

FREE RESOURCE

The AI Readiness Checklist

Twenty questions to answer before you fund an AI initiative — and a score that tells you what to do next.

Drawn from C4 · Judging AI Output and A2 · Building AI-Powered Products

Most AI initiatives do not fail on the model. They fail because nobody could say what “good” meant, nobody owned the output, and the data turned out to be worse than everyone assumed. All three are knowable in an afternoon.

Score one point for each honest yes. Unsure is a no — that is the point of the exercise.

DATA

- 01 ☐ Can you say where the data for this feature comes from, and who owns it?
- 02 ☐ Do you have the rights to use it this way, including anything from customers or vendors?
- 03 ☐ Is the data good enough that a competent person could do the task well using only it?
- 04 ☐ Would you notice if the data silently changed shape or quality?

EVALUATION

- 05 ☐ Do you have a written definition of “good enough” for this feature?
- 06 ☐ Do you have a set of real examples with known-correct answers?
- 07 ☐ Can you tell whether a change made it better or worse, other than by opinion?
- 08 ☐ Do you know what the failure looks like to the user when it goes wrong?

PLATFORM

- 09 ☐ Do you know the cost per use, and what happens if usage multiplies by ten?
- 10 ☐ Do you know the latency, and whether the interface can hide it?
- 11 ☐ Could you switch model or vendor without rebuilding the feature?
- 12 ☐ Do you have monitoring that would catch quality degrading, not just errors?

PEOPLE

- 13 ☐ Is there a named person accountable for the quality of this feature's output?
-
- 14 ☐ Does someone on the team understand how the model behaves, not just how to call it?
-
- 15 ☐ Have the people reviewing the output been told what "good" looks like?
-
- 16 ☐ Would your team say what they actually think if this feature were bad?

POLICY

- 17 ☐ Is there a written rule about what may be sent to a third-party model?
-
- 18 ☐ Do you know which regulations apply to this use, if any?
-
- 19 ☐ Could you explain a particular decision to a customer who asked why?
-
- 20 ☐ Have you decided what you would do if it produced something harmful?

What your score means

16–20

Ready.

Your risk is scope, not capability. Ship something small, measure it against the definition you already have, and resist the urge to make the first one impressive.

11–15

Nearly.

Usually one whole area is empty rather than four being weak. Find which one, fix it, and re-score before you commit budget — it is normally two weeks of work.

6–10

Not yet.

Building anyway will produce a demo that never reaches production, which is the most expensive possible outcome because it costs a quarter and teaches you nothing. Two weeks on the gaps saves that quarter.

0–5

This is not an AI problem.

It is a data and accountability problem wearing an AI costume. AI is not the next thing to fix, and doing it first will make the underlying problem harder to see.

The four most commonly failed

Across the organisations I have worked in and taught, these four are missing far more often than the rest — and each one is cheap now and expensive later.

5 · A written definition of “good enough”

Without it, every review is an argument about taste, and the feature ships when someone gets tired rather than when it is ready.

6 · Real examples with known-correct answers

Twenty of them is enough to start. This is the single highest-value afternoon in the whole list, and almost nobody spends it.

11 · Being able to switch model or vendor

Not for negotiating leverage. For the day your model is deprecated with three months’ notice, which now happens routinely.

17 · A written rule about what may be sent

Your team is already making this decision individually, several times a day. The only question is whether they are making it consistently.

WHAT TO DO WITH A LOW SCORE

Not nothing. Pick the single lowest area, spend two weeks on it, and score again. This is a diagnostic, not a gate — its job is to tell you which two weeks to spend, not to give you permission.

The Bridge is a training academy for product, engineering and data leaders — workshops, programmes, mentoring and advisory on the decisions the three functions have to make together. Taught live by a practising CTO. Everything here is drawn from the catalogue and published openly. thebridge.academy