AI Readiness Assessment

Scoring data quality, process maturity and controls before choosing the first AI use case.

Nov 10, 2026 · @Vijay Srinivas Kothapalli

Purpose

Many businesses start with AI by buying a tool and hoping it will help. Results are usually disappointing, not because the technology failed, but because the data, the process or the people were not ready. This assessment shows where your business stands across five areas, so you choose a first use case that can succeed and build from there.

It does not test any particular AI product. It tests your readiness to use AI well.

When to use it

Use the assessment before:

- Approving an AI pilot or an AI feature purchase.
- Setting an AI budget for the coming year.
- Asking a supplier to add AI features to a system you already use.

Repeat it every 12 months, and whenever the data, the team or the processes change significantly.

The assessment at a glance

1. **Business purpose and use case.** Is there a clear problem worth solving?
2. **Data quality and availability.** Is the data accurate, complete and accessible?
3. **Process maturity and ownership.** Is the process stable and does someone own it?
4. **People, skills and governance.** Do staff have the skills, and are the rules agreed?
5. **Technology, security and suppliers.** Can the systems and suppliers be used safely?

Score each area from 1 (not in place) to 5 (well established). The total and the weakest areas decide the next step.

Area 1: Business purpose and use case

AI should answer a specific business question, not be a general aim.

- **Score 5:** The use case names the problem, the people affected, the measure of success and the current baseline. For example: "reduce the time to answer routine supplier queries from two days to four hours."
- **Score 3:** The problem is clear, but the measure and baseline are not yet agreed.
- **Score 1:** The aim is "use AI" or "be more innovative", with no problem named.

Ask: what decision or task will change, and how will we know it changed for the better?

Area 2: Data quality and availability

AI is only as good as the data it learns from or works on. Check the data for the first use case, not the whole business.

Check	Score 5 looks like	Score 1 looks like
Accuracy	Errors are measured and below an agreed threshold	Nobody knows how accurate the records are
Completeness	Required fields are filled in most of the time	Key fields are often blank
Access	Data can be extracted in a usable format, with permission	Data sits in spreadsheets, paper files or personal folders
Consistency	One definition of each field across systems	The same customer or product has several codes
Rights	It is clear what the data may be used for, including personal data	Nobody has checked whether the use is permitted

If the first use case depends on data that scores 1 on accuracy or rights, fix the data before building anything.

Area 3: Process maturity and ownership

AI works best on a process that is stable, documented and owned.

- **Score 5:** The process is written down, the steps are consistent, and a named manager owns the outcome and its measures.
- **Score 3:** The process is understood, but practice varies between teams or sites.
- **Score 1:** The process lives in one person's head, or nobody owns the outcome.

If the process itself is unclear, map it first. A short process map, as used in other process reviews, is often enough to show where AI could help and where it could not.

Area 4: People, skills and governance

The people who use AI, and the rules they follow, decide whether it is used safely and well.

- **Skills:** Do the users understand what the tool can and cannot do? Score 5 when staff are trained and know how to check outputs.
- **Culture:** Are staff encouraged to report errors and concerns, rather than hide them? Score 5 when they are.
- **Governance:** Is there an agreed AI policy, a named owner and a process for approving new uses? Score 5 when all three are in place.
- **Leadership:** Does the senior team understand the risks and the expected benefits? Score 5 when the budget and the measures are agreed.

If governance scores 1, adopt a basic AI policy before the pilot starts. A short checklist is enough to begin.

Area 5: Technology, security and suppliers

The question here is whether the tools can be used safely with your data.

- **Security:** Are access controls, multi-factor sign-in and logging in place for the systems AI will touch? Score 5 when they are.
- **Integration:** Can the AI tool connect to the systems it needs without copying data into unmanaged places? Score 5 when it can, through approved routes.
- **Supplier terms:** Do the terms say how inputs are used, where they are stored, and whether they train the supplier's models? Score 5 when the answers are written and acceptable.
- **Exit:** Could you stop using the tool, and remove your data, without losing the core system? Score 5 when you could.

A supplier that cannot answer these questions in writing should be scored 1 until it can.

Scoring and the decision

Add the five scores. The total ranges from 5 to 25.

Total	Readiness	Next step
20 to 25	Ready for a pilot	Choose a first use case from the list below and start a controlled pilot
14 to 19	Ready after targeted work	Fix the weakest one or two areas, then reassess before the pilot
5 to 13	Not ready	Build the foundations first: data, process ownership and a basic policy

As with other readiness reviews, any area that scores 1 blocks a pilot until it has a plan. A high total with weak data rights is not ready.

Choosing the first use case

The best first use case is useful, contained and easy to check. Score each candidate from 1 to 5 on these four points:

1. **Value.** Saves measurable time, cuts errors or improves a customer outcome.
2. **Data.** The data needed scores 4 or 5 on the data checks.
3. **Risk.** A wrong output is easy to spot and cheap to correct, and no decision about a person is made by AI alone.
4. **Repetition.** The task happens often enough to show results within a quarter.

Good first use cases include summarising long internal documents, drafting routine replies that a person reviews, and classifying incoming requests for routing. Avoid first use cases that decide credit, hiring, pricing for individuals or access to services.

Common pitfalls to avoid

- **Buying before defining the problem.** A tool with no agreed use case rarely delivers.
- **Assuming the data is ready.** Most projects spend more time on data than on the AI itself.
- **Measuring activity, not outcomes.** "Staff used the tool 500 times" is not a result.

- **Trusting outputs without checks.** Build the review step into the process from day one.
- **Letting one enthusiast run it.** Name an owner and a backup.
- **Ignoring staff concerns.** Explain what the tool is for, and what it will not be used for.

A 90-day path to a first pilot

Period	Action	Owner
Days 1 to 30	Score the five areas and identify the weakest.	Senior manager
Days 1 to 30	Shortlist three candidate use cases and score them.	Business owners
Days 31 to 60	Fix the data gaps for the chosen use case, and agree the data rights.	Operations or IT lead
Days 31 to 60	Adopt a basic AI policy and the staff rules.	Owner, with HR
Days 61 to 90	Run the pilot with a human review step and weekly measures.	Pilot owner
Days 61 to 90	Decide whether to scale, adjust or stop, against the agreed measure.	Board or owner

At day 90, the decision should rest on the measure agreed at the start, not on impressions.

About CRCO LLP

CRCO LLP is a Hyderabad-based advisory firm working with businesses on automation, process improvement, data protection and outsourcing. We help teams assess AI readiness, choose a first use case, fix the data and process foundations, and set up the controls for a safe pilot.

Book a consultation: visit the Contact page on our website, or write to us through the enquiry form.

This assessment provides general guidance only. It is not technical or legal advice. Obligations on personal data and automated decisions depend on your sector and circumstances.