

# The Enterprise AI Readiness Self-Check

*Twelve questions across the five dimensions that decide whether your organisation can move AI from pilot to production — strategy, data, talent, infrastructure, and governance.*

---

Most enterprise AI never ships. It demos well, gets a budget line, and then stalls somewhere between the pilot and the first real user. The pattern is consistent enough to predict: the organisations that cross into production share a named executive sponsor, a success metric tied to a business outcome, and at least one internal champion. The ones that don't, don't.

This is a self-check, not a maturity certificate. Answer each question honestly with **yes**, **no**, or **not sure** — a “not sure” counts as a no. Twelve questions, five dimensions. Score yourself at the end.

## Strategy and sponsorship

1. **Is there a named executive sponsor who owns the outcome — not just the budget?** A sponsor who can unblock a data-access fight or a procurement freeze is worth more than the model.
2. **Is the first use case tied to a metric someone already reports on?** “Reduce handle time,” “recover abandoned carts,” “cut Rx leakage” beat “explore AI.” If it doesn't move a number a director already watches, it won't survive the next planning cycle.
3. **Can you name the internal champion who will use the thing every day?** Pilots without a daily user are demos.

## Data

4. **Do you know where the data lives, who owns it, and whether you're allowed to use it for this?** Data readiness is the single most-cited cause of AI failure — usually access and lineage, not volume.
5. **Is the data good enough to be wrong about in production?** You don't need perfect data. You need to know its failure modes before a customer finds them.

## Talent and literacy

6. **Is there someone on the team who can tell a good output from a plausible-but-wrong one?** Evaluation is a skill, not a dashboard.
7. **Do the people around the system — support, ops, compliance — understand what it can and can't do?** Enterprise-wide AI literacy is what keeps a model from being quietly switched off after the first bad week.

## Infrastructure and integration

8. **Can the system reach the systems it needs to — read and write, by role?** Read access is a demo. Writing records, triggering workflows, and surfacing data by permission is production.
9. **Is there a path to monitor it in production — latency, cost, drift, and a way to roll back?** If you can't see it degrade, you can't keep it alive.

## Governance and the EU AI Act

10. **Have you classified this system under the EU AI Act?** The four risk tiers decide your obligations. High-risk provider duties bind from 2 August 2026 — this is not a 2027 problem.
11. **Is there human oversight built into the architecture, not bolted on after?** Someone must be able to understand, override, and detect anomalies — by design.
12. **Do you have one accountable owner for AI governance?** Not a committee that meets quarterly — a name.

## Read your score

Count your **yes** answers.

- **0 to 4 — Not yet.** You have an idea, not a plan. Start by naming a sponsor and a metric; everything else is downstream of those two.
- **5 to 8 — Pilot-ready, not production-ready.** You can run a credible pilot, but the gaps — usually data lineage, integration depth, or governance — are exactly where pilots die. Close them before you scale.
- **9 to 12 — Ship it.** You have the conditions for production. The remaining risk is execution discipline, not readiness.

Where you scored low is where an embedded product lead earns their keep — turning a stalled pilot into a measured outcome is most of what I do. If three of these questions made you wince, that's usually the conversation worth having.