FREE TOOL

Define the problem before you buy the tool

A workbook for leaders under real pressure to move on AI. It helps you turn that pressure into a clear, defensible decision: where AI genuinely fits, what it's worth, and which product is right for you. Work through it before you sign anything.

> **The pressure is real.** Your customers increasingly expect the businesses they deal with to be using AI. Your competitors are being vocal about having it. Standing still can feel like a disadvantage — so the instinct to act is the right one. This workbook is about pointing that instinct at the right thing, in the right order, so the spend is one you can defend.

01 Separate the pressure from the problem

Both columns are legitimate reasons this is on the table. The left is a problem you can act on; the right is pressure to be seen to act. You need the left before you spend. Tick what's true for you.

A PROBLEM YOU CAN ACT ON

- We can name where we're losing value today
- Someone owns that problem and wants it fixed
- We can describe what good looks like in a year
- We'd still want this fixed if no AI vendor existed

PRESSURE TO BE SEEN TO ACT

- A competitor has been vocal and the board noticed
- Customers are starting to ask what we're doing on AI
- A demo was impressive and we want to use it somewhere
- There's budget to spend before year end

> **What to do with this.** If the right-hand column is doing all the work, you're not stuck — you just don't have the problem defined yet. The rest of this workbook turns the pressure on the right into something you can act on from the left. MIT's 2025 study found roughly 95% of organisations saw no measurable return on GenAI against an estimated $30–40bn of spend; the tools mostly worked, the problem definition didn't.

02 Where is the business losing value?

Before you look at any tool, locate the value. It tends to leak in four places. For each, note where it's happening in your business and what it's costing — in margin, revenue, risk or competitive standing.

WHERE VALUE LEAKS	WHERE THIS SHOWS UP FOR US	WHAT IT'S COSTING
Margin we're giving away Work that's slow, manual or duplicated, quietly eating cost		
Revenue we're leaving Where customers drop out, deals stall or demand goes unmet		
Risk and exposure Where errors, delays or blind spots create liability		
Ground we're losing Where competitors or customer expectations move faster than us		

03 Write the problem down

Pick the one problem from the table that matters most and write it out. Short is good. If you can't fill this in with reasonable confidence, you're ready to go and measure something — not to buy.

WHAT HAPPENS TODAY

WHAT GETS IN THE WAY

WHAT IT COSTS — IN TIME, MONEY, COMPLAINTS OR RISK

THE OUTCOME WE NEED

WHAT GOOD LOOKS LIKE IN TWELVE MONTHS

WHAT WOULD CONVINCE OUR BIGGEST SCEPTIC THAT IT WORKED

If nobody can answer this, go back to the diagnosis. You'll save yourself a pilot that ends in a demo and a shrug.

04 Is AI even the right answer?

Now, and only now, look at solutions — often AI isn't the one. Score each from 1 (no) to 5 (yes) to see what the problem is really asking for before you reach for a model.

Training or people. Would this be fixed faster by teaching people the thing they keep getting wrong? 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐

Process. Is the real problem a broken or missing process that no tool will rescue? 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐

Plain automation. Could rules, a script or existing software do it without AI? 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐

Existing IT. Does a system we already pay for do this, if it were set up or used properly? 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐

Genuinely an AI job. Does this need judgement over messy, unstructured information at a scale people can't manage? 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐

> **Read the scores.** High marks on the first four mean the answer probably isn't AI, and that's a win: those fixes are cheaper, faster and lower risk. Only a clear high score on the last justifies the AI route.

05 What to require before you approve the spend

If AI genuinely fits, these are the answers the board should have before releasing the money — price and platform matter far less. For context: Gartner estimates only around 130 of thousands of "agentic" vendors are the real thing, and that 40%+ of agentic projects will be scrapped by the end of 2027.

- [ ] **What outcome does this deliver, and how will we measure it?** If the answer is "productivity", ask what number moves, by how much, and by when.
- [ ] **Does it actually do the job in our world, and beat the alternative?** Sector, jurisdiction, data, customers, language — a tool built for the general case can be confidently wrong on your specifics. Confirm it performs on yours, and that it genuinely improves on how we work today.
- [ ] **Has it been proven independently, in conditions like ours?** A vendor benchmark on the vendor's own data isn't evidence it will work here.
- [ ] **What does it get wrong, who does it affect, and what's our exposure?** Errors that land on customers or particular groups are a brand and liability risk, not just a defect.
- [ ] **Who is accountable when it's wrong?** Someone owns the output, a decision can be challenged and overridden, and there's a record of what it did.
- [ ] **What has to be true for it to work?** Data, access, integration and adoption. The real cost, and the real risk, usually sits here, not in the licence.
- [ ] **What does it cost us to exit in twelve months?** Lock-in and shelf life. Assume the capability may be obsolete, or far cheaper, well before the contract runs out.

> **The one thing to leave with.** Define the problem first. Decide what you need, who has to change, and what good looks like in a year — then work out whether AI is even the answer. No business case, no purchase. The speed of change makes that discipline matter more, not less.

Weighing up an AI purchase and not sure the problem's clear enough to spend against? That's exactly the conversation to have before you sign. It's the work we do at Deliver Better: define the problem first, then decide whether AI belongs anywhere near it. Start a conversation at deliverbetter.co.uk. First coffee's on David.

A free tool from **David Bain** · Deliver Better

Sources: MIT NANDA, *The GenAI Divide: State of AI in Business 2025*. Gartner, 2025 (agent washing; 40%+ of agentic projects cancelled by end 2027). Problem-diagnosis and sceptic-test framing from the Deliver Better insights series and the LSE AI Leadership Accelerator (N. Soyez).