Anyway Ventures

UK BUSINESSES · 50 TO 1,000 STAFF · 2026 EDITION

AI readiness *scorecard*.

A 25-question AI readiness assessment with the evidence behind every score, a use-case scoring method, a data readiness test, a staff survey on shadow AI, a 90-day plan for each score band and a board summary. Built for UK businesses with 50 to 1,000 staff.

PDF GUIDE | EXCEL SCORECARD

AUTHOR

Brannan Coady, Anyway Ventures

EDITION

Version 1.0 · October 2026

ONLINE, WITH DOWNLOADS

anywayventures.com/resources/ai-readiness-scorecard

CONTACT

brannan@anywayventures.com

Free to use, adapt and share with attribution · Not legal, tax or investment advice

CONTENTS

In this guide

1. Executive summary
2. How to use this scorecard
3. Where UK businesses actually are
4. The scoring rules
5. Strategy and ownership (questions 1 to 4)
6. Processes (questions 5 to 8)
7. Data (questions 9 to 13)
8. Systems and integration (questions 14 to 16)
9. People and skills (questions 17 to 19)
10. Governance and risk (questions 20 to 23)
11. Money (questions 24 to 25)
12. Scoring method and bands
13. Data readiness test
14. Use-case opportunity scoring
15. Shadow AI survey
16. 90-day plan by score band
17. Presenting results to the board
18. How this compares with published maturity models
19. Worked example 1: Holbeck Fabrication Ltd
20. Worked example 2: Aldermoor Insurance Services Ltd
21. Common mistakes
22. Glossary
23. Sources and further reading
24. About this resource

About Anyway Ventures

WHO IT IS FOR

CEOs, COOs, CFOs, operating partners, heads of IT and transformation in UK businesses with 50 to 1,000 staff

ALSO AVAILABLE

PDF guide, Excel scorecard at anywayventures.com/resources/ai-readiness-scorecard

KEY TAKEAWAYS

The short version

1. The Anyway Ventures AI readiness scorecard for UK businesses with 50 to 1,000 staff scores 25 questions from 0 to 2 across seven sections, for a total out of 50.

2. In an AI readiness assessment, four gate questions override the total: a measured problem, a named owner with time, data reachable without manual exports and an AI policy people have been trained on.

3. An AI readiness score under 20 out of 50 means fix foundations and buy nothing, 20 to 34 means one contained first build, and 35 or more means a ranked portfolio of builds.

4. Around 35% of UK businesses with ten or more employees use AI, but only 11% say more than half their workforce has had AI-related training (ONS, July 2026).

5. Microsoft's UK research in October 2025 found 71% of employees had used unapproved consumer AI tools at work, so run an anonymous shadow AI survey before writing policy or buying licences.

This is a free AI readiness scorecard for UK businesses with 50 to 1,000 staff. It takes the same 25 questions, 0 to 2 scoring, four gate questions and score bands as our AI readiness assessment checklist, and adds what a checklist leaves out: the evidence that justifies each score, how to collect it and who to ask, what to do about a 0 or a 1, a method for scoring and ranking your first ten use cases, a data readiness test, a shadow AI survey to send to staff, a 90-day plan for each band and a one-page board summary. It is written for the CEO, COO or CFO who has to decide what the business does about AI next, and for the operating partner or adviser helping them.

1. Executive summary

1. **Readiness belongs to a process, not to a company.** Nobody is "ready for AI" in general. Answer the process, data and systems questions (5 to 16) for the one process you most want to fix, and the rest for the business as a whole.
2. **Score evidence, not opinions.** A 2 means you can show a document, a number or a named person. "Don't know" scores 0. Every question below lists the evidence that earns each score, so two people scoring the same business should land within a few points of each other.
3. **Four gates override the total.** A measured problem (question 1), a named owner with time (2), data reachable without manual exports (10) and a policy people have been trained on (20). A zero on any of them is the first thing to fix, whatever the total says.
4. **Then the weakest section, then the total.** Under 20 out of 50: fix foundations, buy nothing. 20 to 34: one contained first build on your best-prepared process. 35 or more: a ranked portfolio of builds. If you land within two points of a boundary, take the lower band.
5. **Most mid-sized UK firms are earlier than they think.** The ONS reported in July 2026 that around 35% of businesses with ten or more employees use AI, but only 11% say more than half their workforce has had AI-related training. People and data are where self-assessments are most often overscored.
6. **Your staff are already using AI, with or without you.** Microsoft's UK research in October 2025 found 71% of employees had used unapproved consumer AI tools at work. Run the anonymous survey in this guide before you write policy or buy licences; it also produces your best list of use cases.
7. **Rank use cases on five things, in pounds where you can.** Value, feasibility, data availability, risk and time to value, weighted and scored out of 100, with hard flags that park anything with unready data or unmanaged risk. The top of a good list is usually dull back-office work.
8. **Give the board one page, not a maturity level.** Score, gates, weakest section, the top three opportunities with their value and cost, the 90-day plan, the decision you need, and what you will report next quarter.

2. How to use this scorecard

Who should run it. One person owns the exercise, ideally the executive who will own AI outcomes afterwards (usually the COO or CFO, not IT). They need two groups: the leadership team for questions 1 to 4 and 17 to 25, and the owner of the priority process, with one or two people who do the work, for questions 5 to 16.

When. At the start of a planning cycle, before buying AI licences or commissioning a build, after a stalled pilot, or in the first 100 days after an investment. Repeat it every six months; the movement in the score matters more than the first number.

How long. Two hours of meetings to score it, plus two to three weeks of part-time work to collect evidence. Do the evidence work. The fastest way to get a useless score is to fill it in from memory in a meeting room.

STEP	WHAT HAPPENS	TIME	OUTPUT
1. Pick the priority process	Leadership names the one process they would most like to fix	30 minutes	One process, one owner
2. Launch the shadow AI survey	Send the survey in this guide to all staff, anonymously	1 hour to set up, 10 days open	Usage, data exposure, task wish list
3. Collect evidence	Time sample, data readiness test, licence and system checks	2 to 3 weeks part-time	The evidence file for each question
4. Score	Two sessions, scoring each question against the evidence tables	2 hours	Scores, gates, band, weakest section
5. Rank opportunities	Score up to 20 candidate use cases, shortlist ten	2 hours	Ranked opportunity list
6. Plan and present	Choose the 90-day plan for your band, write the board page	Half a day	Board summary and plan

Inputs you will need. An organisation chart, a list of the main systems and who administers them, a copy of any AI or data policy, Microsoft 365 (or Google Workspace) admin usage reports, the software subscription list from finance, a fully loaded hourly cost for the teams in scope, and a week of timing data for the priority process.

The Excel companion. The workbook has a blank scorecard with dropdowns for every question, section subtotals, gate checks and the band calculated by formula, and a chart of section scores. It has a use-case matrix that applies the weights in this guide and ranks the results, a shadow AI survey tab you can copy into Microsoft Forms or Google Forms, and a filled-in example using the second worked example below. Score in the workbook; use this guide to decide what each score should be.

35%

of UK businesses with ten or more staff use AI; 11% say more than half their workforce has had AI training. ONS, 2026

23%

of mid-sized UK businesses (50 to 249 staff) currently use AI, against 36% of large ones. DSIT, 2026

64%

of mid-sized AI users say limited skills have held back wider adoption. DSIT, 2026

71%

of UK employees have used unapproved consumer AI tools at work. Microsoft UK, 2025

3. Where UK businesses actually are

Three recent UK sources, read together, describe the starting point for most businesses in this size range.

The ONS article AI in UK businesses: 2023 to 2026, published in July 2026 from its Business Insights and Conditions Survey, found use of at least one AI technology among businesses with ten or more employees rose from around 12% in September 2023 to around 35% in June 2026, and to 49% among businesses with 250 or more employees. Depth has barely moved: adopters use about 1.6 AI technologies on average, up from about 1.4; only 15% of businesses say more than half their employees use AI in daily work; and only 11% say more than half their workforce has had AI-related training. The ONS also found 55% of employees report using AI for work or education, well above the 35% of businesses reporting use. Staff are ahead of their employers.

DSIT's AI Adoption Research, published in January 2026 from 3,500 telephone interviews with UK private sector businesses of five or more employees (fieldwork February to May 2025), found 16% currently using at least one AI technology: 23% of mid-sized businesses and 36% of large ones. The most common reasons for not adopting were not having identified a use for AI (71%) and limited AI skills (60%); 21% cited data complexity. Among current users, 54% felt ready to scale and only 13% described themselves as completely ready, and 64% of mid-sized users said limited skills had hindered wider adoption. Of businesses planning to adopt in the next year, 47% had no specific budget and 25% didn't know what was budgeted. The report notes that because it interviewed the person responsible for technology, it cannot see shadow AI.

That gap is what Microsoft's UK research, published in October 2025 from a Censuswide survey of 2,003 UK employees, fills. It found 71% had used unapproved consumer AI tools at work and 51% did so every week; 28% said their employer offered no approved option; and only 32% were concerned about the privacy of company or customer data they entered. The most common uses were drafting and responding to workplace communications (49%), drafting reports and presentations (40%) and finance-related tasks (22%).

The practical reading for a business of 50 to 1,000 people: staff use is wide and unmanaged, business use is shallow, and the two most cited blockers, not knowing what to use AI for and lacking skills, are both things a structured assessment fixes. Infrastructure is rarely the constraint at this size.

4. The scoring rules

Score every question 0, 1 or 2, using the evidence tables below.

- **2: yes, and you can show it.** A document, a report, a number or a named person, available today.
- **1: partly, or only in someone's head.** It exists but isn't written down, isn't agreed, or covers only some of what is needed.
- **0: no, or you don't know.** Not knowing is a finding, and scores accordingly.

Three rules keep the score honest.

Score the priority process for questions 5 to 16. If you score the best bits of five processes, you will have a high score and no process you can build on.

Use the lower score when the room disagrees. If the finance director says 2 and the person who does the work says 1, it's a 1. People closest to the work are usually right about the process and data questions.

Write the evidence down. The Excel scorecard has an evidence column for each question. A score with no evidence in that column should be treated as a 1 at most. When you repeat the assessment in six months, the evidence column is what lets you see whether anything has really changed.

The seven sections carry different weights because they have different numbers of questions.

SECTION	QUESTIONS	MAXIMUM	SHARE OF TOTAL
Strategy and ownership	1 to 4	8	16%
Processes	5 to 8	8	16%
Data	9 to 13	10	20%
Systems and integration	14 to 16	6	12%
People and skills	17 to 19	6	12%
Governance and risk	20 to 23	8	16%
Money	24 to 25	4	8%
Total	**25**	**50**	**100%**

Gate questions are marked in the headings below.

5. Strategy and ownership (questions 1 to 4)

These questions test whether there is a measured problem and someone who owns it. Most failed AI projects fail here, long before the technology is chosen.

1. Can you name the one or two business problems you want AI to help with, in numbers? (Gate)

SCORE	EVIDENCE THAT JUSTIFIES IT
2	One or two problems written down with a number attached: hours a month, turnaround time, error rate, backlog or margin. The leadership team can quote them.
1	Problems are named ("quoting is slow") but not measured, or the measures are guesses nobody has checked.
0	The goal is general ("be more efficient", "use AI") or there are ten priorities and no ranking.

Collect it: ask each member of the leadership team, separately and in writing, for the one operational problem they would most like fixed and how big it is. Compare the answers.

Ask: the CEO, COO and CFO; then the manager of the team concerned, who will usually correct the size.

If you scored 0 or 1: run a one-week time sample on the top two candidates (people log time on the task in 15-minute blocks, or you sit with them and time 20 instances). Convert the result to hours a month and pounds at a loaded hourly rate. That turns a 0 into a 2 in a week.

2. Is one named person accountable for AI outcomes, with time set aside? (Gate)

SCORE	EVIDENCE THAT JUSTIFIES IT
2	A named executive with AI outcomes in their objectives, at least half a day a week in their diary, and authority to change processes and budgets in the area.
1	Someone is named but has no protected time, no authority over the affected teams, or ownership sits with a working group chaired by them.
0	"IT is looking at it", a steering group with no single owner, or an enthusiastic junior with no authority.

Collect it: look at the named person's objectives and their diary for the last month. If AI work isn't in either, the honest score is 1 at best.

Ask: the CEO, and the named owner directly: "What have you decided about AI in the last month?"

If you scored 0 or 1: the CEO names one owner this week, usually the COO or CFO, writes the role into their objectives and clears the time. The board or investors should know who it is. Ownership by IT alone tends to produce tools, not changed processes.

3. Does the leadership team agree what AI is for here: cost, capacity, quality or product?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	A one-sentence statement of purpose that the CEO, CFO and operations lead all give unprompted, for example "absorb 20% growth without adding back-office headcount".
1	Broad agreement in conversation, but different emphases when asked separately, or nothing written down.
0	Unspoken disagreement: the CFO expects headcount savings, sales wants a chatbot, operations wants fewer errors.

Collect it: in the same written exercise as question 1, ask each leader "What should AI do for this business in the next 12 months?" Read the answers aloud at the scoring session.

Ask: every member of the executive team, separately.

If you scored 0 or 1: hold a 45-minute session with one output: a single sentence that ranks cost, capacity, quality and new products for the next 12 months. Most mid-sized businesses should start with capacity (doing more with the same people), because it is easier to measure and less threatening to staff than cost.

4. Have you agreed how you will measure success before you start?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	A written baseline for the priority process (volume, time per item, error or rework rate, cycle time) taken over at least a week, and an agreed target.
1	The measures are agreed but no baseline has been taken, or a baseline exists but nobody agreed what good looks like.
0	"We'll know it's working when we see it."

Collect it: check for a baseline document. If one exists, check whether it was measured or estimated.

Ask: the process owner and finance.

If you scored 0 or 1: take a baseline before anything is built. Four weeks of normal trading is better than one; one is far better than none. The method is set out in measuring AI ROI. Without a baseline, a project that works cannot prove it.

6. Processes (questions 5 to 8)

AI pays back on work that is frequent, repetitive and done by several people, where a person can check the output quickly. Answer these for the priority process.

5. Have you found repetitive tasks done at volume, weekly or more often, by several people?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	A list of at least five candidate tasks with volumes ("2,000 supplier invoices a month", "60 quotes a week") and the roles that do them.
1	Candidates are known but volumes are not, or the list is one or two items from one team.
0	No list, or a list of one-off projects and quarterly tasks.

Collect it: a 30-minute walk-through with each team lead asking "What does your team do every day or every week that involves reading something and typing it somewhere else, or writing something similar to last time?" Add the free-text answers from question 11 of the shadow AI survey.

Ask: team leads in finance, operations, customer service, sales support and HR.

If you scored 0 or 1: spend a day collecting candidates and their volumes. Volume often comes straight from systems: invoices from the finance system, tickets from the helpdesk, emails from the shared mailbox report.

6. Is the priority process done the same way whoever does it?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	A current written procedure, or two people who do the work describe the same steps and the same exceptions independently.
1	Broadly consistent, but exceptions and judgement calls vary by person, or the procedure is out of date.
0	Each person has their own method, and the exceptions live in one person's head.

Collect it: ask two people who do the process to each list the steps on paper, separately, then compare. Differences are the finding.

Ask: two or three people who do the work, not their manager.

If you scored 0 or 1: agree one way of doing it before automating it. Write the steps and the top ten exceptions on one page. Automating three different methods gives you three different failure modes.

7. Do you know how long the process takes and how often it goes wrong?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	Measured time per item and an error or rework rate, from a time sample or system timestamps, over at least a week.
1	Estimates from the team that nobody has checked, or time is known but errors are not.
0	No idea, and no way to find out without asking around.

Collect it: the one-week time sample from question 1, plus a count of items that came back, were corrected or escalated.

Ask: the people doing the work, and whoever handles complaints or corrections downstream.

If you scored 0 or 1: time 20 to 30 real instances, including the awkward ones. Expect the measured time to differ from the estimate; in my experience it usually comes in lower than the estimate for routine items and much higher for exceptions.

8. Is someone who does the work willing to help redesign it?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	A named person who does the work has agreed to test early versions, with their manager's agreement and time set aside.
1	There is interest, but nobody has been asked formally or given time.
0	The team hasn't been told, or is hostile because it expects job cuts.

Collect it: ask the team lead who would volunteer, then ask that person.

Ask: the team lead and two team members.

If you scored 0 or 1: explain the purpose from question 3 to the team in person, including what happens to the hours saved. Ask for one volunteer. People who do the work spot the cases that will break a build within ten minutes; without them you find out in production.

7. Data (questions 9 to 13)

This is where good ideas should die, and it is far cheaper for them to die here than four weeks into a build. Answer for the data the priority process needs. Question 10 is a gate. Run the data readiness test further down before scoring this section; self-assessments overrate data more than anything else.

9. Do you know which systems hold the data the priority process needs?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	A short data map for the process: each input, the system it lives in, the format and who owns it.
1	The main systems are known but spreadsheets, inboxes and paper inputs aren't mapped.
0	"Somewhere in the CRM, probably."

Collect it: follow ten real items through the process from arrival to completion and note every place information is read from or written to.

Ask: the person who does the work; then IT to confirm the systems.

If you scored 0 or 1: draw the map. A table with four columns (input, system, format, owner) is enough. It takes an afternoon and is the first thing any builder will ask for.

10. Can that data be reached without someone running a manual export? (Gate)

SCORE	EVIDENCE THAT JUSTIFIES IT
2	Each input is reachable through an API, a supported connector, a database you can query or a shared mailbox, and you have confirmed this with the system administrator or vendor.
1	Some inputs are reachable; others depend on scheduled exports, or access exists in principle but has never been tested.
0	Every month someone exports a file, fixes it by hand and emails it on, or the data is on paper or in PDFs nobody can reach programmatically.

Collect it: for each system on the data map, check the vendor documentation for an API or connector, and ask the administrator whether a read-only service account can be created.

Ask: IT or your managed service provider, and the vendor's support desk.

If you scored 0 or 1: this is a gate, so fix it before anything else in the data section. Often the fix is a licence tier, a connector or a service account rather than a project. If a core system genuinely has no way in, that system, not AI, is the first problem to solve, and the opportunity drops down the ranking.

11. Is it good enough to trust: do two systems agree on the same figure?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	You have checked a key figure (customer count, revenue by customer, open orders) across two systems and it matches, or the differences are known and explained.
1	Differences exist and are tolerated, or nobody has checked but nobody complains.
0	Meetings regularly start with an argument about whose number is right.

Collect it: pick one figure the process depends on and pull it from both systems for the same date. Then run the 20-record test in the data readiness section.

Ask: finance and the process owner.

If you scored 0 or 1: list the differences, decide which system is the source of truth for each field, and fix the top three at source. If you scored 2 without checking, check.

12. Does each important dataset have an owner responsible for its accuracy?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	A named owner for each dataset on the data map, who fixes errors at source and knows it is their job.
1	Ownership is understood informally, or exists for some datasets and not others.
0	Errors are fixed downstream in spreadsheets, differently in each team.

Collect it: add an owner column to the data map and ask each named person if they agree.

Ask: the heads of the functions that own each system.

If you scored 0 or 1: name owners for the data the priority process needs. Ownership of everything can wait. An AI system makes poor source data visible quickly, and someone has to be responsible for fixing it.

13. Do you know which of it is personal, special category or contractually confidential?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	Data is classified (a green, amber and red scheme is enough), the classification is in the policy, and you have checked client contracts for restrictions on where data can be processed.
1	People know personal data when they see it, but there is no scheme, or contracts haven't been checked.
0	Nobody has looked, and staff are pasting client documents into personal accounts.

Collect it: list the data types in the priority process and classify each. Ask legal or the account managers whether the top ten client contracts restrict data location, subprocessors or automated processing.

Ask: your data protection lead, legal or company secretary, and account managers for key clients.

If you scored 0 or 1: adopt a three-colour scheme. Our AI policy template includes one you can copy. Check client contracts before any client data goes into an AI tool, not after.

8. Systems and integration (questions 14 to 16)

You don't need modern systems to start, but you need systems someone can connect to and someone with the authority to connect them.

14. Are the core systems for the priority process current, supported and able to connect to other tools?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	Cloud or current on-premise versions under vendor support, with documented APIs or standard connectors for Microsoft 365, Google Workspace or common automation platforms.
1	Supported but old, or connectable only through a costly add-on or the vendor's own professional services.
0	Out of support, heavily customised, or with no supported way in.

Collect it: list the systems, their version, support end date and integration options. The vendor's website and your support contract usually answer this.

Ask: IT or your managed service provider, and finance for the contract terms.

If you scored 0 or 1: price the connector or upgrade. If the core system is due for replacement within 18 months, make API access a requirement of the replacement and choose a priority process that doesn't depend on it in the meantime.

15. Does someone inside the business administer those systems and can grant access?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	An internal or outsourced IT lead with admin rights who can set up a service account within a few days, and a documented process for doing so.
1	Admin access exists but sits with one person, or requests take weeks through a supplier.
0	Admin access sits with a former employee, or with a supplier who rarely responds.

Collect it: ask IT to show you who holds global admin on Microsoft 365 (or Google Workspace) and on each core system, and how long the last access request took.

Ask: IT or the managed service provider.

If you scored 0 or 1: recover admin access and put two named people on it. This is also a security finding in its own right and belongs on the risk register.

16. Do you know what AI features you already pay for in existing software?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	A list of AI features in Microsoft 365, the CRM, finance, HR and helpdesk systems, whether each is switched on, and usage data for any paid AI licences.
1	You know about the Copilot or ChatGPT licences, but not the AI features bundled into other systems, or you have no usage data.
0	Nobody has looked; teams are buying separate tools to do jobs existing software already does.

Collect it: export 90 days of usage from the Microsoft 365 or Google Workspace admin centre, list software subscriptions from finance, and check each vendor's release notes for AI features.

Ask: IT and finance.

If you scored 0 or 1: run the one-week licence audit in you bought the licences and nothing changed. It is often the quickest saving on the scorecard.

9. People and skills (questions 17 to 19)

Limited skills was the most cited barrier among UK businesses already using AI in DSIT's research, and the gap was widest in mid-sized firms. Score these strictly.

17. Do you know which AI tools staff already use, including personal accounts?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	An anonymous staff survey run in the last six months with a response rate above 40%, the results written up, and a check of network or browser data where IT can provide it.

SCORE	EVIDENCE THAT JUSTIFIES IT
1	Anecdotes, a partial survey, or IT data with no staff view.
0	Leadership assumes nobody uses AI because nothing has been approved.

Collect it: run the shadow AI survey in this guide. Ask IT whether web filtering or endpoint tools can report visits to the main AI services.

Ask: all staff, anonymously; IT for the technical view.

If you scored 0 or 1: run the survey before writing or revising policy. Policies written without it tend to ban what people already rely on.

18. Has more than half your workforce had practical AI training on their own work?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	Training records showing more than half of staff attended hands-on sessions where they used an approved tool on a real task from their own job, covering the data rules.
1	Some teams trained, or training was a lecture, webinar or vendor video.
0	No training, or a link to online courses with no record of completion.

Collect it: training records from HR or the learning platform. Ask three people what they built or changed after the training.

Ask: HR and a sample of staff.

If you scored 0 or 1: start with the team that owns the priority process and the managers of affected teams, not the whole company. The ONS found only 11% of businesses have trained more than half their workforce, so a 2 here is unusual. What training that changes behaviour looks like is in how to train staff on AI. If you sell into the EU, note that the EU AI Act's AI literacy duty has applied since February 2025; see when it applies to UK businesses.

19. Do the managers of the affected teams have the time and willingness to change how their team works?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	The manager has agreed in writing to protect time for testing and adoption, and has a target in their objectives tied to question 4.
1	Supportive in principle but at capacity, or not yet asked.
0	Sees the project as head office's idea, or is actively resisting.

Collect it: ask the manager how many hours a week their team could give to testing over the next eight weeks, and what would have to stop.

Ask: the manager of the team that owns the priority process.

If you scored 0 or 1: agree what comes off the team's plate during the build. A manager with no time will quietly let the new process lapse, and the reasons AI pilots fail are mostly of this kind.

10. Governance and risk (questions 20 to 23)

Governance should be proportionate. A 150-person firm needs a short policy, an approved tool list, a training record and an incident route, not a governance board. Question 20 is a gate.

20. Do you have a written AI policy that people have been trained on? (Gate)

SCORE	EVIDENCE THAT JUSTIFIES IT
2	A short AI acceptable use policy with data rules, signed off by the CEO, with a review date, covered in training, with a record of who has read it.
1	A policy exists but hasn't been trained on, is out of date, or bans everyday use so is ignored.
0	No policy, or a paragraph in the IT policy that predates generative AI.

Collect it: the policy, its sign-off date and the training or acknowledgement record.

Ask: whoever owns HR or IT policy; then three staff, "What does our AI policy say you can't do?"

If you scored 0 or 1: adapt our AI acceptable use policy template in a week, using the shadow AI survey results to decide which tools to approve. Then train on it. A policy nobody has been taught is a PDF.

21. Is there an approved tool list, with business accounts rather than personal ones?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	Named approved tools on business or enterprise plans with suitable data protection terms (no training on your data, a data processing agreement, known data location), issued by the company, and a route to request new ones.
1	One approved tool, but not available to everyone, or no process for requests so people use alternatives.
0	Staff use free personal accounts because there is no alternative.

Collect it: the approved list, the contracts or terms for each tool, and the number of licences against headcount.

Ask: IT and procurement.

If you scored 0 or 1: approve at least one general-purpose assistant on a business plan for everyone who needs one. Microsoft's research found 28% of UK employees said their employer offered no approved option; banning tools without offering one pushes use underground. Copilot vs ChatGPT vs Claude for business compares the main options.

22. Would you know what to do if confidential or personal data went into an AI tool by mistake?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	A reporting route staff know about, a named data protection lead, and an incident procedure that covers AI tools, including the decision on whether to report a personal data breach to the ICO within 72 hours.
1	A general data breach procedure exists but doesn't mention AI tools, or staff wouldn't know to use it.
0	It would be found out by accident, if at all.

Collect it: the incident procedure, and a tabletop test: "An account manager pasted a client's payroll file into a free chatbot yesterday. What happens now?"

Ask: the data protection lead and IT.

If you scored 0 or 1: add an AI section to the incident procedure, tell staff how to report, and make clear that reporting quickly will not be punished. Under UK GDPR a notifiable personal data breach must be reported to the ICO within 72 hours of becoming aware of it, so the route has to be quick.

23. Do you know whether any planned use involves decisions about people, regulated advice or EU customers?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	Every planned and current use has been checked; anything touching hiring, performance, credit, insurance pricing, regulated advice or EU customers has a data protection impact assessment and a named approver.
1	Known high-risk uses are checked, but there is no routine check for new ones.
0	Features have been switched on (CV screening, automated credit decisions) without anyone asking.

Collect it: list current and planned uses from the survey and the opportunity list, and mark any that make or shape decisions about individuals. The ICO's AI and data protection risk toolkit is a useful structured check.

Ask: HR, compliance (and your FCA or other regulatory lead if you have one), and the data protection lead.

If you scored 0 or 1: add a three-question screen to the tool request process: does it make or shape decisions about people, does it involve regulated activity, does it touch EU customers or staff? A yes triggers a DPIA and sign-off. Most mid-market businesses should simply require a person to make decisions about individuals.

11. Money (questions 24 to 25)

24. Is there a budget, even a small one, set aside for a first build and its running costs?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	An agreed amount for a first build plus a monthly line for licences, model usage and upkeep, approved by the CFO, with a named approver for spend.
1	Money is available "if the case is good enough", with no threshold for what good enough means.
0	No budget, and no route to one this financial year.

Collect it: the budget line or board minute. If it doesn't exist in writing, it is a 1 at best.

Ask: the CFO.

If you scored 0 or 1: agree a ring-fenced amount and the payback threshold that releases it (for example, payback inside 12 months at half the expected benefit). How much it costs to build an AI agent gives typical ranges.

25. Do you know what an hour of the affected team's time costs, so you can work out payback?

SCORE	EVIDENCE THAT JUSTIFIES IT
2	A fully loaded hourly cost for each team in scope (salary, employer's National Insurance, pension, benefits, divided by productive hours), agreed with finance.
1	Salaries are known but nobody has worked out a loaded hourly cost, or it uses contracted rather than productive hours.
0	Savings will be argued about after the fact.

Collect it: payroll cost by team from finance, divided by productive hours (typically 1,500 to 1,700 a year per full-time person after holidays, training and absence).

Ask: finance.

If you scored 0 or 1: work it out now and use the same rate for benefits and for internal costs. Using one rate on both sides is what makes the business case credible to a CFO.

12. Scoring method and bands

Add the 25 scores for a total out of 50. Then read the result in this order.

1. Gates. Questions 1, 2, 10 and 20. A zero on any of them means fix it first. A business scoring 38 with no owner will still produce a pilot that stalls. The Excel scorecard flags each gate in red if it scores zero and amber if it scores one.

2. Weakest section. Work out each section as a percentage of its maximum. The lowest is where effort goes first, after the gates. Most mid-sized businesses are weakest on people or data; very few are weakest on strategy.

3. Total and band.

SCORE	BAND	WHAT IT MEANS	WHAT TO DO NEXT
0 to 19	Foundations	Something basic is missing: an owner, a measured problem or reachable data.	Name an owner, pick one process, take a baseline, approve one business AI tool and write a short policy. Buy and build nothing yet. Reassess in 90 days.
20 to 34	First build	Ready for one contained build on the best-prepared process.	Fix that process's specific data gaps, run a fixed-scope build with a measured before and after, and train the team that will use it.
35 to 50	Scale	Ready to run several builds in parallel.	Build a ranked, costed portfolio across teams and run two or three builds at once, each with an owner, baseline and kill criteria.

4. Boundary rule. If you are within two points of a band boundary (18 to 21, or 33 to 36), take the lower band. Self-assessments flatter, and the flattery concentrates in data and training.

5. Section floor. If any section scores below 40% of its maximum, cap the band at first build until that section is fixed, and put that section into weeks 1 to 3 of the plan. A total of 36 with a people score of 1 out of 6 is not ready to scale; it is ready for one build and a training programme.

6. Readiness is per process. If you have scored one process, the band applies to that process. A second process may score very differently on questions 5 to 16. The Excel workbook lets you copy the scorecard tab for each process you assess.

> **Tip.** Score twice: once by the leadership team and once by the people who do the work, then compare. The gap between the two scores is often the most useful finding of the whole exercise.

13. Data readiness test

Questions 9 to 13 are the most overscored on the sheet. This test takes half a day per use case and replaces opinion with a result. Run it with the process owner and someone with system access, on real data, not a tidy sample.

The 20-record test. Pull the last 20 real items the process handled (invoices, enquiries, claims, orders), including whatever arrived last thing on a Friday. For each, check:

☐ **Findable.** Every input the task needs can be located in a known system within two minutes.

☐ **Reachable.** Each input can be read by software without someone exporting, retyping or scanning it. (Question 10.)

☐ **Complete.** At least 18 of the 20 records have every field the task needs.

☐ **Consistent.** The same field means the same thing in every record (one date format, one customer identifier, one set of product codes).

☐ **Agreed.** Key values match across the two systems that hold them, or the difference is explained. (Question 11.)

☐ **Current.**
The data is no more than a day old when the task needs it, or the lag doesn't matter.

☐ **Owned.**
Someone fixes errors at source when they are found. (Question 12.)

☐ **Classified.** You know which records contain personal, special category or contractually restricted data, and which tools may see them. (Question 13.)

☐ **Labelled outcome.**
For each record you can see what the right answer was (the code chosen, the route taken, the reply sent), so a build can be tested against it.

☐ **Exceptions known.**
The process owner can name the five most common exceptions and how often each occurs.

Reading the result.

CHECKS PASSED	DATA READINESS	WHAT IT MEANS FOR THE USE CASE
9 or 10	Ready	Data availability scores 4 or 5 in the opportunity matrix. Proceed to scoping.
6 to 8	Fixable	Score 3. List the failed checks as tasks with an owner and a date; most take days, not months.
3 to 5	Weak	Score 2. The use case is flagged "Fix data first" in the matrix. Price the data work before going further.
0 to 2	Not ready	Score 1. Park the use case and choose another.

Two checks deserve special attention. "Reachable" is the gate in question 10, and a failure usually means a licence, a connector or a service account, which is cheap. "Labelled outcome" is the one most often missed: without examples of the right answer, nobody can measure whether a build is accurate, and the pilot ends in opinions.

14. Use-case opportunity scoring

The scorecard tells you whether you are ready. This method tells you what to do first. Score every candidate on five criteria, from 1 to 5, using the definitions below. The Excel matrix applies the weights, calculates a priority score out of 100, ranks the list and applies the flags.

The five criteria

CRITERION	WEIGHT	1	3	5
Value	30%	Under £5k a year	£20k to £50k a year	Over £100k a year
Feasibility	20%	Unproven technique, complex judgement, no way to check output quickly	Proven technique, some judgement, output checkable in minutes	Common pattern (extraction, triage, drafting), output checkable at a glance
Data availability	20%	Data readiness test: 0 to 2 checks passed	6 to 8 checks passed	9 or 10 checks passed
Risk (5 = lowest)	15%	Decisions about people, regulated advice, or red data, with no DPIA	Amber data, human review in place, contained impact if wrong	Green or amber data, internal use, a person checks every output
Time to value	15%	More than six months to first measured benefit	Two to four months	Under six weeks

Use these value bands for the in-between scores: 2 is £5k to £20k a year and 4 is £50k to £100k. For the other criteria, 2 and 4 sit between the descriptions either side.

Value in pounds. Work out value as hours returned a month, times the loaded hourly rate from question 25, times 12. Use a realistic share of the hours, not all of them: for a well-defined extraction or triage task, 50% to 70% of the measured time is a reasonable planning assumption; for drafting tasks where people edit the output, 30% to 50%. If the value is revenue or margin rather than hours, put an annual figure in the "other value" column and discount it by half, because revenue claims are the ones that most often fail to appear.

The formula. Priority score = (Value × 0.30 + Feasibility × 0.20 + Data × 0.20 + Risk × 0.15 + Time to value × 0.15) × 20. The result runs from 20 to 100. This is consistent with the value and readiness ranking in our AI value creation plan: value carries the most weight, and the readiness criteria decide what can start now.

The flags. Weighted scores hide fatal weaknesses, so three rules override the ranking:

CONDITION	FLAG	WHAT IT MEANS
Risk scores 1	Park: risk	Don't proceed until a DPIA, legal view or regulatory check is done and the risk score rises.
Data scores 1 or 2	Fix data first	Keep it on the list with the data work priced as its first task.
Time to value scores 1 or 2	Later	A second-wave candidate; don't start it in the first 90 days.
None of the above	Go	Eligible for the first build or quick-win list.

How to rank the first ten

1. **Long-list 15 to 30 candidates.** Sources: the team walk-throughs from question 5, free-text answers from the shadow AI survey, unused features from question 16, and the leadership problems from question 1.
2. **Kill the obvious misfits early.** Anything done less than monthly, anything nobody can measure, anything with no owner. Write down why each was killed; a list with no rejections means nobody tested anything.
3. **Run the data readiness test** on the top 10 to 12 by rough value. Don't score data availability from memory.
4. **Score in one session** with the process owners in the room. Score value from numbers, not enthusiasm.
5. **Sort by priority score, then apply the flags.** Your first build is the highest-scoring "Go" item in the band your scorecard allows.
6. **Balance the top ten.** Include at least one quick win that uses tools you already pay for (visible within weeks) and one larger build with real value. Don't put three builds on the same team at once.
7. **Write the top three as one-paragraph business cases:** hours, value, build cost, monthly running cost, payback at full and half capture, owner, kill criteria.
8. **Revisit the ranking every quarter.** Fixed data moves items up; disappointing results move them down.

> **Red flag.** If the top of your list is a customer-facing chatbot, a pricing model or anything an investor asked about at the last board meeting, rescore value from measured hours. Visible ideas are consistently overscored on value and underscored on risk.

15. Shadow AI survey

The DSIT research interviewed the person responsible for technology, so it couldn't see shadow AI; neither can your IT team from licence data alone. This survey is the cheapest way to find out what is really happening, and it produces the best long-list of use cases you will get. It is in the Excel workbook in a form you can paste into Microsoft Forms or Google Forms.

How to run it.

- Make it anonymous and say so. Collect department, but only report groups of five or more responses.
- Send it from the CEO or COO with a one-paragraph note: we want to understand how people use AI so we can provide better tools and rules; nobody will be in trouble for anything they report.
- Keep it open for ten working days with one reminder. Aim for a response rate of at least 40%; below 25%, the result is mostly the enthusiasts.
- It takes about seven minutes to complete.

Suggested introduction. "We want to give everyone AI tools that are safe and useful. To do that, we need to know how you use AI now, including free and personal tools. This survey is anonymous. Nothing you say will be used against you, and there are no wrong answers."

The questions.

#	QUESTION	RESPONSE OPTIONS
1	Which department or team do you work in?	List of departments (report only groups of five or more)
2	In the last month, how often have you used an AI tool for work?	Never; once or twice; weekly; most days
3	Which AI tools have you used for work? Tick all that apply.	ChatGPT; Microsoft Copilot (work account); Copilot or Bing (personal); Google Gemini; Claude; Perplexity; AI features in other apps (Grammarly, Canva, Notion and similar); AI meeting note-takers; AI browser extensions; other (please say)
4	Which account did you use?	Account provided by the company; free personal account; paid personal account I pay for myself; not sure
5	What do you use AI for? Tick all that apply.	Drafting emails or messages; writing reports or presentations; summarising documents; research; spreadsheets, formulas or data analysis; writing code; meeting notes; translating; images or design; other
6	What kinds of information have you put into an AI tool? Tick all that apply.	Only public information; internal documents; customer or supplier names and contact details; pricing or financial information; information about colleagues (HR, performance, pay); client documents or confidential material; none of these; not sure
7	In a typical week, roughly how much time does AI save you?	None; under 1 hour; 1 to 3 hours; 3 to 5 hours; more than 5 hours
8	If you use a tool the company hasn't provided, why?	There is no approved tool; the approved tool doesn't do what I need; I didn't know there was one; I use it at home and I'm used to it; other

#	QUESTION	RESPONSE OPTIONS
9	How often do you check AI output before using it?	Always; mostly; sometimes; rarely
10	Have you had any training on using AI at work?	None; I taught myself; online videos or courses; a session run by the company
11	Which task in your job would you most like AI to take off you? Please describe it and roughly how long it takes each week.	Free text
12	What, if anything, worries you about AI at work?	Free text

How to read the results. Report six numbers, by department where groups are large enough:

MEASURE	FROM	WHAT IT TELLS YOU	SCORECARD LINK
Weekly or more often users	Q2	How embedded AI already is	Q17
Share using personal accounts	Q3, Q4	Size of the shadow AI exposure	Q21
Share entering amber or red data	Q6	The real data protection risk	Q13, Q22
Share with no training	Q10	The training gap	Q18
Share who check output only sometimes or rarely	Q9	Quality risk in customer-facing work	Q20
Most requested tasks	Q11	Your use-case long-list	Q5

For comparison, Microsoft's national figures were 71% having used unapproved tools and 51% doing so weekly. If your numbers are much lower, check your response rate before celebrating. Close the loop within a month: tell staff what you found and what you are changing, usually an approved tool, a short policy and training. A survey that disappears makes the next one worthless.

16. 90-day plan by score band

Each plan follows the same three phases as our 90-day AI implementation plan: weeks 1 to 3, weeks 4 to 9, weeks 10 to 13. Pick the plan for your band, after applying the boundary and section floor rules.

Foundations band (0 to 19)

WEEKS	ACTIONS	OWNER	EXIT TEST
1 to 3	Name the accountable owner (Q2); agree the purpose sentence (Q3); run the shadow AI survey (Q17); pick the priority process and start a four-week baseline (Q4, Q7).	CEO, then the owner	Owner named in objectives; survey closed
4 to 9	Approve one business-grade assistant and switch off routes to free accounts where practical (Q21); write and sign off a short policy and train every manager on it (Q20); map the priority process data and fix access (Q9, Q10); work out loaded hourly costs (Q25).	Owner, IT, HR, finance	Policy live and trained; data reachable

WEEKS	ACTIONS	OWNER	EXIT TEST
10 to 13	Hands-on training for the priority team on their own tasks (Q18); long-list and score use cases; reassess with this scorecard.	Owner, process owner	Score of 20 or more with no gate at zero

Spend in this band is mostly internal time plus licences. Don't commission a build until the reassessment clears 20.

First build band (20 to 34)

WEEKS	ACTIONS	OWNER	EXIT TEST
1 to 3	Fix any gate at 1; run the data readiness test on the top three candidates; score and rank opportunities; write business cases for the top three; agree budget and kill criteria for one first build (Q24).	Owner, process owner, finance	One build approved with scope, price, baseline and kill criteria
4 to 9	Build inside the real systems with the people who do the work; DPIA if personal data is involved (Q23); start one quick win using licences you already pay for (Q16); train the affected team as the build goes live (Q18).	Delivery lead, process owner, IT	Build live on real work, handling the straightforward majority of cases with human review
10 to 13	Measure against the baseline; first monthly hours-returned report to the leadership team; fix the weakest section; reassess and choose the second build.	Process owner, finance	Measured result against baseline; decision to keep, change or stop

This is the band most mid-sized businesses land in. The temptation is to start three things. Start one, and do it properly.

Scale band (35 to 50)

WEEKS	ACTIONS	OWNER	EXIT TEST
1 to 3	Rank opportunities across all functions; choose two or three builds in different teams, each with its own owner, baseline and kill criteria; set up a monthly portfolio review.	Owner, function heads	Portfolio of two or three approved builds
4 to 9	Run the builds in parallel; standardise the approval route for new tools and use cases; extend hands-on training beyond the first teams; decide what to build, buy or leave to vendor roadmaps.	Function heads, delivery lead, IT	Each build live on real work
10 to 13	Report hours returned and where they went (redeployed, not backfilled, growth absorbed) for each build; stop anything that misses its kill criteria; publish the next quarter's ranked list.	Owner, finance	A board-ready portfolio report, including what was stopped

At this band the risk shifts from "nothing happens" to "too many things happen with no measurement". The portfolio review is what prevents that.

17. Presenting results to the board

Boards don't need the 25 questions. They need to know where the business stands, what it will do, what it will cost and how they will know it worked. One page, presented in ten minutes, is enough. The same page should be repeated every quarter with the numbers updated; the change is what tells the board whether management has a grip. Our article on board questions about AI covers what directors should ask in return.

The one-page structure.

1. **Headline.** Total score, band, and one sentence: "We are ready for one contained first build, on renewal preparation, starting in November."
2. **Gates.** The four gate questions with their scores, red, amber or green. If any is red, what will fix it and by when.
3. **Section chart.** The seven section percentages as a bar or radar chart (the Excel workbook draws this), with the weakest section named.
4. **What staff are already doing.** Three numbers from the shadow AI survey: weekly users, share on personal accounts, share entering amber or red data. This is usually the number that gets the board's attention.
5. **Top three opportunities.** For each: hours a month, annual value at half capture, build cost, monthly running cost, payback, owner.
6. **The 90-day plan.** Three lines, one per phase, with the exit test for each.
7. **The decision needed.** Budget, the owner's time, and any policy or tool approval.
8. **What we will report next quarter.** The reassessed score, the measured result against baseline and anything stopped.

What to leave out. A maturity level with no evidence behind it; comparisons with vendor survey "industry averages" that don't match your size or sector; a long list of use cases with no numbers; and vendor names in the headline. Directors remember the vendor and forget the decision.

Questions to prepare for. Who owns this, and how much of their time does it take? What is the baseline? What happens to the hours saved? What data are staff putting into tools today? What would make us stop? If the answers to those five are on the page, the meeting will be about the decision rather than the deck.

18. How this compares with published maturity models

Published maturity models are useful, and most were designed for large enterprises or to support a vendor's sales process. Here is how the main ones compare, and where they leave gaps for a business of 50 to 1,000 people.

MODEL	STRUCTURE	BEST FOR	GAP FOR A MID-SIZED UK BUSINESS
Gartner AI Maturity Model and assessment	Seven pillars (strategy, data, governance, engineering, operating model, culture, AI product and value); organisations are grouped in five stages from foundational to transformational	Large enterprises benchmarking against peers, with a Gartner subscription	The full diagnostic and benchmarks are Gartner research products; questions assume a dedicated AI function
Microsoft AI Readiness Wizard	Ten questions across five areas; places you in one of five stages: exploring, planning, implementing, scaling, realizing (Microsoft, 2024)	A quick, free first view, especially for Microsoft 365 customers	Ten questions can't test evidence; recommendations point to Microsoft adoption resources
MIT CISR Enterprise AI Maturity Model	Four stages: experiment and prepare; build pilots and capabilities; develop AI ways of working; become AI future ready	Explaining to a board why moving from pilots to ways of working matters financially	Built on a 2022 survey of large enterprises, before generative AI was widespread; not a scoring tool
Cisco AI Readiness Index	Six pillars: strategy, infrastructure, data, governance, talent, culture; respondents grouped from pacesetters to laggards	Infrastructure and security planning in large IT estates	Infrastructure-heavy; questions about networks and data centres rarely apply to a 200-person firm on Microsoft 365
Alan Turing Institute (AI Skills for Business Competency Framework; AI Ethics and Governance in Practice)	Skills framework with four learner personas (AI citizens, workers, professionals, leaders); governance workbooks built on a process-based governance framework	Designing training (Q18) and governance for AI you build yourself	Not a readiness or maturity score; the governance workbooks were written mainly for public sector teams
DSIT AI Management Essentials (draft)	Draft self-assessment of AI governance processes, consulted on from November 2024 to January 2025	A second pass on questions 20 to 23 if you build or sell AI	A consultation draft, not a live tool; the government response (dated December 2025) said it would develop SME-focused guidance instead
Innovate UK BridgeAI and Digital Catapult toolkit	Digital maturity, data readiness, data ethics and MLOps tools; a four-stage adoption framework (strategy, data, build, implement)	UK SMEs and startups building AI products	Aimed more at building AI than adopting it in operations; less on staff use, policy and payback
Google Cloud AI Adoption Framework	Six themes (learn, lead, access, scale, automate, secure) across three phases: tactical, strategic, transformational	Organisations with data science and machine learning teams	Assumes in-house ML engineering; little on generative AI use by ordinary staff

A few findings from these models are worth taking into a board discussion. Gartner's June 2025 survey of 432 organisations found 45% of leaders in high-maturity organisations said their AI initiatives stayed in production for three years or more, against 20% in low-maturity ones, and 91% of high-maturity organisations had appointed dedicated AI leaders: question 2, in other words. MIT CISR's model, published in December 2024, found 28% of

enterprises in stage 1 and 34% in stage 2, both with growth and profit below their industry average, while the 31% in stage 3 and 7% in stage 4 were above it. Cisco's 2025 index, from a survey of more than 8,000 leaders reported by The Register, found about 13% of organisations qualified as pacesetters for the third year running.

What this scorecard does differently. It is written for one size of business and one country. Each question has evidence for every score, so the result can be checked. It is scored per process, because that is how builds succeed or fail. It includes the staff view that leadership-only surveys miss, a use-case ranking in pounds, a data test you run on real records, UK-specific governance (UK GDPR, the ICO, client contracts) and a plan for each band. It produces a decision, not a label.

A rough mapping, for those who need one. If an investor or group function uses a maturity model, this table translates. Treat it as indicative: the models measure different things.

THIS SCORECARD	MIT CISR STAGE	MICROSOFT STAGE	GARTNER STAGE
0 to 19, foundations	1: experiment and prepare	Exploring or planning	Foundational
20 to 34, first build	2: build pilots and capabilities	Implementing	Emerging to operational
35 to 50, scale	3: develop AI ways of working	Scaling	Operational to scaled

19. Worked example 1: Holbeck Fabrication Ltd

> **Illustrative example.** Holbeck Fabrication Ltd is a fictional company. The figures are invented to show the method, not drawn from any real business.

Holbeck is a fictional structural steel fabrication and machining business near Wakefield with 160 staff and revenue of about £28m. The managing director wants to "do something with AI" after a competitor mentioned it at a trade event. The leadership team picks estimating as the priority process: five estimators turn customer drawings and tender documents into quotes.

Evidence collected. A one-week time sample shows estimators spend about 22 hours each a week on quotes, roughly 475 hours a month across the team, and win about one in six. Quote turnaround averages six working days. Drawings arrive as PDFs by email; rates and material prices sit in a spreadsheet one senior estimator maintains; job history is in an old on-premise ERP installed in 2011 with no API, and the vendor offers only a monthly CSV export. The survey (52% response) found 58% of respondents used AI tools at least weekly, 46% of them on personal accounts, and that estimators were pasting tender documents into a free chatbot to summarise specifications. There is no AI policy and no approved tool.

Scores.

SECTION	SCORE	MAX	%	NOTES
Strategy and ownership	1	8	13%	Problem named but unmeasured until this exercise (Q1 = 1); no owner (Q2 = 0); no agreed purpose or baseline
Processes	5	8	63%	High volume (Q5 = 2); each estimator prices differently (Q6 = 0); time now measured, wins not analysed (Q7 = 1); senior estimator keen to help (Q8 = 2)

SECTION	SCORE	MAX	%	NOTES
Data	3	10	30%	Systems known roughly; ERP data only by manual export (Q10 = 0); rates spreadsheet unowned (Q12 = 0)
Systems and integration	2	6	33%	ERP out of support (Q14 = 0); IT outsourced and slow (Q15 = 1); Microsoft 365 in place (Q16 = 1)
People and skills	2	6	33%	Survey now run, but no IT view of usage (Q17 = 1); no training (Q18 = 0); production managers willing but stretched (Q19 = 1)
Governance and risk	1	8	13%	No policy (Q20 = 0); personal accounts (Q21 = 0); no incident route (Q22 = 0)
Money	2	4	50%	No budget (Q24 = 0); hourly costs known (Q25 = 2)
Total	**16**	**50**	**32%**	**Foundations band; three gates at zero (Q2, Q10, Q20)**

Reading it. Sixteen points puts Holbeck in the foundations band, with three of the four gates at zero. The estimating opportunity is real (475 hours a month at a loaded £36 an hour is about £205,000 a year of estimator time) but not buildable yet: estimators price differently, the ERP can't be reached, and confidential tender documents are already going into personal chatbot accounts, which is the most urgent finding.

The plan. In weeks 1 to 3 the operations director is named owner with a day a week, and the business approves a business-grade assistant under its existing Microsoft 365 work accounts as the only route for tender documents, with a two-page policy. In weeks 4 to 9 the senior estimator writes one estimating method with the top ten exceptions, finance takes ownership of the rates spreadsheet and IT prices a supported connector or a scheduled database export from the ERP. In weeks 10 to 13 all five estimators get hands-on training on summarising tender specifications with the approved tool, which is a quick win needing no build, and the business reassesses. A projected score of 26 would move it into the first build band, with estimating as the first build.

20. Worked example 2: Aldermoor Insurance Services Ltd

> **Illustrative example.** Aldermoor Insurance Services Ltd is a fictional company. The figures are invented to show the method, not drawn from any real business. This is the company used in the Example tab of the Excel workbook.

Aldermoor is a fictional commercial insurance broker with 420 staff across offices in Leeds and Bristol, regulated by the FCA. Its COO has been asked by the board for an AI plan. The priority process is renewal preparation: account handlers assemble a renewal pack for each client from the expiring policy, claims history and updated client information, then present it to insurers. The team handles about 650 renewals a month.

Evidence collected. A four-week baseline found renewal pack preparation takes about 2.4 hours per renewal on average, roughly 1,560 hours a month across the account handling teams, of which about 260 hours is spent rekeying and summarising information already held elsewhere. The broking platform has a documented API; claims data arrives from insurers as spreadsheets in varying formats. Microsoft 365 Copilot licences were bought for 60 staff a year ago; the admin report shows 14 active weekly. The survey (61% response) found 63% used AI at least weekly, 31% via personal accounts, and 18% had entered client information into personal tools. There is an AI policy from 2024 that nobody has been trained on.

Scores.

SECTION	SCORE	MAX	%	NOTES
Strategy and ownership	5	8	63%	Problem measured (Q1 = 2); COO owns it but without protected time (Q2 = 1); purpose and success measures partly agreed
Processes	6	8	75%	High volume (Q5 = 2); method varies by office (Q6 = 1); baseline taken, error rate estimated (Q7 = 1); account handlers keen (Q8 = 2)
Data	7	10	70%	Systems mapped (Q9 = 2); broking platform reachable, insurer claims spreadsheets are not (Q10 = 1); data classified (Q13 = 2)
Systems and integration	4	6	67%	Platform current but integration needs an add-on (Q14 = 1); IT capable (Q15 = 2); Copilot usage known, other AI features not (Q16 = 1)
People and skills	2	6	33%	Survey now run, but IT has not checked usage data (Q17 = 1); no practical training (Q18 = 0); managers stretched (Q19 = 1)
Governance and risk	5	8	63%	Policy untrained (Q20 = 1); Copilot approved but not for everyone (Q21 = 1); incident route not AI-specific (Q22 = 1); regulated uses screened by compliance (Q23 = 2)
Money	3	4	75%	Budget available but no threshold (Q24 = 1); loaded costs agreed (Q25 = 2)
Total	**32**	**50**	**64%**	**First build band; no gate at zero; weakest section people**

Reading it. Thirty-two is in the first build band, just outside the boundary zone of 33 to 36. No gate is at zero, though three of the four are at 1. People and skills, at 33%, is below the 40% section floor, so even if the total were higher Aldermoor would be held to one build until training is fixed. The board finding that lands hardest is that 18% of respondents have put client information into personal AI tools while 46 of the 60 paid Copilot licences sit idle.

Use-case ranking. Ten candidates, scored in one session with the account handling and finance leads. Value uses a loaded rate of £38 an hour and hours at realistic capture.

RANK	USE CASE	SCORE	FLAG
1	Client service inbox triage and draft acknowledgements (180 hours a month)	87	Go
2	Renewal pack first draft from expiring policy and claims history (260 hours a month)	86	Go
3	Account handler file notes and meeting summaries with existing Copilot licences (150 hours a month)	84	Go
4	Prospect research briefs for the new business team (40 hours a month)	75	Go
5	Credit control chaser drafts (30 hours a month)	68	Go
6	Claims first notification summaries (60 hours a month)	67	Go
7	Expiring versus renewal policy wording comparison (90 hours a month)	64	Go
8	Insurer bordereaux reconciliation in finance (70 hours a month)	59	Fix data first
9	Risk selection and pricing recommendations (margin, about £160k a year before the 50% discount)	46	Park: risk
10	Website chatbot for client queries (20 hours a month)	44	Fix data first

The renewal pack draft scores 86, with value of 5 (260 hours a month, about £118,600 a year) and a data score of 4 after the claims spreadsheets failed one check (consistency) in the data readiness test. Inbox triage scores 87 on lower value but faster time to value. Aldermoor chooses the renewal pack as its first build because it is the measured problem in question 1, and treats the file notes item, which needs no build, as the quick win to get the idle Copilot licences used. The pricing recommendation is parked: it shapes decisions about which risks to place and on what terms, and needs compliance and Consumer Duty review before it can be scored above 1 on risk. The website chatbot sits at the bottom, as it usually does.

The 90-day plan and board page. The COO's diary is cleared for a day a week (fixing Q2). Weeks 1 to 3: retrain everyone on the existing policy, reissue the 46 idle Copilot licences to the account handlers who asked for them in the survey, and approve a fixed-scope renewal pack build with kill criteria (fewer than half of drafts usable with light edits by week 8). Weeks 4 to 9: build inside the broking platform with two account handlers testing, and standardise the claims spreadsheet intake with the three largest insurers. Weeks 10 to 13: measure against the baseline and report. The board page leads with the score, the 18% client data finding and the renewal case: at half capture, about £59,000 a year against a build of £25,000 to £35,000 and around £1,200 a month to run, a payback of roughly eight months at the midpoint.

21. Common mistakes

Scoring the company instead of a process. Averaging across functions produces a respectable score and no buildable process. Pick one and score that.

Scoring from the boardroom. Leadership scores data and training higher than the people doing the work. If only executives were in the room, take two points off the data section and one off people, then go and check.

Treating "don't know" as a 1. It is a 0. Not knowing who administers the ERP is a finding, not a gap in the questionnaire.

Ignoring the gates because the total looks good. A 36 with no owner stalls. A 32 with no reachable data builds on exports and breaks the first month someone is on holiday.

Running the staff survey after writing the policy. You end up banning the tools people rely on, and use goes further underground.

Ranking use cases by enthusiasm. Value must come from measured hours or a revenue figure you would defend to the CFO. A score of 5 on value with no number beside it is a 3.

Counting hours saved as cash. Hours returned become money only when they are redeployed, not backfilled, or absorb growth. Say which, or the board will discount the whole case.

Buying a platform to fix a readiness gap. No tool fixes missing ownership, an unmeasured process or an untrained team. Below 20, buy nothing but a business-grade assistant.

Assessing once. A single score is a snapshot. The value is in the movement, so reassess at 90 days and then every six months, with the evidence column updated.

Letting a maturity label replace a decision. "We are at stage two" is not a plan. "We will build the renewal pack draft by February, and stop if it misses these numbers" is.

22. Glossary

Baseline. A measurement of volume, time per item, error rate and cycle time for a process before any change, taken over at least a week of normal trading.

DPIA (data protection impact assessment). A UK GDPR assessment required for processing likely to result in a high risk to individuals, which includes many uses of AI on personal data.

Gate question. One of the four questions (1, 2, 10, 20) where a zero overrides the total score.

Hours returned. The reduction in hours a month spent on a process after a change, measured against the baseline.

Loaded hourly cost. Salary plus employer's National Insurance, pension and benefits, divided by productive hours rather than contracted hours.

Priority process. The one process used to answer questions 5 to 16; usually the one leadership most wants to fix.

Section floor. The rule that a section below 40% of its maximum caps the business at the first build band, whatever its total.

Shadow AI. Use of AI tools at work that the organisation has not approved or provided, often through personal accounts.

Time to value. The time from deciding to proceed to the first measured benefit in real work.

23. Sources and further reading

UK statistics cited.

- Office for National Statistics, AI in UK businesses: 2023 to 2026, July 2026. Business Insights and Conditions Survey, wave 159 fieldwork June 2026.
- Department for Science, Innovation and Technology, AI Adoption Research, January 2026. 3,500 interviews with UK private sector businesses with five or more employees, February to May 2025.
- Microsoft UK, Rise in shadow AI tools raising security concerns for UK organisations, October 2025. Censuswide survey of 2,003 UK employees.

Maturity models and readiness tools reviewed.

- Gartner, AI Maturity Assessment, and survey press release, June 2025.
- Microsoft, AI Readiness Wizard, and A strategic approach to assessing your AI readiness, November 2024.
- MIT CISR, Weill, Woerner and Sebastian, Building Enterprise AI Maturity, December 2024.
- Cisco, AI Readiness Index, 2025, with coverage in The Register, October 2025.
- The Alan Turing Institute, AI Ethics and Governance in Practice workbook series, and the AI Skills for Business Competency Framework developed for Innovate UK BridgeAI.
- DSIT, AI Management Essentials consultation, November 2024, and government response, dated December 2025.

- Innovate UK Business Connect and Digital Catapult, AI Adoption Assessment Toolkit and AI Adoption Stages Framework.
- Google Cloud, AI Adoption Framework.
- ICO, AI and data protection risk toolkit. The best free structured check for UK GDPR risk in AI use.

Related Anyway Ventures material. AI readiness assessment checklist, what an AI audit is, how to implement AI in a mid-market business, measuring AI ROI, AI acceptable use policy template and why AI pilots fail. If you want the answers tested against your systems and people, with a first build scoped and priced, that is the purpose of an AI audit.

24. About this resource

Written by Anyway Ventures, the York technology and AI practice led by Brannan Coady, former CEO of Netsells and YourParkingSpace. Version 1.0, October 2026. Free to use, adapt and share with attribution. Not legal, tax or investment advice.

ABOUT ANYWAY VENTURES

Senior hands for technology and AI decisions

Technology and AI work for private equity funds, venture investors and UK mid-market management teams: before a deal, during the hold and at exit. Every engagement is led by Brannan Coady, an exited founder and former CEO. Fixed price, built to be handed over.

FOR INVESTORS

Technical due diligence

Code, architecture, security, IP and one-to-one team interviews, every material finding costed. Red-flag review in five to seven days.

ANYWAYVENTURES.COM/TECHNICAL-DUE-DILIGENCE

AI due diligence

Whether a target's AI claims are real, defensible and theirs, and a costed AI value-creation plan for the IC.

ANYWAYVENTURES.COM/AI-DUE-DILIGENCE

AI disruption risk review

Is AI a threat to the business or its moat? One deal, or every holding in a portfolio ranked by exposure.

ANYWAYVENTURES.COM/AI-DISRUPTION-RISK

Technical due diligence for VCs

A two to three day check before the term sheet: the product tested on evidence, not the demo.

ANYWAYVENTURES.COM/VC-TECHNICAL-DUE-DILIGENCE

Retained AI operating partner

Senior AI capacity held by the fund and deployed into the company with the strongest AI value case.

ANYWAYVENTURES.COM/AI-OPERATING-PARTNER

Exit readiness and vendor technical due diligence

Find what a buyer's diligence will find, 12 to 24 months out, fix it, then publish a report buyers can rely on.

ANYWAYVENTURES.COM/EXIT-READINESS

FOR OPERATORS

AI audit and readiness assessment

Two to three weeks to a ranked, costed list of AI opportunities, the gaps to close and a 90-day plan.

ANYWAYVENTURES.COM/AI-AUDIT

Embedded AI delivery

Internal tooling, automation and AI agents in production inside your business, built with your team.

ANYWAYVENTURES.COM/EMBEDDED-AI-DELIVERY

Fractional Chief AI Officer

One senior person who owns AI on your leadership team, two to four days a month.

ANYWAYVENTURES.COM/FRACTIONAL-CHIEF-AI-OFFICER

Hands-on AI training

In-person training where every attendee builds a working tool on their own job.

ANYWAYVENTURES.COM/AI-TRAINING

Board and NED roles

Non-executive seats where technology, product and AI decide the outcome.

ANYWAYVENTURES.COM/BOARD-AND-NED

Every enquiry starts with the same free 30-minute call.

A straight answer on whether we can help, what it would cost and how long it would take.

EMAIL
brannan@anywayventures.com

BOOK A CALL
anywayventures.com/contact

MORE RESOURCES
anywayventures.com/resources

Anyway Ventures Limited · Registered in England and Wales, company number 16318119 · Registered office: Newbury House, Main Street, Little Ouseburn, York, YO26 9TD. This resource is general guidance, not legal, tax or investment advice. Free to use, adapt and share with attribution to Anyway Ventures. Version 1.0, October 2026.