

GUIDE

AI in 90 Days

A step-by-step guide to starting AI in your organisation the right way



THE OPENING CASE

Only 5% of organisations are achieving AI value at scale. Around 60% report minimal returns despite substantial investment.

Boston Consulting Group, 2025 | Survey of 1,250 senior executives and AI decision-makers

Ninety Days Is Longer Than You Think

Most organisations don't fail at AI because they chose the wrong tool. They fail because they never decided what they were trying to change, never gave anyone ownership of it and never set a date by which something had to be working. A year later there's a subscription, a few enthusiastic users and nothing anyone can point to.

It doesn't have to go that way, and the fix isn't a bigger budget. Organisations getting real value from AI don't do more than everyone else. They do considerably less, aimed far more carefully, with a small number of decisions settled at the start that everyone else leaves open.

This guide sets out those decisions in the order they need making, on a ninety-day timeline you can hold people to. It assumes no technical background, no existing strategy and no appetite for a transformation programme. One piece of work, chosen well, live and measurable inside three months.

16%

of UK businesses have deliberately deployed at least one AI technology (DSIT, 2026)

5%

of organisations worldwide are achieving AI value at scale (BCG, 2025)

21%


have redesigned any workflow around AI, the single biggest predictor of returns (McKinsey, 2025)

What's Inside

- 01 Why Most AI Spending Disappears
- 02 Before You Begin: What Has to Be Settled
- 03 Choosing the First Project
- 04 Days 1 to 90
- 05 Redesign the Work, Don't Just Add a Tool
- 06 Measure It, Then Do the Next One

Why Most AI Spending Disappears

Before deciding what to do, it's worth understanding what usually goes wrong. The pattern is remarkably consistent, and almost none of it is about the technology being immature.



Using Something, Almost Nobody Has Deployed Anything

78% of organisations worldwide were using AI in at least one business function when McKinsey surveyed them in mid-2024, and more than 80% reported no tangible effect on enterprise-level earnings from generative AI (McKinsey, 2025). Adoption has climbed since. The distance between using AI and seeing it in the accounts is where most AI budgets go.

The UK picture sharpens it. When the government asked 3,500 businesses whether they had deliberately deployed a recognised AI technology, 16% said yes, 5% planned to and 80% had neither adopted nor made plans (DSIT, 2026). Informal use is everywhere. Deliberate deployment is rare, which means the bar for being genuinely ahead is far lower than the noise around this subject suggests.

The Gap Between the Top and Everyone Else Is Widening

Only 5% of organisations are achieving value at scale, while around 60% report minimal returns despite substantial spending (BCG, 2025). The top group reaches impact in nine to twelve months against twelve to eighteen for everyone else, and 62% of their initiatives are already deployed against 12% among the laggards.

What separates them is not budget size or technical sophistication. It's a short list of habits, all of which are available to an organisation of any size and most of which cost nothing but attention.

The Single Biggest Predictor Is Whether Anyone Redesigned the Work

McKinsey tested 25 organisational attributes against whether AI showed up in earnings. Redesigning workflows came out on top. Only 21% of organisations using AI have done it (McKinsey, 2025).

That statistic is worth sitting with. The most effective thing an organisation can do is also the thing four out of five don't attempt. Buying licences and hoping people work out the rest remains the most common AI strategy in business, and it's precisely why so little value appears.

Nobody Budgets for What Nobody Owns

Among UK businesses planning to adopt AI within the year, 47% have not set a specific budget and a further 25% don't know what has been set aside (DSIT, 2026). Separately, among businesses already using or planning to use AI that could put a figure on it, median spending on AI across 2024 was £2,000.

Read generously, that shows how cheap starting has become. Read honestly, it shows AI being bought the way stationery is bought, out of whichever budget is nearest, with nobody accountable for what comes back. The first is fine. The second is what the next ninety days are designed to prevent.

Four Questions to Answer Before Anything Else

If you can't answer these today, that's fine. Answering them is what the first thirty days are for.


- What specific outcome would make this worth doing? Name it in a way someone could measure
- Who owns this, by name, and how much of their week is it worth?
- What are we prepared to change about how the work currently happens?
- How will we know in ninety days whether it worked?

MOVING FROM PILOT TO LIVE USE

Only one quarter of organisations have moved 40% or more of their AI experiments into production. Work scoped at three months can stretch to 18 once integration, security and ongoing ownership arrive.

Deloitte AI Institute, 2026 | Survey of 3,235 senior leaders across 24 countries

Before You Begin: What Has to Be Settled



Four decisions determine most of what follows. None require technical knowledge, and all of them are yours to make rather than something to delegate.

Leadership Attention Is the Strongest Signal in the Data

Among organisations achieving value at scale, close to 100% report deeply engaged senior leadership. Among those getting nothing back, the figure is 8% (BCG, 2025). It's the widest single gap in the research.

This doesn't mean chairing weekly meetings. It means the organisation can see that you've decided this matters, that you've named the outcome you want and that you'll ask how it's going. Where that signal is missing, middle managers are left guessing what's expected and reasonably cautious about a technology that touches their own roles.

Don't Hand It to IT and Walk Away

Handing AI entirely to IT is one of the clearest markers of an organisation that gets nothing back. Those achieving value are 1.5 times more likely to share ownership between IT and the part of the business the work belongs to, and more than 40% write that shared ownership into their governance against 19% of the weakest performers (BCG, 2025).

The reasoning is straightforward. IT knows what's possible, what's secure and what will actually integrate. Only the people doing the work know which parts of it are worth changing. Neither has the full picture alone, and a project owned by one of them tends to solve that group's problem rather than the organisation's.

Name One Owner and Give Them Real Time

Someone needs this in their objectives, not on top of them. That person doesn't have to be technical. They need to understand the work being changed, have the standing to ask people to do things differently and have a genuine share of their week protected for it.

Readiness follows ownership more closely than it follows technology. Just over half of UK organisations already using AI feel ready to scale it, against a third of those still planning their first move (DSIT, 2026). What separates those two groups is experience, and experience starts with one person being responsible for finishing one thing.

Settle Governance Now, in Four Sentences

Governance sounds like a policy document and a committee. At this stage it isn't. It's four answers, written down, that everyone can find. Agree them in week one and you remove the ambiguity that stops people using anything.

Most organisations already check their outputs in some form: 84% of UK AI users report at least some human checking and 67% report significant checking, though 2% report none at all (DSIT, 2026). Far fewer have written any of it down. Writing it down is what turns a habit that depends on one conscientious person into something the organisation can rely on.

The Four Governance Questions

Answer each in a sentence or two. One page, circulated, is enough to start.

Question	What a good answer looks like
What data can go in, what can't?	A short list of what's fine, what needs approval and what must never be entered. Name the categories that matter to you, such as personal data, commercial terms or anything covered by a client agreement
Where must output be checked before use?	Name the situations. Anything leaving the organisation, anything informing a decision that carries consequences, anything a regulator might ask about
Where does AI draft, and where does it decide?	At this stage the honest answer is usually that AI drafts and a person decides, everywhere. Say so plainly
Who owns errors and escalation?	One named person, and the route someone takes when something goes wrong. People will forgive a mistake. They won't forgive not knowing who to tell

Choosing the First Project

This is the decision that determines everything else. Organisations achieving value aren't working harder, they're aiming better: they achieve a 76% higher match between where they deploy AI and where it delivers impact (BCG, 2025).

CHOOSING THE FIRST PROJECT

The organisations pulling ahead don't run more AI projects. They place them better: 62% of their initiatives are deployed, against 12% among laggards, and their choices align 76% more closely with where AI delivers impact.

Boston Consulting Group, 2025 | The Widening AI Value Gap

Start With the Work, Not With a Demo

The most common reason UK businesses give for not adopting AI isn't cost, regulation or fear. It's that they haven't identified a use for it, named by 71% (DSIT, 2026). A vendor conversation answers that question for you, which is how organisations end up with capable software solving a problem nobody was losing sleep over.

Turn it around and ask the people doing the work where their week disappears. You'll hear the same handful of answers, and the encouraging part is that the tasks named most often, pulling information out of documents, summarising what's come in, keeping records current, writing up notes and drafting routine correspondence, are exactly the ones this technology handles well today.

What Makes a Good First Candidate

The best first project is boring, frequent and low-stakes. Boring because the task is well understood and nobody's precious about how it's done. Frequent because the saving repeats. Low-stakes because your first attempt will need adjusting and you want that to be a learning moment rather than an incident.

It also needs an owner who wants it. Enthusiasm at the coalface is worth more than executive sponsorship for a first project, because the person doing the work will find the edge cases nobody anticipated.

What to Avoid the First Time

Leave anything customer-facing until you've learned how the technology behaves on internal work. Leave anything requiring judgement across conflicting information, because that's where AI is least reliable and the failure is hardest to spot. Leave anything that needs integration with a system nobody wants to touch.

And resist the instinct to pick the biggest problem you have. Big problems are usually big because they're tangled, and a tangled first project teaches you very little about AI while consuming all of your goodwill.

Where Organisations Put AI First

Business areas named by UK organisations that use or plan to use AI (DSIT, 2026). The right-hand column is what that work typically looks like in practice.

Business area	Named by	What the work usually is
Marketing	72%	Drafting copy, repurposing material, first-pass research
Administration	72%	Document handling, records, correspondence, notes
IT	64%	Code assistance, first-line queries, documentation
Operations	53%	Scheduling, routine reporting, exception handling
Customer service	46%	Drafting responses, summarising cases, routing
Finance	29%	Reconciliation support, variance commentary, reporting

First Project Selection Checklist

A good candidate should tick most of these. If it ticks fewer than five, keep looking.

- The task happens at least weekly, ideally daily
- You could write the instructions for it on a single page
- It's currently done by people whose time is worth more than the task
- Getting it wrong once would be annoying rather than serious
- The information it needs already exists somewhere accessible
- It doesn't require connecting to a system nobody wants to touch
- Someone doing the work is genuinely keen to try
- You could describe the benefit in a number, even roughly
- It's internal rather than customer-facing, for this first round

Days 1 to 90

Three blocks of thirty days. The dates matter more than the detail, because a deadline is what converts an experiment into a decision.

Days 1 to 30: Decide

This month is entirely about choosing and agreeing. Map where the time goes, shortlist three candidate tasks, pick one, name the owner, answer the four governance questions and write down what success looks like as a number.

It's tempting to compress this and start building in week two. Don't. Almost every stalled AI project can be traced back to something that was never actually decided in month one.

Days 31 to 60: Build and Test

Now you configure, connect and try it on real work with a small group of real users. Expect the first version to need work. That's normal, and it's why this block is thirty days rather than five.

Test with the awkward cases deliberately, not just the clean ones. The gap between a pilot that works and a live system that works is almost always made of edge cases nobody thought to try.

Buy What's Standard, Build What Isn't

Most organisations do both. Of UK organisations investing in AI this year, 65% are implementing off-the-shelf applications and 59% are embedding AI into systems they already run, while 22% plan something built specifically for them (DSIT, 2026).

Each is right in the right place. Buy where the workflow is standard, the data isn't sensitive and speed matters most. Build where the work is genuinely particular to you, where the data can't leave, or where configuring a general tool would mean paying for a great deal you'd never use. Custom sounds like the expensive option, and set against the licences, workarounds and quiet hours a poor fit consumes, it frequently isn't.

Days 61 to 90: Go Live and Measure

Move it into daily use for the group who need it, with the check step in place and the owner watching. Measure against the number you wrote down on day thirty. Ask users what's still awkward, then fix the two most common answers.

At day ninety you make a decision: scale it, adjust it or stop it. Deciding to stop is a perfectly good outcome and considerably better than the alternative, which is letting it drift on unmeasured for another year.

Plan for Live From Day One

Only a quarter of organisations have moved even 40% of their AI experiments into live use, though 54% expect to reach that level within three to six months (Deloitte, 2026). The reason isn't that pilots fail. It's that going live means integration, security sign-off, monitoring and someone owning it on a Monday morning, and work scoped at three months can stretch to eighteen once that arrives.

None of it is insurmountable, but it does have to be anticipated. Ask the going-live questions in month one when they're cheap, rather than in month three when they're a surprise.

The Ninety Day Plan

Block	What happens	What exists at the end
Days 1 to 30 Decide	Map where time goes. Shortlist three tasks. Choose one. Name the owner. Answer the four governance questions. Define the measure of success	A one-page brief naming the task, the owner, the change and the number you're aiming at
Days 31 to 60 Build and test	Configure and connect. Run it on real work with a small group. Test awkward cases deliberately. Redesign the surrounding steps rather than inserting a tool into them	A working version used by real people, with the review step defined and the manual steps it replaces removed
Days 61 to 90 Go live and measure	Move into daily use. Measure against the day-thirty number. Gather what's still awkward. Fix the top two. Decide	A live workflow, a measured result and a decision to scale, adjust or stop


Questions to Ask Any Vendor

Eight questions that separate a tool you'll still be using next year from one that quietly lapses.

- What does this do that our current process doesn't, in one sentence?
- What does it cost at ten users, and at a hundred, once usage grows?
- Where does our data go, who can see it and is any of it used to train models?
- What happens to our data and our outputs if we leave?
- Which parts of this are configuration and which parts need development?
- What does implementation involve on our side, counted in days of our people's time?
- Who else of our size and in our sector runs this, and may we speak to them?
- What does it do badly? A vendor with no answer hasn't deployed it very often

Redesign the Work, Don't Just Add a Tool

This is the step almost everyone skips and the one the evidence says matters most. It's also far less daunting than it sounds.



What Redesign Actually Means

It doesn't mean restructuring a department. It means deciding which step AI drafts and which step a person owns, agreeing what gets checked before anything goes anywhere and removing the manual stage that only existed because the old way needed it.

That last part is where the value hides. Organisations routinely add AI to a process and keep every step AI made unnecessary, then wonder why nothing got faster. If a report is now generated automatically, the meeting that existed to assemble it needs to change too.

Where the Effort Should Go

BCG's rule of thumb for technology change puts 70% of the effort into people and process, 20% into technology and 10% into algorithms (BCG, 2025). Nearly every stalled AI project has that ratio upside down.

It makes a useful test of your own plan. If the ninety days is mostly about choosing a tool, you're spending your attention on the 20% and hoping the 70% happens by itself.

Keep People Where the Judgement Is

In a controlled experiment with 758 consultants, AI improved output on suitable tasks by more than 30%, but on a task requiring judgement across conflicting information it caused a 19 percentage point drop in correct answers (Dell'Acqua et al., 2026).

The same study found something more important still. AI-assisted work scored higher on persuasiveness and coherence whether the underlying answer was right or wrong. Reading the output and judging whether it reads well is therefore not a check. Design your review around whether the reasoning holds, not whether the writing flows.

Involve the People Doing the Work

Organisations achieving value involve their workforce in reshaping workflows roughly twice as often as everyone else (BCG, 2025). This isn't consultation for its own sake. The person who does the task daily knows which exception breaks everything, and they'll tell you in week two rather than month six if you ask.

There's an adoption benefit too. People who helped design something tend to use it. People handed a finished system tend to work around it.

Redesign Prompts

Walk through the workflow with these five questions and note every answer.

- Which step does AI draft, and which step does a person own?
- What must be checked before this leaves the desk, and who checks it?
- Which existing steps are now unnecessary, and who has authority to remove them?
- What happens when the output is wrong, and how would we know?
- What does the person freed from this task do instead?

CHOOSING WHERE PEOPLE STAY IN CONTROL

On suitable professional tasks, AI lifted output quality by more than 30%. On a task requiring judgement across conflicting evidence, it reduced correct answers by 19 percentage points.

Dell'Acqua et al., 2026 | Controlled experiment with 758 consultants

Measure It, Then Do the Next One

Measurement is the difference between a project you can build on and an anecdote. It's also the most commonly skipped step in the whole process.

Almost Nobody Tracks This, and It Costs Them

Fewer than one in five organisations track defined measures for their AI work, despite tracking having the largest measured effect on returns of any practice examined (McKinsey, 2025).

Among organisations achieving value at scale, more than 60% track rigorously, against 17% of those getting nothing back (BCG, 2025).

The measure doesn't need to be sophisticated. Hours per week on the task before and after. Number of items processed. Errors caught. Days from request to delivery. Pick one or two, capture the starting position before you change anything, and you'll have something to show at day ninety.

Measure Use, Not Just Licences

Access to AI tools grew by 50% in a single year, yet among people who have access, fewer than 60% use it in their daily work, a figure that hasn't shifted (Deloitte, 2026). Counting licences tells you what you bought. Counting daily use tells you what happened.

If use is low, that's information rather than failure. It usually means the tool sits beside the work rather than inside it, which is a design problem you can solve.

Expect Efficiency First and Revenue Later

Among UK organisations already using AI, 75% report improved workforce productivity and 57% report new or improved processes, while 77% have seen no change in revenue yet (DSIT, 2026). The international picture agrees: 66% of organisations say they're achieving efficiency gains today, against 20% seeing revenue growth from a group where 74% hoped for it (Deloitte, 2026).

That order is worth setting out to your board before you start rather than afterwards. Efficiency arrives reliably and early. Revenue arrives later, indirectly and only where the capacity you've freed is pointed at something that earns.

The Second Project Is Cheaper Than the First

Everything the first project taught you carries forward. The data questions are answered, the review habits exist, the governance page is written and the people who were sceptical at the start have watched something work.

This is why starting small isn't a compromise. It's the fastest route to a second project that costs less and lands quicker, and then a third. The organisations pulling ahead didn't leap. They compounded.

Worth Tracking From Day One

Measure	How to capture it	Why it matters
Time on the task	Ask the people doing it before you change anything, then again at day 90	The clearest evidence of benefit and the easiest to explain to anyone
Daily use	Count people using it on real work each week, not licences issued	Separates adoption from availability, which is where most value is lost
Output quality	Sample a set of outputs and review the reasoning rather than the writing	Guards against work that reads well while being wrong
Cycle time	Days from request to delivery	Often improves more than headline time saved and is felt by everyone downstream
Exceptions	Log every case the workflow couldn't handle	Tells you what to fix next and what the second project should be

One Thing, Done Properly

Nothing in this guide requires a transformation programme, a data science team or a leap of faith. It requires choosing one piece of work, deciding who owns it, agreeing four sentences of governance and holding yourself to a date.

That's a modest amount of discipline for something that puts you in a small group. The organisations seeing real returns aren't further along a technical curve than you. They settled the questions everyone else left open, then finished something.

And when it lands, the thing you'll notice first won't be a number on a report. It'll be someone in your organisation getting an afternoon back and spending it on the work they came here to do.

Want help choosing that first project? A short conversation is usually all it takes to find it.

Jacobs & Xu

Applied AI. Results You Can Measure.

contact@jacobsandxu.com | +44 (0)203 105 7544

Sources

References are given in the text by author and year. Full details are below, in Harvard format.

Boston Consulting Group (BCG) (2025) *The Widening AI Value Gap: Build for the Future 2025*. Boston: Boston Consulting Group. Survey of 1,250 senior executives and AI decision makers across 68 countries and 25 or more sectors. Published September 2025.

Deloitte AI Institute (2026) *State of AI in the Enterprise: The Untapped Edge*. Deloitte Insights. Survey of 3,235 board, C-suite and director-level respondents across 24 countries, fielded August to September 2025. Published January 2026.

Dell'Acqua, F., McFowland, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayner, L., Candelon, F. and Lakhani, K. (2026) 'Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality', *Organization Science*, 37(2), pp. 403 to 423. Preregistered randomised controlled experiment with 758 consultants.

Department for Science, Innovation and Technology (DSIT) (2026) *AI Adoption Research*. London: DSIT. Research conducted by IFF Research and Technopolis Group. Telephone survey of 3,500 UK businesses with five or more employees, fielded 12 February to 2 May 2025, with 100 follow-up interviews. Published January 2026, updated February 2026. Available at: <https://www.gov.uk/government/publications/ai-adoption-research> (Accessed: 12 August 2026).

McKinsey & Company (2025) *The State of AI: How Organizations Are Rewiring to Capture Value*. Survey of 1,491 participants across 101 nations, fielded July 2024. Published March 2025.

Figures drawn from executive surveys are self-reported by the organisations taking part. Where this guide describes what high-performing organisations do differently, it describes a pattern found among them rather than a guaranteed result. Four points of context are worth carrying with the numbers. Every Deloitte respondent already runs at least one AI implementation, so its access and usage figures are an upper bound rather than an average. BCG's sample is weighted towards large enterprises, so its findings read better as principles than as benchmarks. McKinsey's fieldwork was carried out in July 2024 and adoption has risen since. DSIT figures cover UK businesses with five or more employees and measure deliberate deployment of a recognised AI technology, which is a narrower definition than most commercial surveys use.