?


If you take nothing else from this piece, take this: **Don't give up on AI. Give up on Band-Aid AI.**

Get the ball rolling on a 90-day action plan by kick starting each workstream with these practical actions:



Value realisation – Run a value chain and hotspot workshop

Get your C-suite and operators in a room. Map where value is created and lost. Pick 1–2 catalyst hot-spot opportunities using the four tests. Then stand up a real product pod, give them a clear outcome, decision rights and access to the right data and platform.



Strategy & operating model – Decide how you'll work, not just what you'll build

Define your AI operating model: roles, responsibilities, standards, review points. Don't leave it to chance.



Architecture – Invest in AI-ready data and governance on a unified platform

Whether you're already on Databricks or choosing your core, make sure data products, models and agents can live – and be governed, and choose an open data platform as a core part of your ecosystem.

Perhaps most importantly: make sure the story you tell about AI is one your board actually cares about. Start from the strategic objectives in your annual report. Show how specific AI-enabled changes in the value chain could unlock real money or impact – X million of cost avoided, Y million of revenue unlocked, Z improvement in critical outcomes. That's the story that justifies serious investment in platforms like Databricks, AI and the operating model around them.



At Manuka, we're pure-play Databricks specialists. We help organisations re-engineer business processes and build intelligent, agentic systems on a platform that can actually carry the weight.

AI isn't going away. The question is whether you spend the next three years applying Band-Aids – or **join the 5% and build something that lasts.**

**We're Manuka.
We build AI on Databricks.
Let's build together.**

📍 UK – Manchester – London – Liverpool

🌐 www.linkedin.com/company/manukaai-emea

📍 ANZ – Sydney – Auckland

🌐 linkedin.com/company/manukaaianz

🌐 www.manuka-ai.com