

XAMUN POSITION PAPER NO. 4

# Your AI strategy is a list of tools. It should be a list of operations.

How a 50-to-500-person company should think about AI — without a CIO, a data lake, or a transformation program.

---

Arup Maity, Founder and CEO, Xamun Technologies

Version 1.0 · September 2026

**Summary.** Almost every mid-sized company now has an AI strategy document. Almost none of them changes how the company runs. The documents share a shape: a list of tools, a list of use cases, a training plan, a committee. What they lack is the only thing a strategy for a 50-to-500-person company needs — a short, ordered list of the operations the company is going to re-derive, and the numbers each one is expected to move.

This paper sets out how to write that list. It argues that the frameworks built for 5,000-headcount enterprises are actively harmful at this scale; that the correct unit of AI strategy is an operation, not a use case and not a department; that the first act of strategy is a constraint inventory, not a vendor evaluation; and that governance is not a policy document but a property of how a system is built. It closes with the questions a board or owner should ask before approving anything.

It follows Papers No. 1 (*You have AI. It isn't in your P&L.*) and No. 2 (*Don't renovate the work. Re-derive it.*), and assumes their argument: that generation must be placed inside a governed process, not patched onto the side of one.

## 1. The wrong question

Ask a hundred owners and managing directors of mid-sized companies what their AI strategy is and you will hear the same answer in a hundred accents: "We're rolling out Copilot," or "We've set up a working group," or "We've done some pilots." Press further and the strategy resolves into three things — a licence purchase, a list of use cases someone in marketing drafted, and a plan to train people.

None of that is a strategy. It is procurement with a mission statement.

The evidence that this approach fails is no longer anecdotal. MIT's NANDA study found that roughly 95% of enterprise AI pilots produce no measurable P&L impact. McKinsey's survey of AI adoption found only about 6% of companies attributing more than 5% of EBIT to AI, and found that the single strongest correlate of that impact was not model choice or spend but workflow redesign. Gartner expects more than 40% of agentic AI projects to be cancelled by the end of 2027. These are large-company numbers, gathered from companies with CIOs, data teams and budgets. If the well-resourced fail at this rate, the assumption that a 200-person company can succeed by copying their playbook at a smaller scale is not just optimistic. It is backwards.

The mid-sized company's problem is not that it lacks the resources of a large one. It is that it has been handed a question shaped for a large one. "What is our AI strategy?" invites an answer at the level of the whole company — a vision, a roadmap, a centre of excellence. At 5,000 people that abstraction is a necessity. At 200 it is a delay.

The right question is narrower and harder: *which operations in this company are we going to re-derive, in what order, and what number will each one move?* Everything else in this paper is a method for answering it.

*A strategy that cannot be written as an ordered list of operations, each with a baseline and an owner, is not yet a strategy. It is a mood.*

## 2. What a mid-sized company actually is

Before deciding what to do with AI it is worth being honest about the object it is being applied to. The 50-to-500-person company has a particular anatomy, and most of it is invisible in the org chart.

### Operational gravity

In a company of this size, the way work actually gets done is held in perhaps a dozen people. Not in the ERP, not in the process manual, not in the job descriptions. The credit controller who knows which customers to chase on which day. The operations manager who knows that a certain supplier's delivery note never matches its invoice and quietly corrects it. The scheduler who keeps the real plan in a spreadsheet because the system's plan is fiction by Tuesday.

This is operational gravity: the unrecorded weight of how things are really done. It is what makes the company work, and it is what makes the company hard to change. Any AI strategy that begins from the documented process is beginning from a fiction, and will fail in the same place every previous system implementation failed — the moment it meets the real work.

## The systems landscape

The typical estate is an ERP or accounting system that was implemented once and never quite finished; a CRM that is trusted by nobody except the person who set it up; a payroll system; twenty to two hundred spreadsheets, some of which are load-bearing; email; and a set of WhatsApp or Teams groups where the actual coordination happens. Integration between these is performed by humans copying values from one screen to another.

It is tempting to read this as a data problem to be solved first — clean the data, build a warehouse, then do AI. That sequence is how large companies have spent three years and produced nothing. At this scale the data is not clean because the operations that generate it are not clean. Fix the operation and the data follows. Attempt the reverse and the project will still be cleaning data when the budget ends.

## The people

There is no CIO. There may be an IT manager whose job is keeping laptops working and the ERP vendor honest. There is no data scientist and there will not be one; even if hired, they would have no data team to sit inside and no platform to build on. Decision-making is concentrated in an owner or a small executive group who are close enough to the operations to know where the pain is and far enough from the technology to be sold almost anything.

This is not a weakness to be apologised for. It is the mid-sized company's single greatest structural advantage in AI adoption, and the rest of this paper depends on it. In a large company, the person who knows where the operational waste is and the person who can authorise fixing it are five layers apart, and the distance is filled with committees. In a 150-person company they are often the same person, or two people who share a corridor. The mid-sized company can decide, baseline, build and switch on one re-derived operation in a quarter. A large one cannot.

*Readiness for AI is not adoption capacity. It is margin leverage: how much of the P&L sits inside operations that a re-derivation would move. On that measure the mid-sized company is not behind. It is the primary market.*

## 3. Three strategies that are not strategies

Most mid-sized companies have already adopted one of three postures. Each is worth naming so that it can be recognised and retired.

## The licence strategy

Buy a productivity-suite AI add-on for every seat, hold a training session, and let a thousand flowers bloom. The logic is that if every employee is ten percent more productive the company will be ten percent more productive.

It does not work that way. Individual productivity gains do not aggregate into operational gains unless the operation is redesigned to absorb them. The credit controller who drafts chasing emails twice as fast still waits the same number of days for the aged-debt report, still works from the same spreadsheet, still chases in the same order. The bottleneck is the operation's design, and no quantity of faster typing moves it. The licence strategy produces enthusiasm in month one, a handful of genuinely useful personal habits by month three, and no change in any number the board tracks by month twelve. It is the source of a large share of MIT's 95%.

This is not an argument against buying the licences. Buy them; they are cheap and people like them. It is an argument against calling the purchase a strategy.

## The pilot strategy

Choose a few use cases, run small experiments, "learn," and scale what works. It sounds prudent. Its failure mode is that nothing is ever scaled, because the pilots were designed to be safe rather than to be real. A pilot that runs alongside the existing process, that nobody is required to use, that touches no system of record and moves no number, generates a report that says the technology "shows promise" and a decision to run another pilot. This is pilot purgatory, and the way out is not a better pilot. It is to stop piloting and start replacing.

## The transformation strategy

Hire a consultancy, produce a three-year roadmap, establish a centre of excellence, build a data platform, and then deploy AI across the enterprise. This is the enterprise playbook scaled down, and the scaling down is exactly what breaks it. The roadmap costs more than the company's IT budget. The data platform is a two-year project preceding any benefit. The centre of excellence is one person with a title. The consultancy, meanwhile, is paid for the roadmap, not for any operation that moves — the consultant paradox in its purest form.

There is a version of this that works at 5,000 people because the scale amortises the overhead. At 200 people the overhead is the whole budget.

# 4. The unit of strategy is an operation

Discard the use case. A use case is a technology looking for a home: "AI for customer service," "AI for document summarisation," "AI for forecasting." It describes a capability, not a piece of work, and so it can never be baselined, never be owned by anyone in particular, and never be declared finished.

The unit of AI strategy in a mid-sized company is an *operation*: a bounded, recurring piece of work with a trigger, a set of decisions, a system of record it touches, and an output someone downstream depends on. Order-to-cash. Purchase-to-pay. Month-end close. Quote-to-contract. Inbound enquiry to booked job. Claim intake to settlement. Goods receipt to putaway. Every company has between eight and twenty of these. They are the load-bearing walls.

An operation has properties a use case never has:

- **It is countable.** It happens a known number of times per week or month. Orders, closes, quotes, claims, shipments.
- **It has a cycle time.** From trigger to completed output there is a measurable elapsed time, and the company already knows, roughly, what it is.
- **It has an error rate.** A proportion of instances come back, get corrected, get escalated, or get quietly fixed by the person who carries the operational gravity.
- **It has an owner.** Somebody is blamed when it goes wrong. That person exists today, and their name is known.
- **It touches a system of record.** Which means the re-derived version must too, or it is a parallel process — and a parallel process is a pilot wearing a disguise.

Strategy, then, is the act of choosing which of these operations to re-derive, in what order. Paper No. 2 covers what re-derivation involves in detail — the three design moves of inverting the trigger, going exception-first, and placing every decision — and this paper will not repeat it. The point here is prior: the strategy is the selection and the sequence, and the strategy document is a table with perhaps six rows.

*A company that has re-derived three operations in eighteen months has an AI strategy that is working. A company with a forty-page roadmap and a working group has a document.*

## 5. The first act: a constraint inventory

The first thing a mid-sized company should do is not evaluate vendors, not run a pilot, and not buy licences. It is to build a constraint inventory: a plain list of the company's recurring operations, with the friction in each written down honestly.

This takes two to three weeks and should be done by the people who carry the operational gravity, not by the IT manager and not by a consultant. The owner or managing director must be in the room for at least part of it, because the inventory will surface things nobody has written down before, and some of them will be uncomfortable.

### What goes in it

For each recurring operation, record the following. Approximate numbers are fine at this stage; the purpose is comparison, not precision.

| Field                  | What it captures                                                                                                                    |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| Operation              | The bounded piece of work, named the way people actually refer to it.                                                               |
| Trigger                | What starts an instance — an email, a system event, a calendar date, a phone call, somebody remembering.                            |
| Volume                 | Instances per week or month.                                                                                                        |
| Cycle time             | Trigger to completed output, typical and worst case.                                                                                |
| Handoffs               | How many people touch an instance, and how many times it waits in someone's inbox.                                                  |
| Decisions              | The judgement points inside the operation: approve, reject, price, route, escalate, classify.                                       |
| Error and rework       | What proportion comes back, is corrected, or is silently fixed.                                                                     |
| Systems                | Which systems of record are read or written, and where the spreadsheets sit.                                                        |
| Owner                  | The person accountable today.                                                                                                       |
| Constraint             | In one sentence: what makes this slow, wrong, or expensive.                                                                         |
| Cost of the constraint | The best honest guess at what the friction costs per year — in margin, in lost revenue, in headcount that exists only to absorb it. |

The last row is the one that matters and the one people resist filling in. It is almost always underestimated, because the cost of a constraint is not the salary of the people working inside it. It is the orders not taken because quoting is slow, the discounts given because invoices were late, the customers who left because the enquiry sat for four days, the working capital tied up in a close that takes fifteen days instead of five. A rough number, argued over for twenty minutes, is worth more than a precise one nobody believes.

## What it produces

The inventory produces two things. First, a ranked list: the operations whose constraint costs the most, sorted. Second, a shared vocabulary. Once the executive group has agreed that quote-to-contract takes nine days and costs the company a fifth of its pipeline, the conversation about AI stops being about technology and starts being about that nine days. That shift in conversation is the actual beginning of the strategy.

## 6. Choosing the first operation

The inventory ranks operations by the cost of their constraint. The first operation to re-derive is not automatically the one at the top. It is the highest-ranked one that passes four tests.

1. **Countable.** The operation produces a unit that can be metered — a completed order, a closed month, an issued quote, a settled claim. If the output cannot be counted, the improvement cannot be measured, and an unmeasured operation will slide back into pilot purgatory.
2. **Constrained.** The friction is structural, not personal. If the operation is slow because one individual is slow, that is a management problem, not a re-derivation candidate. If it is slow because every instance waits for three approvals that are given identically ninety-five percent of the time, that is structural.
3. **Reversible.** The company can switch the re-derived operation off and return to the old one within a day if something goes wrong. Operations that touch cash, regulatory filings, or safety should not be first — not because they cannot be re-derived, but because the company has not yet earned the confidence to leave them running unattended. Reversibility is what lets the company move fast; irreversibility is what forces caution.
4. **Owned.** A specific operator — not IT, not a project manager — will own the re-derived operation and be measured on the number it is meant to move. If nobody wants it, it will not be used, and an unused system is indistinguishable from no system.

The operation that scores highest on cost and passes all four tests is the first one. In most mid-sized companies it turns out to be something unglamorous: order entry, quotation, inbound enquiry handling, supplier invoice processing, a scheduling operation. It is rarely the thing the executive group was excited about when the AI conversation started. That is expected. Excitement clusters around visible outputs; cost clusters around invisible friction.

## A word on the operations not to start with

Two categories of operation are consistently proposed first and should consistently be deferred.

The first is customer-facing conversation — chatbots, AI on the website, AI answering the phone. The reasoning is that this is where the customer is. But it is also where the language model is most exposed, least governed, and least connected to the operation that actually fulfils the customer's need. A chatbot that can answer questions about delivery times but cannot see the order system is a source of confident misinformation. Put the operation right first; put a conversational surface on it later, if at all.

The second is "insight" — dashboards, forecasts, reports. These feel strategic and are low-risk, which is exactly why they produce nothing. A forecast that nobody's operation is wired to act on is a decoration. Re-derive the operation that would consume the forecast, and the forecast becomes worth building.

## 7. Baseline before build

Before anything is built, before a vendor is chosen, the first operation is measured. Three numbers, at minimum, taken over four to eight weeks of normal running:

- **Cycle time** — trigger to completed output, as a distribution, not an average. The tail matters more than the mean.
- **Error rate** — the proportion of instances that are returned, corrected, escalated, or fixed by hand after completion.
- **Throughput per person** — completed instances per week divided by the people whose time is inside the operation.

These are the numbers the re-derived operation will be judged against, and they are the only numbers that should appear in any business case. Series readers will recognise the stance from Paper No. 1: no projected figures, no "AI will save 30% of costs," no ROI model built on a vendor's benchmark. Projections are how projects get approved and how they get quietly abandoned. A baseline is how a project is held to account.

The discipline has a second benefit that is easy to underestimate. Measuring the operation for six weeks tells the company things about it that nobody knew. The cycle time turns out to be bimodal — half the instances take a day and half take two weeks, and the two-week ones share a cause. The error rate is concentrated in one customer segment. The throughput is bounded not by the people but by a weekly batch run in a system nobody has looked at in years. Each of these findings reshapes the re-derivation before a line of it has been designed.

*If the baseline cannot be measured, the operation is not ready. If the vendor will not commit to the baseline, the vendor is not ready.*

## 8. Governance is a property of the build, not a policy

Every AI strategy document contains a section on governance. It usually describes a committee, an acceptable-use policy, and a commitment to "responsible AI." None of that is governance. It is the paperwork that a company produces when it does not yet know what governance would mean operationally.

Governance in an operation means one thing: the operation can be left running unattended, and when something goes wrong, the company can see what happened, why, and who or what decided it. This is the gap Paper No. 1 called governance blindness — the distance between an AI producing an answer and an AI being trusted to run. Closing it is a design problem, and it has a known shape.

### Two minds, one memory

The shape is that every re-derived operation is built from two separate components, and they are not equal.

The first is a **deterministic mind**: rules, entitlements, calculations, thresholds, compliance checks, and an audit log. It runs first on every instance and it has the final say. It does not generate; it



decides. If an instance violates a rule, the deterministic mind stops it, regardless of what any language model thinks. Its logic is readable by a human, testable in advance, and identical every time.

The second is a **governed language model**: reading documents, extracting fields, classifying, drafting, summarising, explaining. It is used for the things language models are good at, and it is never the last word. Its output is proposed, not enacted; the deterministic mind checks it before anything reaches a system of record.

Between them sits **one graph memory**: the entities, obligations, precedents and lineage of the operation, read and written by both minds. It is what allows the operation to know that this supplier has a standing credit hold, that this customer's contract has a non-standard payment term, that this decision was made last quarter for this reason. It is workflow-scoped — it holds what the operation needs, not everything the company has ever recorded.

The user touches neither mind. There is no prompt and no chat window. There is a purpose-built application that does the operation, surfaces a human at a genuine judgement fork, and logs everything else.

## Why this is the governance question

A company that has built its first operation this way does not need a responsible-AI policy to answer the board's question about risk. It can show the rules, show the log, and show the forks where a human was called. The regulator, the auditor and the insurer can be shown the same things. Governance was never going to come from a document. It comes from an operation that cannot act outside its rules and cannot hide what it did.

The alternative — a language model with a prompt, connected to a system of record, with a policy document asking people to use it responsibly — is not governable, and the mid-sized company should decline it however it is dressed.

## The regulatory floor

For companies operating in the Gulf specifically, the binding surface today is narrower than the marketing suggests. Data protection law governs automated decision-making, impact assessment and data protection roles; the UAE's AI ethics principles are soft law; sector regulators — notably in financial services — have their own AI guidance. There is no binding UAE AI Act and no self-assessment deadline, and any vendor asserting one should be asked for the citation. What does exist is a policy signal: Dubai's May 2026 direction to the private sector on agentic AI adoption, with a two-year horizon, and the federal target of half of government services delivered through AI agents by 2028. The direction of travel is unambiguous. The obligations are, for now, mostly the ones that already applied to any automated system handling personal data.

The practical consequence is that an operation built on the two-minds shape already meets the floor: decisions are explainable because the deterministic mind made them; data flows are logged

because the graph records them; a human is in the loop where judgement is required because the design put one there. Compliance falls out of good construction. It does not need to be retrofitted.

## 9. Data: you do not need a lake

The enterprise playbook says: get your data house in order first. Build the warehouse, establish master data, then apply AI. The mid-sized company that follows this advice will spend eighteen months and its whole discretionary budget on a platform, and have no operation to show for it.

The correct scope of data work is the operation being re-derived. If the first operation is supplier invoice processing, the data that matters is supplier master, purchase orders, goods receipts, the invoice documents themselves, and the approval history. That is a bounded set. It lives in perhaps three systems and a shared drive. Extracting it, cleaning it to the standard that operation needs, and holding it in a workflow-scoped graph is a matter of weeks, not quarters.

The second operation will need some of the same data and some new data. Its graph extends the first. By the fourth operation the company has, almost by accident, the connected operational memory that the data-platform project was supposed to deliver — except that every node in it exists because an operation needed it, and every one of those operations is already running.

Data quality follows the same logic. The reason the supplier master is dirty is that nothing depended on it being clean. Once an operation runs unattended against it, and stops with an exception every time a record is wrong, the master gets cleaned — by the operation, one exception at a time, in the course of doing real work. This is the only data-cleaning programme that has ever finished.

### On sovereignty and residency

Mid-sized companies in regulated sectors or government supply chains will be asked where their data goes. This is a legitimate question and it should be settled at the design stage, not discovered at procurement. The two-minds shape makes it tractable: the deterministic mind and the graph run wherever the company's data must reside; the language model component can run on open-weight models hosted in-region for most tasks, with frontier models reserved for the narrow set of tasks where output quality demands them. A company should ask its vendor which tasks go where, and should expect a specific answer.

## 10. People: who owns this

The single most reliable predictor of whether a re-derived operation survives its first year is who owns it. The consistent failure mode is IT ownership: the project is given to the IT manager because it involves technology, the IT manager delivers something that works technically, and the

operation's actual staff continue doing the work the old way because nobody asked them to stop and nobody is measuring whether they did.

### **The operator-owner**

Every re-derived operation has an operator-owner: the person who runs that operation today, or their manager, whose performance is measured on the baseline number moving. They specify what "correct" means. They sit in the exception queue in the first weeks. They decide when the autonomy level goes up. They are the person who kills the parallel process — who says, on a specific date, that the old way is no longer permitted.

This is a real job, and for the first two or three operations it will consume a significant fraction of that person's time. That is the investment. It is smaller than a data platform and much larger than a licence purchase, and the strategy should say so explicitly.

### **The executive sponsor**

Above the operator-owner sits an executive — the owner, the managing director, the COO — whose role is not to run the project but to protect it. Protection means: refusing to let the operation be scoped down to a safe pilot; refusing to let the parallel process run indefinitely; refusing to let the baseline be replaced with a projection; and refusing to let the second operation start until the first is running unattended. This is a small amount of work and an enormous amount of willpower.

### **Everyone else**

The people inside the operation — the ones with the operational gravity — are not being replaced by the re-derivation. They are being moved from doing every instance to handling the exceptions and improving the rules. This should be said plainly and early, because if it is not, they will assume the opposite and will, reasonably, protect themselves by ensuring the new system does not work. The company that wants its operators to make the re-derived operation succeed must make it obviously in their interest.

What about headcount? The honest answer is that the first re-derived operation rarely reduces it, because the people freed from routine instances are absorbed into the exception queue, the next operation's inventory, and the work that was never being done because everyone was too busy. The third and fourth operations are where the company discovers that it can grow without hiring in proportion. Headcount-shaped thinking — designing the re-derivation around who will be removed — produces bad designs and worse morale. Design around the number, and let the org chart follow.

### **Skills**

The mid-sized company does not need to hire AI engineers and should not try. It needs three things it can develop internally: operators who can describe a decision precisely enough to write a

rule for it; a manager who can read an exception log and see a pattern; and an executive who can hold a baseline. None of these are technical skills. All of them are rarer than they sound.

## 11. Vendors and the three doors

At some point the company will engage a vendor, because it cannot build this alone and should not try. There are, broadly, three doors, and the strategy should say which one the company is walking through and why.

### Door one: the SaaS with AI in it

The company's existing ERP, CRM or vertical SaaS vendor adds AI features. This is the path of least resistance and it is legitimate for some operations. Its limitation is that the vendor's AI is designed for the vendor's generic process, not the company's actual one, and cannot reach across the boundary into the other systems and spreadsheets where the operation really lives. For an operation that sits entirely inside one system, take this door. For an operation that crosses systems — which is most of the expensive ones — it will not be enough.

### Door two: build your own, or have someone build it by the hour

The company hires developers, or a development firm, and builds. This is the traditional path and it fails in the traditional way: the company must write the scope before it understands the problem; the estimate is wrong; the scope changes; the budget and the timeline stretch; and the vendor is paid for the effort regardless of whether the operation's number moved. The man-day is the unit of sale, and the man-day is indifferent to outcome. Companies with genuine internal engineering capability can take this door for some operations. Companies without it should not, and the absence of internal capability is precisely the mid-sized company's condition.

### Door three: pay for the outcome

The company engages a partner who is paid when the operation runs and judged on whether the baseline number moved. The partner does the constraint work with the company, builds on the two-minds shape, and takes the risk that the scope was wrong. Pricing is per unit of the countable output, or a fixed price with no charge for scope changes, because the partner has accepted that the problem, not the specification, is what it has been engaged to solve.

This door is not always available and not always appropriate. It requires a countable output, which is why the four tests in section 6 matter. It requires a partner willing to be forward-deployed inside the operation rather than sitting behind a ticket queue. And it requires the company to accept that the partner will push back on scope, because the partner's economics depend on the operation actually working.

### How to read a proposal

Whichever door, the company should read every vendor proposal against the same six questions:

- Does it name the operation, or does it name a capability?
- Does it commit to the baseline the company measured, or does it substitute a projection?
- Does it describe what the deterministic component does and what the language model does — separately?
- Does it say where a human is called, and on what condition?
- Does it say what happens to the old process, and when?
- Does it say what the vendor is paid for — effort, or the number moving?

A proposal that fails three of these is a licence strategy or a pilot strategy in a different font.

## 12. Sequence and cadence

The strategy is the ordered list. The cadence is how the company works through it. The recommended rhythm for a company of this size is one operation per quarter, staged as follows.

| Weeks   | Activity                                                                                                                                                                               |
|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1–3     | Constraint inventory (first quarter only; thereafter a two-day refresh). Select the operation against the four tests. Name the operator-owner.                                         |
| 3–8     | Baseline. Measure cycle time, error rate and throughput on the live operation. Extract and scope the data. Design the re-derivation: trigger, decisions, rules, exceptions, forks.     |
| 8–12    | Build and switch on at the lowest autonomy level: the system proposes, the operator disposes. Every decision is logged and reviewed.                                                   |
| 12–13   | Raise autonomy for the reversible decision classes. Kill the parallel process on a named date. Begin the next quarter's inventory.                                                     |
| Ongoing | The operator-owner reviews the exception log weekly and promotes decision classes to unattended as the log earns it. The baseline number is reported monthly to the executive sponsor. |

### Staged autonomy, paced by reversibility

The re-derived operation does not go from human-does-everything to machine-does-everything in a step. Each class of decision inside the operation is promoted separately — from proposed-and-reviewed, to enacted-and-logged, to unattended — and the pace of promotion is set by how reversible the decision is and how clean its log has been. A classification decision that has been right for four hundred consecutive instances gets promoted. A credit decision gets promoted more slowly, or never. The company sets the pace; the design makes the pace visible.

### The second operation is easier

The company that has re-derived one operation properly has learned how to run a constraint inventory, how to hold a baseline, how to read an exception log and how to kill a parallel process. It has a graph with real entities in it. It has an operator who has done this once. The second operation takes less time and less willpower than the first. By the fourth, the company has a method, and the method is the strategy.

## 13. What the board should ask

The owner, board or investment committee that is asked to approve an AI initiative should ask the following. If the answers are not specific, the initiative is not ready.

1. Which operation? Name it the way the people who do it name it.
2. What is its baseline — cycle time, error rate, throughput — and over what period was it measured?
3. What is the cost of its constraint, and how was that estimated?
4. Who is the operator-owner, and what are they measured on?
5. What does the deterministic component decide, and what does the language model produce? Show them separately.
6. Where does a human get called, and on what condition?
7. On what date does the old process stop being permitted?
8. How is the vendor paid — for effort, or for the number moving?
9. If we switch it off on day thirty, what does it cost us and how long does it take to go back?
10. What is the second operation, and what did we learn from the first that changed our view of it?

Ten questions. A page. If they can be answered, the company has a strategy. If they cannot, it has a list of tools, and it should not yet spend.

## 14. Closing

The mid-sized company has been told for three years that it is not ready for AI: too small, too messy, too under-resourced. The evidence says the opposite. The large companies with the resources are failing at rates above ninety percent, and they are failing because they are too far from their own operations to re-derive them. The mid-sized company is not too far. It is close enough to do this in a quarter.

What it needs is not more tools and not a bigger roadmap. It needs a list of six operations, ranked by the cost of their friction, each with a baseline and an owner; a refusal to pilot; a refusal to project; and a construction discipline that makes governance a property of the build rather than a paragraph in a policy. That is the whole of the strategy. It fits on a page, and it should.

*Strategy is not what the company says about AI. It is which operations it re-derives, in what order, and whether the numbers moved.*

## About this series

This is the fourth paper in a series from Xamun Technologies on the conditions under which AI shows up in a company's P&L. Paper No. 1, *You have AI. It isn't in your P&L.*, sets out the two-minds architecture and the five-level maturity ladder. Paper No. 2, *Don't renovate the work. Re-derive it.*, describes how an operation becomes AI-native. Paper No. 5 applies the method of this paper to a single sector: the accounting and advisory practice.

Xamun Technologies, licensed in the Dubai International Financial Centre, builds purpose-built operational systems for mid-market and government-related enterprises in the Gulf and Southeast Asia, and is paid on outcomes rather than effort. The author has spent twenty-six years building enterprise software, most of it sold by the man-day, and holds the view expressed here as a conclusion rather than a preference.

## Sources

- MIT NANDA, *The GenAI Divide: State of AI in Business 2025* — finding that approximately 95% of enterprise generative-AI pilots deliver no measurable P&L impact.
- McKinsey & Company, *The State of AI (2025)* — approximately 6% of respondents attributing more than 5% of EBIT to AI; workflow redesign identified as the strongest correlate of EBIT impact.
- Gartner (2025) — forecast that over 40% of agentic AI projects will be cancelled by end-2027.
- Government of Dubai, May 2026 private-sector direction on agentic AI adoption; UAE federal target of 50% of government services delivered via AI agents by 2028.
- Xamun Position Papers No. 1 and No. 2 (August 2026).