

THE IGNITE ALMANACK OF **THE 90-DAY PROOF**

How to know your AI programme is working before you have paid for all of it.

A FIELD GUIDE FOR UK SMES

IGNITEAI SOLUTIONS

THE PROBLEM

You will be asked to commit a year's budget on a promise.

Every AI proposal you read has a number at the bottom and a story about what happens next. The story is usually good. The trouble starts at month 3, when someone asks whether it is working and the room goes quiet.

Nobody can answer, because nobody wrote down what "working" looked like on day one. There is no baseline. So the conversation turns into opinion. The keen say it feels faster. The sceptics say they have not seen anything. Both are guessing, and the loudest voice wins.

Then the budget conversation arrives. A programme with no evidence behind it loses to anything that turns up with a spreadsheet.

WHY THIS MATTERS

A programme that cannot show its own value gets killed for being unproven, which is a common failure and an avoidable one. The evidence has to be designed in on day zero. You cannot go back and take a baseline in month 4.


THE IDEA

90 days is a gate.

At day 90 the programme stops and gets measured. If the numbers are there, it scales. If they are missing, something upstream is wrong and it gets corrected before another pound goes in.

The gate has one job: to make failure cheap. A programme that fails at day 90 has cost you a quarter. The same programme discovered at month 14 has cost you the year, and it has taught your whole company that AI does not work here. That belief is the expensive part.

For the gate to work, three things have to be true before you start. You need a baseline taken by measurement. You need success written as a number, agreed by the people who will be judged on it. You need a named owner with a named deputy.



The gate at day 90 compares against the baseline taken at day 0, using the same measurement both times.

WHAT SUCCESS LOOKS LIKE

Five things you should be holding at day 90.

Success at day 90 is a short list of things you either have or do not have.

- 1 A number that moved, taken the same way twice.** The day 90 measurement uses the day 0 instrument. Same log, same people, same length of run.
- 2 One workflow running daily that nobody babysits.** The real test is whether it survives the person who built it being on holiday.
- 3 A signed policy with names against it, and more than one administrator.** Signed, circulated, and acknowledged by the people it governs.
- 4 Two or three people who can build the next one.** If every improvement still routes through one person, you have bought a dependency.
- 5 A decision you can defend to whoever holds the budget.** With the working shown, in their language, on one page.

WHAT A FAILED GATE IS WORTH

If day 90 says the numbers are not there, the gate has done its job. You have spent a quarter to learn something that costs most businesses a year. Correct the workflow choice, the training, or the design, and run the next 90 days from a better start.

What the 90 days actually contain.

When	What happens	What you hold at the end of it
Day 0	14 days of real task logging across the roles in scope.	A measured picture of where the hours go, and the agreed number that defines success.
Days 1 to 30	Governance signed. Policy, data classification, red lines, a named owner and a named deputy.	A signed policy people can follow. Training starts after this.
Days 30 to 60	The highest-friction workflow goes live. The first group is trained on it.	One workflow running daily, and people who can run it without you.
Days 60 to 90	The workflow runs in anger. Time and quality are tracked against the baseline.	Evidence, collected the same way it was collected on day zero.
Day 90	The proof report. Time saved, quality changed, risk reduced, in real numbers.	A decision you can defend in front of the people paying for it.

THE ROW PEOPLE GET WRONG


Days 1 to 30 ends at a signature. A policy sitting at 70% governs nothing. Every programme in this guide that ran into trouble had a document in progress and a team already using the tools.

WHAT STOPS IT

Four blockers, and one of them causes the other three.

Across the engagements behind this guide, the same four things stop a programme reaching a day 90 it can defend. They do not arrive separately.

The root is almost always the same. One person is carrying the programme. When that person is on leave, overloaded, or gone, the problem definition never gets agreed, the baseline never gets taken, and the policy sits unfinished while the team gets on with it anyway.



Fix the root and the other three become tractable. Fix the three without the root and they come back.

BLOCKER ONE

The problem was never agreed.

"We want to use AI" is a budget line looking for a justification. The businesses that get a defensible day 90 start from a business problem specific enough that you could measure it before anyone touched a tool.

The tell is easy to spot. Ask four people in the same company what the programme is for and count the answers.

SEEN IN THE FIELD

An engineering consultancy of about 60 people. Three engineers in the same discovery room. The principal engineer, 40 years in, worked on pen and paper and could not see a problem worth solving. A junior wanted help writing report paragraphs and could not see past the fact that no two jobs were alike. The third would not let juniors near it at all, because his worry was professional indemnity and skill atrophy. Elsewhere in the same business, finance, operations and sales were each investigating automation separately, with no single view of what mattered most. Four definitions of the programme, one company.

What this costs you is a pilot built to solve a problem nobody quite agreed on, which lands well with nobody and hardens into "we tried AI and it did not really work for us."


BLOCKER TWO

Nobody measured the before.

Ask an owner how long the team spends quoting each week and you get a number in about 4 seconds. Ask the team to log it for a fortnight and the number moves. Sometimes it halves. More often it doubles, because the estimate covers the work people notice and misses the work they have stopped noticing: the chasing, the re-checking, the same figure typed into a fourth system.

SEEN IN THE FIELD

A manufacturer whose ERP costed labour at £22 an hour. The real figure ran between £30 and £33. They were pricing to a loss in places and could not say where or why, and a strategic decision about which product lines to specialise in had been waiting on data the business already owned. In a service business the same season, a small team believed it handled around 10 upgrade requests a week. Fourteen days of logging measured a peak of 43. More than 4 times the assumed load, invisible for years.



Illustrative shape, drawn from the pattern that repeats across engagements. Your own numbers will differ. The gap between the two bars is the point, and the orange portion is usually the cheapest thing to fix.

BLOCKER THREE

One person was carrying it.

This is the root, and it is rarely a lack of enthusiasm. The programmes that stall usually have a committed, capable champion. They have exactly one.

A single owner looks efficient in month one. By month three it is the reason nothing moves, because everything routes through one diary, one inbox, one signature.

SEEN IN THE FIELD

A consultancy across 3 offices had one champion. She went on extended leave. The plan slipped 2 months while two colleagues tried to carry it around their day jobs. At a larger client the same year, 3 separate workstreams stalled for 3 different single-person reasons: a month to secure one meeting and then 3 weeks of silence on the follow-up questions, a solution approved and waiting on one person's budget signature, and a workstream whose only named lead left with no handover. A travel business froze company-wide because one director was the sole system administrator and went away for a fortnight, which meant nobody could approve a permission, adopt the policy, or set a spend cap until he returned.

The fix is unglamorous and takes an afternoon. Name a deputy for every owner, and make sure at least two people hold every administrative permission the programme depends on.

Governance lost the race.

Almost nobody skips governance on purpose. What happens is that the policy is in flight while the team is already using the tools, and the tools move faster than the document. By the time the policy is ready, the behaviour it was meant to shape has already set.

SEEN IN THE FIELD

At one firm the chief executive told the company to learn AI or get out, then engaged with none of the governance that followed. A field marketing manager spent roughly \$4,000 of model tokens in a month building something that turned out to be useless. At a UK engineering practice the same quarter, a fire alarm specification was produced from a public chatbot before any policy existed, and one member of staff listed 10 different AI apps in personal use. At a third, a developer connected an AI agent to a 25-year-old system with no API and let it write dates back into 150 live projects, with no change control, no rollback and no human in the loop.

The spread is the part people miss. Training opens a gap. Some race ahead, some never start. In one 7-person firm the positions were named person by person before the training day: two excited, one neutral, one resistant. The fastest adopter is where the risk concentrates, because that is the person putting real client data through a tool nobody has approved yet.

TRY THIS ON MONDAY

Run what one client called an AI amnesty, explicitly not a witch hunt. Ask everyone, in writing, which AI tools they already use for work. **You are buying an accurate picture, so promise in advance that nobody gets in trouble for answering honestly.**

One problem. One pilot. One measurable outcome.

- 1 Pick the roles, not the company.** 3 to 6 people whose work you are actually trying to change. A baseline across everyone is a survey. A baseline across the roles in scope is evidence.
- 2 Log for 14 days.** Task, minutes, and whether it involved rekeying something that already existed somewhere else. 14 days catches the weekly and fortnightly work that a single week misses.
- 3 Write the success number down and get it agreed.** "Quoting drops from 72 hours a fortnight to 50" is a target. "The team uses AI more" is a mood.
- 4 Name the owner and the deputy.** Two people, both senior enough to change how a team works. Check that no single permission, licence or approval sits with only one of them.
- 5 Sign the policy before anybody is trained.** Signed and acknowledged, with the red lines written in plain words people can actually follow.
- 6 Build one thing.** Take the workflow with the highest hours and the highest irritation. Irritation matters, because the people doing that job will forgive an imperfect first version.
- 7 Re-measure the same way.** Same log, same fortnight, same people. A different method at day 90 gives you a number you cannot defend.

Three prompts to take your own baseline.

These run in Claude, ChatGPT or Gemini. Start with the first one.

PROMPT 1: DESIGN THE LOG

I run a [size]-person [sector] business in the UK. Before I invest in any AI tooling I want a 2-week baseline of where my team's time actually goes. Roles in scope: [list the roles].

Design a daily task log they can fill in under 3 minutes. Ask what our main systems are first. Then give me the columns, one worked example row, and the 3 questions I should be able to answer at the end of the fortnight.

Builds the instrument before you collect anything, so the data answers a question you chose.

PROMPT 2: FIND THE NUMBER THAT MATTERS

Here is our 2-week task log: [paste the data].

Group the tasks by category and total the hours for each. Then find every task where the same information goes into more than one system, and total those separately. Show me the 3 categories with the highest hours, and for each one a realistic 90-day reduction and what you assumed to get there.

Separates volume from duplication. Duplication is usually the cheapest win.

PROMPT 3: WRITE THE GATE

Our baseline says [category] takes [X] hours a fortnight across [N] people. We want that down by [Y]% in 90 days.

Write me a one-page proof plan: the exact measurement we repeat at day 90, who owns it, who deputises, what evidence we will hold, and the 3 signs at day 45 that would tell us this is not going to land.

The day-45 warning signs are the point. A gate you only look at once is a post-mortem.

THE NEXT STEP

TAKE THE BASELINE

The **AI Readiness Scorecard** takes a few minutes and gives you a structured read on where your business actually is, across culture, governance and process, before you spend anything.

Skills that stay when we go.

If it says you are ready to move, a **SPARK discovery** takes the baseline properly and sets the 90-day number with you.

IGNITEAI SOLUTIONS

igniteaisolutions.co.uk · chris.duffy@igniteaisolutions.co.uk

THE IGNITE ALMANACK
FIELD GUIDE 09