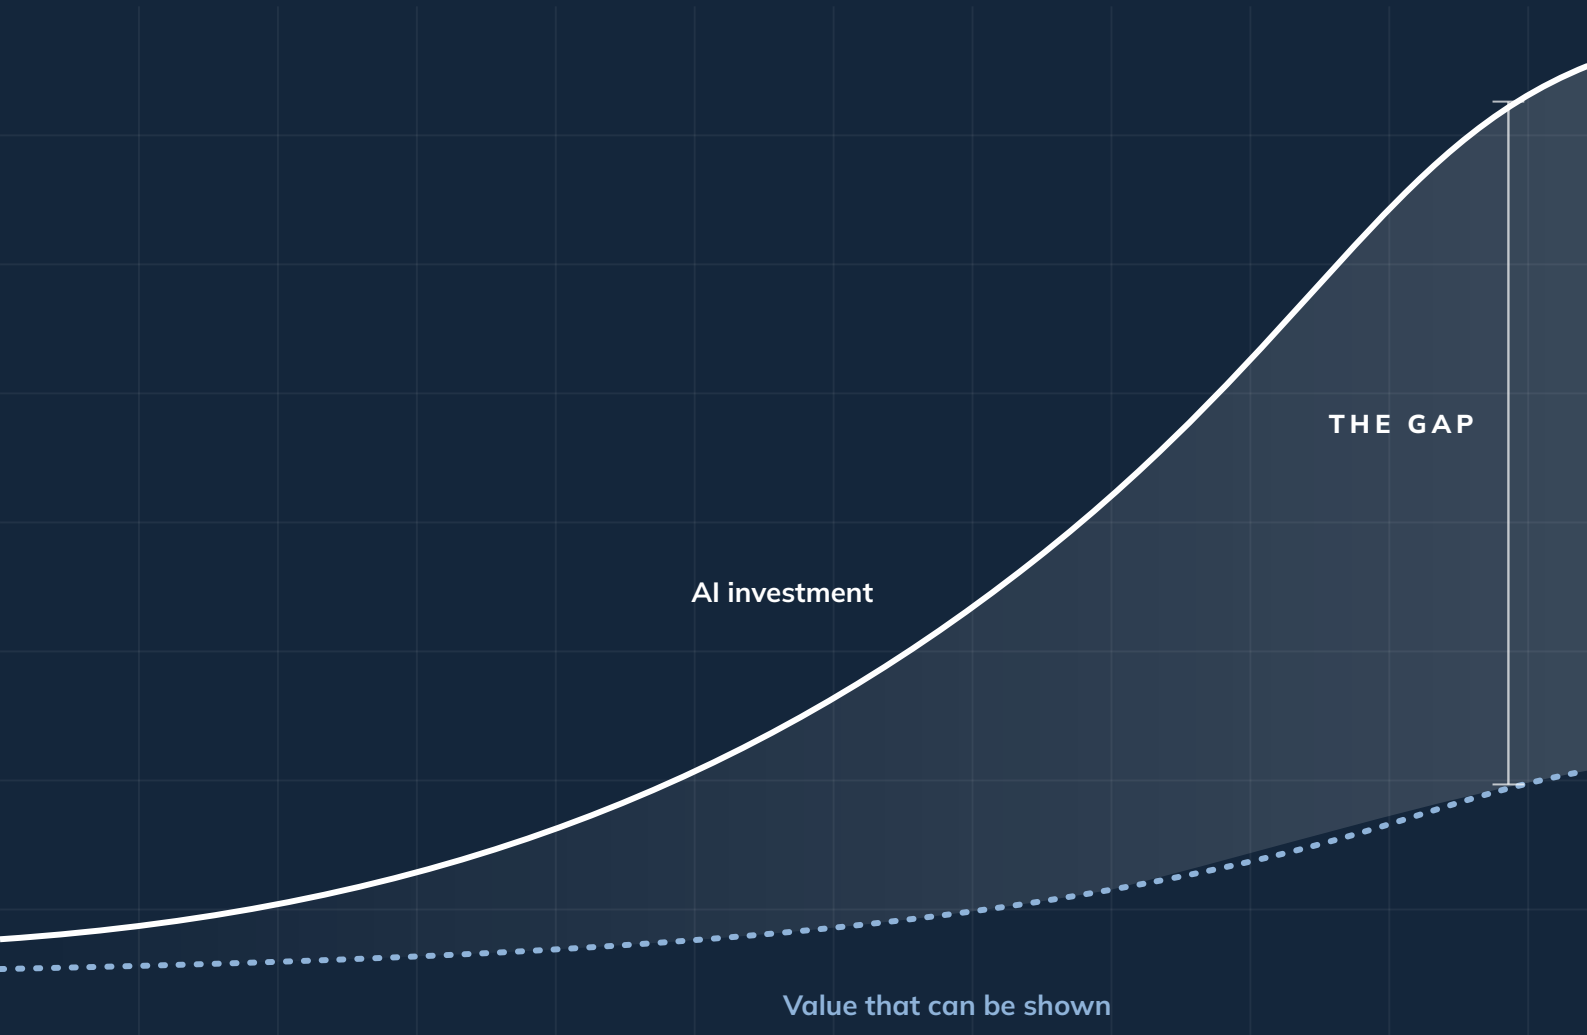
A DATAZUUM WHITE PAPER

NOVEMBER 2026

THE AI INVESTMENT GAP

Why organisations are funding AI initiatives without knowing whether they will deliver business value.



Samir Sharma
Founder & CEO, datazuum

DATAZUUM.COM

01 — THE INVESTMENT PARADOX

More AI activity has not meant more business value

IN BRIEF

1. AI spending and adoption keep rising. The share of executives reporting an enterprise-level financial impact from AI has barely moved.
2. The deeper problem is accountability. Strategy, investment approval, delivery, adoption and benefits are owned by different people, and the business outcome often belongs to nobody.
3. Business cases routinely treat released capacity as cash and leave out much of the cost of change.
4. Executive teams should stop managing AI initiatives and start managing the value those initiatives are meant to create.

Few boardrooms are short of AI activity. There are pilots in most functions, a pipeline of use cases, an AI council or steering group, and a budget line that has grown every year since 2023. Gartner forecasts worldwide AI spending of $2.67 trillion in 2026, almost 50% more than the year before[1]. Stanford's AI Index puts global corporate AI investment at $581.7 billion in 2025, more than double the previous year[2]. Inside companies, BCG estimates AI spending has reached 3.3% of revenue, roughly double the level of late 2025, with more than 80% of it now sitting outside the IT budget[3].

What executives report about the returns is more awkward. In McKinsey's 2026 survey, 44% of respondents say AI is scaling across their enterprise, yet the share reporting that AI has contributed to their organisation's EBIT is 37%, essentially unchanged from a year earlier. The share of “high performers”, who attribute at least 5% of EBIT to AI and describe its impact as significant, has stayed at about 6%[4]. In PwC's survey of 4,454 chief executives, 56% said their company had seen neither higher revenue nor lower costs from AI over the past 12 months; 12% reported both[5]. In a large NBER study, senior executives in seven countries estimated that AI had added around 0.2% a year to labour productivity growth over the past three years, and expect roughly three times that over the next three[6].

These are survey results: what executives say, not audited financial returns. Each study asks a different question of a different sample. When executives across several large surveys, asked different questions in different markets, describe a broadly similar pattern, it deserves attention. And the picture is not uniformly bleak. BCG's 2026 Applied AI Index reports that almost half of companies now say they generate value from AI, up sharply from 2025, although its published summary does not define what counts as value[3]. Deloitte found the share of organisations reporting revenue gains from AI rose from 20% to 28% during 2026[7]. Value is arriving. It is arriving unevenly, and many organisations still struggle to show whether their investments are delivering meaningful financial returns.

Nearly 9 in 10

of respondents report regular use of AI in at least one business function.

McKinsey Global Survey, 2026 · 1,719 respondents [4]

SELF-REPORTED

44%

of respondents say AI is scaling across their enterprise.

McKinsey Global Survey, 2026 · 1,719 respondents [4]

SELF-REPORTED

37%

of respondents report that AI has contributed to their organisation's EBIT, essentially unchanged from 2025.

McKinsey Global Survey, 2026 · 1,719 respondents [4]

SELF-REPORTED, NOT AUDITED

56%

of CEOs say their company has seen neither higher revenue nor lower costs from AI over the past 12 months. 12% report both.

PwC Global CEO Survey, January 2026 · 4,454 CEOs [5]

SELF-REPORTED

14%

of CEOs have clearly defined the P&L impact of all their AI initiatives.

BCG, July 2026 · 152 CEOs of companies with revenue of $500m+ [8]

SELF-REPORTED

0.2%

a year: executives' estimate of AI's contribution to labour productivity growth over the past three years.

NBER Working Paper 34836, 2026 · nearly 9,000 executives [6]

SELF-REPORTED ESTIMATE

Figure 1. Six findings from four independent studies

Each card reports what a different survey asked a different sample. The figures are not comparable with each other and should not be read as stages of a funnel. All are self-reported; none is an independently verified financial return.

Here is the paradox. Organisations have become much better at generating AI activity: finding opportunities, launching pilots, buying platforms, training people. They have not become proportionally better at deciding which investments deserve funding, what has to change in the business for those investments to pay back, or whether the expected value ever turned up. That is the AI investment gap.

Technology success is not business success

Part of the problem is language. Adoption, experimentation, deployment, transformation and financial return are routinely used as if they meant the same thing. They do not. Adoption means people have access to a tool. Experimentation means someone has tested whether it works. Deployment means it runs in production. Transformation means the way work is done, decisions are made or customers are served has changed. A return means the P&L, the balance sheet or a measurable operational outcome is different because of it. Each step is harder than the one before, and each tends to be owned by different people.

A model can be accurate, an assistant can be popular and a platform can be delivered on time and on budget while the business stays exactly where it was. That is a technology success. A business success is when something the organisation cares about has moved and you can show why. Which raises the question I would put to every leadership team spending serious money on AI: **are you managing AI initiatives, or are you managing the business value those initiatives are supposed to create?** They are very different jobs.

02 THE ACCOUNTABILITY GAP

Everyone delivered. Nobody owned the outcome.

Most explanations of disappointing AI returns focus on the arithmetic: weak business cases, poor ROI models, benefits that were never measured. Those problems are real, and the next sections deal with them. But they are symptoms of something more basic. **As an AI initiative moves through the organisation, accountability for its outcome breaks into pieces.** I call this the investment-to-value accountability gap.

Follow a typical initiative from idea to result. The ambition is set by the executive team. The investment is approved by a committee. The use case is shaped by a business sponsor and a data or AI team. The platforms and data are provided by the CIO and CDO. The solution is built by a programme team. The new way of working has to be adopted by line managers who were not in the room when it was designed. And the benefits are tracked, if at all, by finance or a programme office.

	LINK IN THE CHAIN	TYPICALLY OWNED BY	JUDGED ON	WHERE ACCOUNTABILITY FOR THE OUTCOME IS LOST
1	**Strategic ambition**	CEO and executive team	Setting direction; putting AI on the agenda	The ambition is “use AI” rather than an outcome with a measure attached.
2	**Investment decision**	Investment committee, CFO	Approving spend within the envelope	The committee approves a budget, not a benefit. Nobody is named to deliver the return.
3	**Use case definition**	Business sponsor with an AI or data team	A healthy pipeline of use cases	Use cases are defined around what the technology can do, not the decision that has to change.
4	**Data and technology enablement**	CIO, CDO, platform teams	Platforms live, data available, controls in place	Capability is built for the portfolio in general, with no line of sight to a specific outcome.
5	**Implementation**	Programme or delivery team	On time, on budget, in scope; go-live	The programme closes at go-live, and success is defined as “delivered”.
6	**Business adoption and process change**	Line managers, operations	Existing operational targets	Targets, budgets and roles are not changed, so the business inherits the change but not the benefit.
7	**Benefits realisation**	Finance or a PMO, often by default	Reporting what others submit	Nobody owns the number. Benefits are assumed in a budget or never measured.

Who owns the business outcome across all seven links? In many organisations, nobody does.

Figure 2. The accountability chain: where ownership of the outcome is lost

Typical patterns from practice; roles vary by organisation.

Each of those people can succeed on their own terms. The committee approved a sensible budget. The platform team delivered. The programme went live on time. Operations kept hitting its targets. Every stage reports green, and the business outcome the investment was approved for still does not appear, because it sits between the stages rather than inside any of them. Nobody failed. Nobody was accountable either.

The evidence suggests this is common rather than unusual. BCG found that only 14% of CEOs have clearly defined the P&L impact of all their AI initiatives, even though more than half see the link between AI and the P&L as a key barrier[8]. Deloitte found that when AI produces a bad outcome, 23% of organisations cannot say who is accountable[7]. If ownership is unclear when something goes wrong, it is rarely clearer when the question is whether something went right.

Where the gap opens matters. In my experience the two most expensive breaks are at the investment decision, where money is approved without anyone in the business being named to deliver the return, and at go-live, where the programme hands over to an operation whose targets, budgets and headcount plans have not changed. The first means nobody is on the hook for the benefit. The second means the benefit has nowhere to land.

What executive teams should do differently

1. **Name the outcome owner before the money is approved.**
   A business leader whose budget or targets will carry the benefit, not a sponsor who attends the steering committee. Ownership means owning what is measured, making sure the value shows up and reporting it to the board, good or bad. No owner, no funding.

2. **Approve outcomes and conditions, not just budgets.**
   The investment decision should record the expected benefit, the baseline, the conditions that must be true and who will put each one in place.

3. **Test value at every gate, not just progress.**
   Each stage gate should ask whether the initiative is still on a credible path to its outcome, not only whether that stage is complete.

4. **Change the operation when the capability goes live.**
   Budgets, headcount plans and targets should move to reflect the expected benefit. If they do not, the benefit has nowhere to land.

5. **Close the initiative when the benefit is reviewed, not when the system is live.**
   A benefits review six to twelve months after go-live, agreed with finance and reported to the executive, should be the real finish line.

03 THE USE CASE PROBLEM

Starting with the answer

Most AI portfolios I review were built from the technology outwards. A vendor demonstration, a competitor's announcement or an enthusiastic internal team produced an idea. The idea became a use case, and the use case acquired a business case to justify it. The order matters. Once a use case exists it gathers a constituency: people are assigned to it, budget is committed, a sponsor's name is attached. From that point the organisation stops asking whether it should solve this problem and starts asking how to make this solution succeed.

The symptoms are familiar. Generative AI automates a process nobody has mapped, so it speeds up steps that should have been removed. A model predicts churn, demand or fraud, but nobody has decided what anyone will do differently when the prediction changes, so the forecast is accurate and ignored. An assistant is rolled out with a target for licences activated rather than outcomes achieved. A dozen use cases run in parallel with no view of which matter most to the strategy, competing for the same engineers, experts and executive attention.

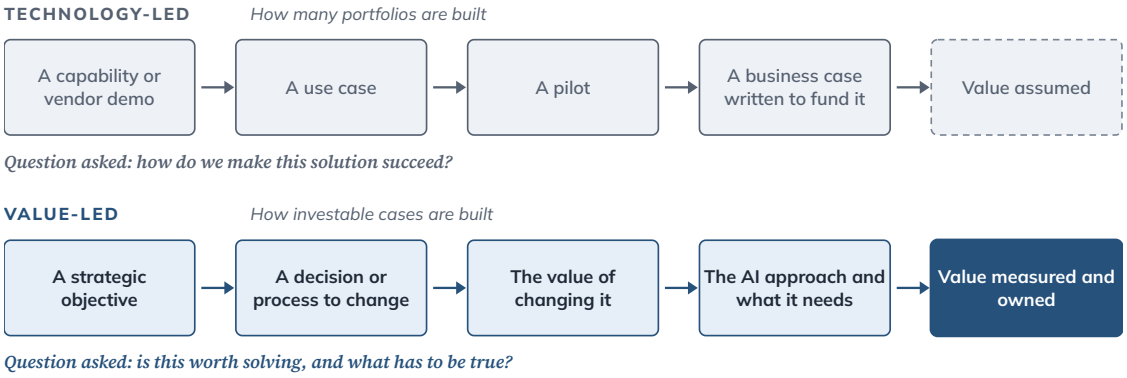


Figure 3. The order in which a use case is built decides the question the organisation ends up answering.

None of these initiatives is technically foolish. Many are technically impressive, which is exactly the trap. A technically interesting use case is not automatically a worthwhile investment, and the people best placed to judge the first are rarely the people best placed to judge the second. BCG's July 2026 research among CEOs of large companies found that nearly two-thirds are running AI pilots, but only 26% say they have embedded AI in a broader business transformation[8].

Start with the decision, not the model

The question that separates a value-led use case from a technology-led one is easy to ask and surprisingly hard to answer: *what decision, process or customer outcome will change, and who will change it?* Value is almost never created by a model on its own. It is created when a person or a system acts differently: a claim is settled faster, a price is set more accurately, a loan is approved or declined on better evidence, a store is replenished before it runs out. If you cannot name that change, you do not yet have a use case. You have a capability looking for a purpose.

Agentic AI makes this more pressing. Gartner predicts that more than 40% of agentic AI projects will be cancelled by the end of 2027 because of escalating costs, unclear business value or inadequate risk controls[9]. That is a forecast, not an observed failure rate, but its reasons are the classic ones: the use case was chosen before the value was understood.

04 THE BUSINESS CASE ILLUSION

A number nobody owns

When an AI initiative reaches an investment committee, it usually arrives with a business case. In my experience it is often the weakest document in the pack. Not because the people who wrote it lack skill, but because it was written to secure funding rather than to test whether funding was deserved. Gartner made the point plainly in its January 2026 forecast: "The improved predictability of ROI must occur before AI can truly be scaled up by the enterprise"[10].

The flaws repeat. Benefits rest on the best result from a pilot run by enthusiasts. There is no credible baseline, so nobody can later say what changed. Hours saved are converted straight into pounds, as if capacity released were cash in the bank. Technology costs are counted, while data remediation, integration, process redesign, training, change management and ongoing operation are left out or parked in someone else's budget. Benefits are owned by the programme team, who will have moved on by the time they are due. And the ROI cannot be substantiated afterwards, because nobody agreed how it would be measured.

The cost side is moving faster than most business cases assume. Deloitte found that 71% of organisations track their total AI inference and API spend, but only 44% track the cost per unit of inference, and it warns that AI running costs can quietly exceed the cost of the labour they replace[7]. In McKinsey's survey, about one in five respondents say AI operating costs are already limiting their organisation's use of AI[4]. Agentic systems, which call models repeatedly, make this harder to forecast.

Productivity is not profit

The evidence on individual productivity is real: a study of 5,179 customer-support agents found a 14% increase in issues resolved per hour, and 34% for the least experienced[11]. But when economists tracked AI chatbot use across 25,000 workers in Denmark, average reported time savings were 2.8% of working hours, and most of the time saved went into other tasks, some of them created by the AI itself[12]. Released time only becomes value when management decides what to do with it. That is why an investable business case separates five different kinds of benefit, which most AI business cases blur into one.

KIND OF BENEFIT	WHAT IT IS	HOW TO TREAT IT
Productivity improvement	The same work done in less time.	A measure of efficiency. Not money in itself.
Capacity released	Hours freed up, usually valued at the cost of the people who work them.	Notional until management decides what the time will be used for.
Cashable saving	Spend genuinely removed from today's cost base: posts taken out of the establishment, lower agency or supplier costs.	Shows up as lower cost in the P&L. Not replacing leavers counts only if the posts are removed and the budget is cut.
Avoided future expenditure	Spend that would have risen but does not, such as growth absorbed without recruiting.	Real economic value, but costs do not fall against today. Count it only if the increase was planned and is removed from the budget.
Additional revenue	New or retained income made possible by the change.	Valuable, but the hardest to attribute. Needs its own evidence.

Figure 4. Five kinds of benefit that should never be added together without saying which is which

That is the difference between a theoretical opportunity and an investable business case. The first says “this could save 30% of handling time”. The second says what will reduce, by how much, by when, what has to be true for that to happen, who is accountable and how everyone will know.

ILLUSTRATIVE EXAMPLE: FIGURES ARE HYPOTHETICAL, NOT INDUSTRY RESEARCH

An AI assistant for claims handlers

A general insurer with 400 claims handlers, each costing £45,000 a year fully loaded (£18m in total), considers an AI assistant that drafts correspondence, summarises files and suggests next actions. The vendor's pilot showed a 30% reduction in handling time. The headline case is compelling: 120 people's worth of capacity, presented as savings of £5.4m a year or £16.2m over three years, against £0.78m of licences and build.

ASSUMPTION	HEADLINE CASE	RE-EXAMINED
Time saved per active user	30% (vendor pilot)	12% (measured in a representative team)
Handlers actively using it	100%	60%
Benefit ramp-up, years 1–3	Full from day one	50% · 100% · 100%
Capacity released at full run rate	120 FTE · £5.4m a year	28.8 FTE · £1.3m a year
Treated as cashable saving	All of it	Half (14.4 FTE): posts removed and payroll budget cut as people leave
Treated as avoided expenditure	Not separated	The other half (14.4 FTE): only if growth is absorbed without budgeted hires
Additional revenue	Not quantified	Not quantified
Three-year cost	£0.78m: licences £0.43m, build £0.35m	£2.18m: adds data and integration £0.45m, change and training £0.20m, running costs £0.75m

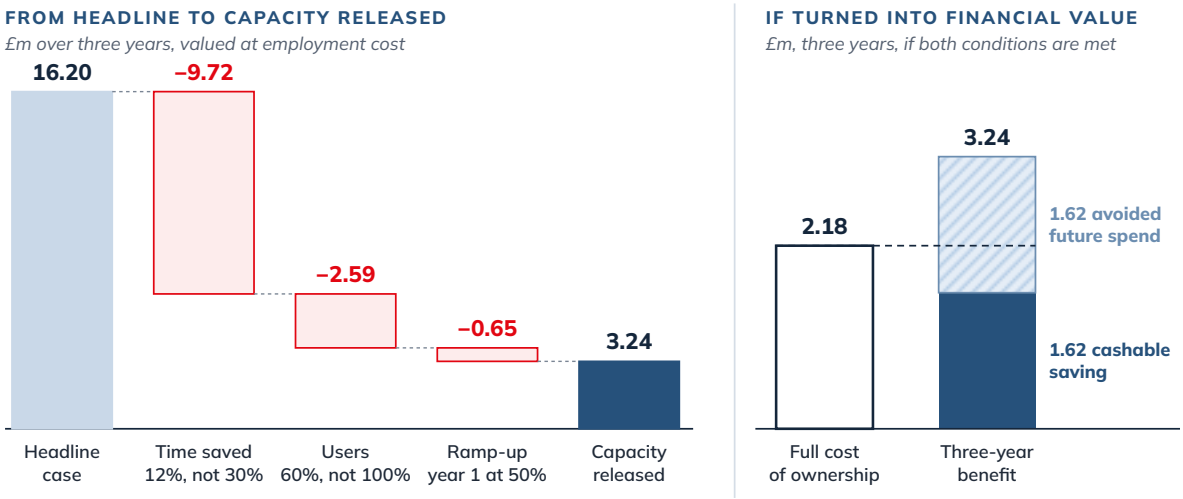


Figure 5. The illustrative business case over three years (£m). Illustrative figures only.

£3.24m

CAPACITY RELEASED

28.8 FTE of time, valued at employment cost. Not a financial benefit until someone decides what to do with it.

£1.62m

CASHABLE SAVING

Only if 14.4 posts are removed as people leave and the payroll budget falls against today's baseline.

£1.62m

AVOIDED FUTURE EXPENDITURE

Only if forecast growth is absorbed without 14.4 budgeted hires. Spend does not fall; it does not rise.

This is the distinction a CFO will press on. Not replacing leavers is not, by itself, a cashable saving: if the posts stay on the establishment and the budget is unchanged, the money has not left the cost base. The £1.62m counts as cashable only if the claims director agrees, before funding, to remove 14.4 posts and cut the payroll budget as people leave. Whether turnover allows that within the ramp-up period is itself an assumption to test.

On cashable savings alone, the case does not pay back: £1.62m against a full cost of £2.18m, a shortfall of £0.56m over three years. Adding the avoided expenditure takes the financial benefit to £3.24m and the net position to £1.06m, but only if the growth is forecast, the hires are in the approved budget and they are removed from it. If neither decision is taken, the investment releases capacity and nothing more.

What would need to be true. To pay back on cashable savings alone, effective capacity (time saved multiplied by adoption) would need to reach about 9.7% of the team's cost rather than 7.2%, for example 16% time saved at 60% adoption, with half the released capacity still removed from the establishment. Otherwise the case depends on the avoided recruitment being real and written into the budget. Either way, the investment decision is really a decision about how the claims operation will be run, and its owner should be the claims director, not the AI programme.

05 — THE MISSING CONNECTIONS

Ten things that have to be true at the same time

AI investment is usually discussed as a technology decision with a business case attached. It is better understood as a business decision with technology attached. A use case only delivers value when a whole chain of conditions holds at the same time, from a strategic objective worth pursuing through to risks that are understood and managed.

This is the thinking behind the Data & AI Strategy Canvas, which I developed after watching too many well-funded initiatives fail for reasons that were visible at the start. Its ten points sit beneath one anchor, the strategic objective, and fall into three groups: Strategic Framing, which defines the decision and the value; Strategic Enablement, which sets out what needs to be true for that value to be realised; and Strategic Readiness, which asks whether the organisation can sustain and scale it. The ten points are not a form to complete once. They are dependencies. A weakness in any one of them can cap the value of all the others, and it is rarely where the investment committee is looking.

An investment test, not a maturity score

Used properly, the Canvas is not a maturity assessment and not a template to fill in. It is how an executive team tests a specific investment before committing to it. For each of the ten points the questions are practical. Is this condition in place? If not, what would it take to put it in place, who will do it, and what will it cost? And once those costs and risks are included, does the investment still make sense?

This is where the Canvas changes the conversation. A gap in data quality (5) stops being a technical caveat and becomes a remediation cost in the value case (4). A missing owner (9) stops being a line on a risk register and becomes a decision the executive has to make. The result is an investment decision with four legitimate answers: proceed, fix the gaps first, redesign, or stop. Stopping early is not failure. It is the cheapest outcome available.

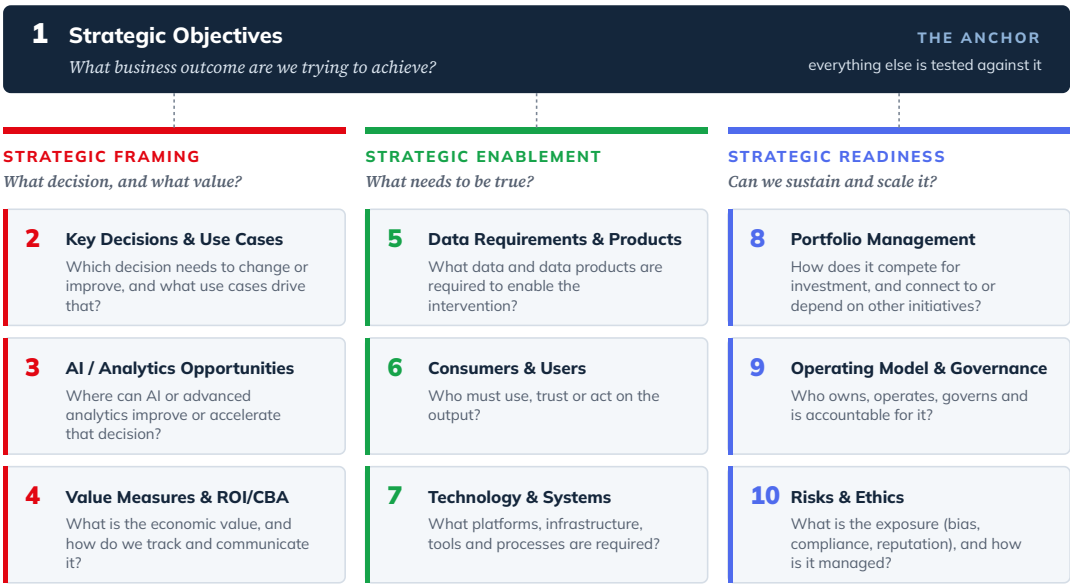


Figure 6. The ten points of the Data & AI Strategy Canvas.

Source: Sharma, S., The Strategy Canvas: A Field Guide for Data & AI. Canvas licensed CC BY-NC-ND 4.0.

How one weak link undoes a good model

Take a composite example drawn from patterns that recur in retail, rather than any single organisation. A retailer wants to protect gross margin **(1)** by improving end-of-season markdown decisions **(2)**. A machine learning model that sets markdown depth store by store **(3)** is valued at a 1.5-point margin improvement on clearance lines **(4)**. In testing, the model performs well.

Then the connections start to fail. Store stock records are not accurate enough, so the model prices stock that is not there **(5)**. Trading managers override recommendations they do not trust, and nobody tracks the overrides **(6)**. The pricing system can only hold one national price, so store-level recommendations are averaged away **(7)**. The initiative competes for the same engineers as the loyalty programme, which has a louder sponsor **(8)**. Nobody in the commercial team owns the margin target; the data team owns the model **(9)**. And nobody has asked how customers will react when the same product carries different prices in neighbouring towns **(10)**.

The model was right. The investment still failed. Every one of those weaknesses could have been found before a pound was spent. Had they been, the retailer would have seen that fixing stock accuracy and replacing the pricing engine might cost more than the margin the model was expected

to deliver. That is a perfectly good reason not to proceed, and a far cheaper one to discover before the build than after it.

06 — THE EXECUTION GAP

Go-live is the halfway point

A successful pilot answers a narrower question than many steering committees assume. It shows that the technology can work under controlled conditions, with willing users and prepared data. It does not show that it will work at scale, with every user, or that the business will change around it. The mistake is not running pilots. It is treating a pilot's technical success as evidence of business value.

It helps to separate three jobs that are often managed as one. **Designing the value** happens before implementation: deciding what will change, setting the baseline, naming the owner and confirming the conditions for value are in place. **Delivering the technology** is the build: integration, testing, security, go-live. **Realising the value** happens after implementation: adoption, process and role change, adjusted targets and budgets, and benefits measured against the baseline. Most organisations govern the second job well. The first and third are where the value is won or lost.

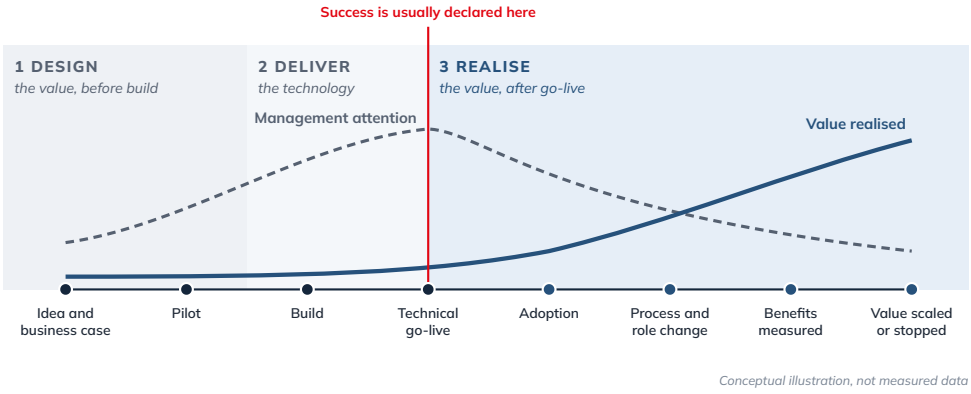


Figure 7. Technical delivery is one of three jobs, and most governance effort is spent on it.

The research describes the same pattern from different angles. In Deloitte's January 2026 survey, only a quarter of organisations said they had moved 40% or more of their AI pilots into production[13]. Its October update found that while 60% had redesigned jobs around AI to some degree, only 15% described the changes as extensive[7]. In IBM's survey of 2,000 CEOs, respondents said only a quarter of their AI initiatives had delivered the expected return, and 16% had scaled across the enterprise[14]. The usual causes are unglamorous: an unchanged process, users given access but not a reason, and no funded plan to scale.

Delivering a solution is not the same as changing the business. In McKinsey's survey, high performers are far more likely to have fundamentally redesigned workflows around AI[4]; BCG finds them about seven times more likely to redesign workflows and reshape the business end to end[8]. That is a business change programme with a technology component, and should be governed as one.

This is the difference between delivery assurance and value assurance. A programme can be green on schedule, cost and scope while becoming steadily less likely to deliver the value it was funded for. Declaring success at go-live moves attention and budget away just as value is supposed to start. The team disbands, the benefits are assumed, and a year later nobody can say if it worked.

07 — SEVEN QUESTIONS

Seven questions every executive team should ask

Ask them before funding, at every stage gate and twelve months after go-live. Vague answers mean the investment is not ready, however good the technology looks.

1 What specific business problem are we trying to solve?

WHY IT MATTERS
Without a defined problem, success cannot be defined and scope drifts towards what the technology does well.

A CONVINCING ANSWER
The problem in business terms, with its current cost or impact quantified, recognised by the leader who lives with it.

WARNING SIGN
The problem is described as “using AI in…”.

2 Which strategic objective does this investment support?

WHY IT MATTERS
Funding that cannot be traced to strategy competes on enthusiasm, not importance.

A CONVINCING ANSWER
A named objective from the strategy, and a view of how this ranks against other investments chasing the same objective.

WARNING SIGN
Every initiative supports every objective.

3 What business decision, process or customer outcome will change?

WHY IT MATTERS
Value is created when someone acts differently, not when a model is built.

A CONVINCING ANSWER
The specific decision or step, who makes it today, how it will be made in future, and what happens to the old way.

WARNING SIGN
The answer is a description of the technology.

4 How have we established the financial value of the opportunity?

WHY IT MATTERS
Optimistic, unbaselined benefits are among the most common reasons AI value fails to appear.

A CONVINCING ANSWER
A baseline, assumptions tested against evidence, cashable savings, avoided spend and capacity shown separately, and finance's agreement to the method.

WARNING SIGN
Hours saved are multiplied by salary and called a saving.

5 What data, technology, people and organisational capabilities are required?

WHY IT MATTERS
The cost of what has to be true is usually larger than the cost of the technology.

A CONVINCING ANSWER
A full cost of ownership, including data, integration, change, training and running costs, with each gap owned.

WARNING SIGN
The cost line is licences plus build.

6 Who is accountable for delivering the business benefits, not just the technology?

WHY IT MATTERS
Benefits owned by a programme team leave when the programme does.

A CONVINCING ANSWER
A named business leader whose targets include the benefit and who reports it to the executive, good or bad.

WARNING SIGN
The accountable owner is a committee, or the data team.

7 How will we know whether the investment has delivered the expected value?

WHY IT MATTERS
If value is not measured after go-live, it cannot be proven, improved or stopped.

A CONVINCING ANSWER
Agreed measures, a measurement date, a review cadence and a pre-agreed point to scale, change or stop.

WARNING SIGN
Success is defined as “live”.

USING THE QUESTIONS

Ask them in order. Questions 1–3 test whether there is a real business reason to invest. 4–5 test whether the case is honest. 6–7 test whether anyone will make sure the value arrives.

Ask the business owner, not the delivery team. If only the technologists can answer, the business has not yet bought in.

Write the answers down. Revisit them at every gate. When an answer changes, the investment case has changed.

08 — FROM AI ACTIVITY TO BUSINESS VALUE

Invest in outcomes, not initiatives

Nothing in this paper is an argument for spending less on AI. In McKinsey's survey, the high performers commit a larger share of their technology spending to AI, not a smaller one[4]. The argument is about how the money is allocated, owned and held to account.

Organisations do not need to pursue every AI opportunity. They need to identify the right ones, make informed investment decisions, create the conditions for success and stay accountable for the results. In practice that comes down to seven disciplines.

01	**Business-first thinking**	Every initiative starts from a business problem and a decision to change, not from a capability looking for a home.
02	**Disciplined prioritisation**	A small number of investments, ranked against strategy and funded properly, rather than a long list funded thinly.
03	**Evidence-based business cases**	Baselines, tested assumptions, full cost of ownership, and cashable savings, avoided spend and capacity shown separately.
04	**Clear ownership**	A business leader owns the outcome from investment decision to benefits review: what is measured, how the value is surfaced and how it is reported to the board.
05	**Readiness and operating model alignment**	Data, people, process, technology and governance are treated as conditions for value, assessed and costed before the money goes in.
06	**Benefits measurement**	Measures agreed upfront and tracked after go-live, with finance validating the results.
07	**Ongoing value assurance**	Regular reviews that ask whether each investment is still on a credible path to its value, and the willingness to stop the ones that are not.

The organisations that look back on this period with satisfaction will not be the ones that ran the most pilots or bought the most licences. They will be the ones that could tell their board, investment by investment, what each one was meant to change, what it cost, who owned it and what it actually returned.

Funding AI is a technology decision.
Funding outcomes is a leadership one.

SOURCES

References and notes

1. Gartner (16 September 2026). Gartner Forecasts Worldwide AI Spending to Grow 49.5% in 2026. Press release. FORECAST
   https://www.gartner.com/en/newsroom/press-releases/2026-09-16-gartner-forecasts-worldwide-ai-spending-to-grow-49-point-5-percent-in-2026
2. Stanford Institute for Human-Centered AI (April 2026). The 2026 AI Index Report, Economy chapter. COMPILED DATA
   https://hai.stanford.edu/ai-index/2026-ai-index-report/economy
3. Boston Consulting Group (30 September 2026). AI Is Starting to Pay Off: Almost 50% of Companies Now Generate Value with It. Press release on the Applied AI Index 2026; survey of 1,330 CxOs and senior leaders. SURVEY
   https://www.bcg.com/press/30september2026-ai-starting-to-pay-off-companies-generate-value
4. McKinsey & Company (2026). The state of AI in 2026: On the road to ROI. McKinsey Global Survey; 1,719 respondents in 97 countries; fieldwork 4 May–8 June 2026. SURVEY
   https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai
5. PwC (19 January 2026). 29th Annual Global CEO Survey: Leading through uncertainty in the age of AI. 4,454 CEOs in 95 countries and territories. SURVEY
   https://www.pwc.com/gx/en/ceo-survey/2026/pwc-ceo-survey-2026.pdf
6. Yotzov, I., Barrero, J. M., Bloom, N., Bunn, P., Davis, S. J. et al. (2026). Firm Data on AI. NBER Working Paper 34836, revised 2026; survey of nearly 9,000 senior executives in seven countries. SURVEY (EXECUTIVE ESTIMATES)
   https://www.nber.org/papers/w34836
7. Deloitte AI Institute (October 2026). The State of AI in the Enterprise. Built differently. 3,500 leaders in 25 countries; fieldwork May–June 2026. SURVEY
   https://www.deloitte.com/content/dam/assets-zone1/au/en/docs/about/2026/state-of-ai-in-the-enterprise-2026-report-built-differently.pdf
8. Boston Consulting Group (22 July 2026). CEOs Are Starting to See Value from AI. Now Comes Execution. 152 CEOs of companies with revenue of at least $500m. SURVEY
   https://www.bcg.com/press/22july2026-ceos-cost-revenue-benefits-ai-struggling-scale
9. Gartner (25 June 2025). Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027. Press release. PREDICTION
   https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027
10. Gartner (15 January 2026). Gartner Says Worldwide AI Spending Will Total $2.5 Trillion in 2026. Press release. FORECAST
    https://www.gartner.com/en/newsroom/press-releases/2026-1-15-gartner-says-worldwide-ai-spending-will-total-2-point-5-trillion-dollars-in-2026
11. Brynjolfsson, E., Li, D. and Raymond, L. (2025). Generative AI at Work. The Quarterly Journal of Economics, 140(2), 889–942. OBSERVED (FIELD STUDY)
    https://www.nber.org/papers/w31161
12. Humlum, A. and Vestergaard, E. (2025). Large Language Models, Small Labor Market Effects. NBER Working Paper 33777; revised 2026. OBSERVED (ADMINISTRATIVE DATA AND SURVEYS)
    https://www.nber.org/papers/w33777
13. Deloitte AI Institute (January 2026). The State of AI in the Enterprise 2026. 3,235 leaders in 24 countries. SURVEY
    https://www.deloitte.com/global/en/issues/generative-ai/state-of-ai-in-enterprise.html
14. IBM Institute for Business Value (May 2025). 2025 CEO Study. 2,000 CEOs. SURVEY
    https://newsroom.ibm.com/2025-05-06-ibm-study-ceos-double-down-on-ai-while-navigating-enterprise-hurdles

A NOTE ON THE EVIDENCE

Survey findings are not audited returns. Most figures in this paper come from surveys of executives. They record what respondents report or estimate about their own organisations, not independently verified financial results. Where the text says “report”, “say” or “estimate”, that is deliberate.

Different studies, different definitions. Samples, questions and definitions of “value”, “scaling” and “impact” differ between studies, so figures from different sources should not be added together or read as stages of a single funnel. Forecasts and predictions are labelled as such. Notes 11 and 12 are observed results from field and administrative data.

What is not used. This paper does not use the widely quoted claim that 95% of generative AI pilots deliver no return. It comes from a 2025 MIT NANDA report whose small interview sample and definitions have been widely questioned.

Illustrations. The claims-handling and retail examples are built to show how investment logic works. They are not client case studies, and their figures are not industry research.

© 2026 datazuum. All sources accessed November 2026. The Data & AI Strategy Canvas is licensed under CC BY-NC-ND 4.0.

CONTINUE THE CONVERSATION

Can your executive team say who owns the return on every AI investment?

Put your AI investments under the spotlight.

Samir Sharma and the datazuum team work with executive teams to assess whether their data and AI investments are strategically aligned, commercially justified and set up to deliver measurable business value. If your leadership team is struggling to answer the seven questions in this paper, let's discuss where to start.

Book a 60-minute AI Investment Review →

datazuum.com

ABOUT THE AUTHOR

Samir Sharma

Founder & CEO, datazuum

20+ years with executive teams and boards across five regions

Creator of the Data & AI Strategy Canvas

Author of The Strategy Canvas: A Field Guide for Data & AI

Host of The Data Strategy Show podcast

Samir Sharma is the Founder and CEO of datazuum. For more than 20 years he has worked with executive teams and boards across the UK, Europe, the US, Africa, Latin America and the Caribbean, in sectors including financial services, insurance, retail, consumer goods, media, technology, logistics, telecoms and the public sector. He helps leadership teams connect business objectives to data and AI investment, design the operating models and governance that make it work, and stay accountable for the value.

ABOUT DATAZUUM

datazuum is a data and AI strategy consultancy that helps organisations make their data and AI investments pay back: deciding what to invest in, putting the conditions for success in place, and tracking whether the value arrives.