

# 1Tech — AI Readiness Scorecard

Twelve questions that decide whether to invest, prepare, or walk away

## How to use this

Score each question 0 (no), 1 (partially) or 2 (confidently yes). Be harsh — the scorecard only works if the answers are honest. Twenty minutes with the people who own the process beats a month of workshops. Full reasoning for every question: [www.1tech.co/blog-ai-readiness-assessment.html](http://www.1tech.co/blog-ai-readiness-assessment.html)

## Group 1 — The Use Case

|                                                                                                                                                                                    | 0 | 1 | 2 |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|---|---|
| 1. <b>Can you state the problem without the word “AI”?</b><br>“Invoice processing takes 11 minutes each and errors cost £180K a year” is fundable. “Explore AI in finance” is not. |   |   |   |
| 2. <b>Is the task high-volume and pattern-rich?</b><br>AI compounds on repetition. A judgement made 10,000 times a month is a candidate; a bespoke quarterly decision is not.      |   |   |   |
| 3. <b>Would 90% accuracy be valuable?</b><br>If only perfection will do, AI drafts and a human decides — still valuable, but a different design and business case.                 |   |   |   |

## Group 2 — The Data

|                                                                                                                                                                                           | 0 | 1 | 2 |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|---|---|
| 4. <b>Does the data the task needs exist, digitally, today?</b><br>Not “could be collected” — exists now, in systems, with history.                                                       |   |   |   |
| 5. <b>Can you get to it?</b><br>APIs or database access, with permissions and a privacy basis. An integration project hiding inside an AI project should be named as such.                |   |   |   |
| 6. <b>Would you trust it in front of a customer?</b><br>If the CRM is 40% duplicates, the AI will be confidently wrong at scale. Cleanup is often phase one — but it must be in the plan. |   |   |   |

# lTech — AI Readiness Scorecard

Twelve questions that decide whether to invest, prepare, or walk away

## Group 3 — The Workflow

|                                                                                                                                                           | 0 | 1 | 2 |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------|---|---|---|
| <b>7. Is there a named owner whose process this is?</b><br>AI adopted into a workflow succeeds; AI installed next to one dies. No owner, no project.      |   |   |   |
| <b>8. Do you know where the human sits?</b><br>Reviewing everything, reviewing exceptions, or spot-checking — decide it; don't let it happen by accident. |   |   |   |
| <b>9. Can success be measured against today's baseline?</b><br>If you don't know current cost-per-case, error rate and cycle time, measure that first.    |   |   |   |

## Group 4 — The Risk

|                                                                                                                                                                                       | 0 | 1 | 2 |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|---|---|
| <b>10. What happens when it's wrong?</b><br>A mis-sorted email is cheap; a mis-priced quote is not. Map the error cost; size the guardrails to it.                                    |   |   |   |
| <b>11. Is the use regulator- and customer-explainable?</b><br>Decisions about people carry duties — explainability, appeal routes, EU AI Act categories. Know yours before you build. |   |   |   |
| <b>12. Can you run it where your data must live?</b><br>Sovereignty and confidentiality constraints decide the architecture — cheaper decided now than retrofitted.                   |   |   |   |

**Your total \_\_\_\_\_ / 24**

|                 |                                                                                                                                                                                                  |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>18–24</b>    | Ready. Pick your first use case and move — the risk now lives in delivery, not readiness. A two-week AI Opportunity Sprint turns the top candidate into an architecture and a fixed-price quote. |
| <b>10–17</b>    | Nearly. Invest 4–8 weeks fixing your two weakest answers first — usually data access or workflow ownership. It will halve the project risk for a fraction of a build's cost.                     |
| <b>Under 10</b> | Not yet — and discovering that today is the cheapest possible outcome. Fix the foundations; the AI opportunity will still be there, larger, when you return.                                     |

## What next

Whatever you scored, the numbers behind questions 2, 9 and 10 are worth two minutes in our ROI calculator: [www.ltech.co/roi.html](http://www.ltech.co/roi.html). And if you'd like the assessment run properly — with your systems mapped and your top-3 opportunities ranked by ROI — that's the AI Opportunity Sprint: two weeks, £4,950 fixed, credited in full against any build. Book 15 minutes at [www.ltech.co/sprint.html](http://www.ltech.co/sprint.html).