

ONE PAGE · FIVE QUESTIONS

The AI Readiness Lens

Five questions to put to any process before it gets called an AI opportunity.

Most stalled AI pilots were never intelligence problems. They were data, integration or process problems with an AI budget attached. Run these five questions across whatever is currently on your candidate list. Fair warning: most candidates fail question one. That is the point of question one.

Q1 Is intelligence actually the bottleneck?

A GOOD ANSWER Skilled people apply judgement, repeatedly, to unstructured input. The exception rate is high and mostly undocumented.

A BAD ANSWER The SOP covers nearly everything. What actually delays the work is a broken handover, an integration nobody owns, or a decision nobody wants to sign.

Q2 Is the data available?

A GOOD ANSWER It can be reached in a usable form inside the assessment window, and one named person can authorise access to a sanitised sample.

A BAD ANSWER It exists somewhere. Getting to it needs legal, security and a data owner, and nobody has started that conversation.

PERCEPTION TASKS Labelled examples of what goes wrong are rare. Check they exist before pricing a build.

Q3 What needs deterministic software instead?

A GOOD ANSWER The rules and the judgement have been separated, and only the judgement is being priced as AI.

A BAD ANSWER The whole process is in scope, including the part the ERP already does and the part a rules engine would do more cheaply and more predictably.

Q4 How hard is the integration, honestly?

A GOOD ANSWER You know which system consumes the output, who owns it, and roughly what the change costs.

A BAD ANSWER The output lands in a spreadsheet and someone re-keys it into the system of record. Net saving: zero.

Q5 Will the people doing the work trust it enough to use it?

A GOOD ANSWER They can see why it reached its answer in language they already use, override it without raising a ticket, and they know who carries a wrong output.

A BAD ANSWER Adoption is listed as a change-management workstream starting after build. You end up with an accurate system at 12% utilisation.

Scoring

Five yeses is rare. Four is worth proving with a prototype. Three or fewer and it is not an AI opportunity yet — which is a seven-figure initiative you did not start.

One nuance worth holding on to: a no on Q2 or Q4 is usually a timing answer. A no on Q1 is a permanent one.

WORKSHEET

Score your own candidates

Start with the list you already have — the business case, the pilot, the thing someone keeps mentioning. Then add the candidates that were never on it. The best ones rarely are: they are too small and too boring to justify a slide, which is also why nobody has costed them.

THE QUESTION THAT FINDS THEM

What does your organisation do that is easy for your people, impossible to write down, and happening several hundred times a day?

CANDIDATE PROCESS	Q1	Q2	Q3	Q4	Q5	SCORE
ON THE CURRENT AI LIST						
.....	☐	☐	☐	☐	☐	_____
.....	☐	☐	☐	☐	☐	_____
.....	☐	☐	☐	☐	☐	_____
NOT ON ANY LIST — the two-second, several-hundred-times-a-day work						
.....	☐	☐	☐	☐	☐	_____
.....	☐	☐	☐	☐	☐	_____
.....	☐	☐	☐	☐	☐	_____

Record the reason for every no

This is the half that survives the meeting. A candidate ruled out with a written reason is a decision; a candidate ruled out in conversation comes back next quarter.

.....

.....

If two of them look real, the next step is half an hour.

Bring three or four candidates and we will tell you honestly which look real and which look like plumbing. The call is free and it is not a pitch.

opta6.com/ai-opportunity-sprint