

GUIDE

Enterprise AI Readiness Checklist

A checklist for deciding whether the whole organization is ready to fund AI beyond a pilot: data, ownership, process, sponsorship, operations and governance.

Version	1.0
Date	20 October 2026
Published by	XONIK
For	A CIO or executive deciding whether to fund AI work.
Use with	The Enterprise AI Readiness Assessment on the XONIK website, which scores the same six areas.

How to use this checklist

Work through the six areas with the people who own each one: someone from data or technology, someone from the business function that would use AI, and someone who answers for risk. Tick a box only when the statement is true today and you could show the evidence to a colleague who asked. A statement that is true in one team and not in another is not yet true.

Do not add the ticks into a single score. Readiness is uneven, and the useful result is the area with the fewest ticks, because that is where funded AI work will stall first. The last page lists the evidence to collect for each area, so the next review starts from facts and not from memory.

Data

AI work stalls when the data it needs is reachable only by a project of its own, or when two systems disagree about the same record and nobody owns the difference.

- ☐ The data a first use case needs can be reached by the people building it without a separate project.
- ☐ Each source has a named owner who can say what it contains and how current it is.
- ☐ Where two systems hold the same record, a rule says which one wins and the disagreements are logged.
- ☐ Someone has checked a sample of the data against the live source, not only against an extract.
- ☐ The data a use case needs may be used for that purpose, and the permission is written down.

If you cannot tick most of these

Choose one well-documented source for the first use case and make that one reachable. Do not start with a platform.

Ownership

A pilot has a team. A live system needs an owner, a person with the authority to change it or switch it off.

- ☐ Every AI system in use, or planned, has a named owner.
- ☐ The owner has agreed, and has the authority to change the system or stop it.
- ☐ It is clear who acts, and how, if a system gives a customer a wrong answer.
- ☐ A named role decides whether a use case moves from pilot to daily use, against criteria agreed in advance.
- ☐ The owner has time for the work, and it appears in their objectives.

If you cannot tick most of these

Name the owner before the pilot starts, and make the owner a condition of approval.

Process

A model placed in front of a process nobody has drawn tends to automate its confusion.

- ☐ The processes AI would change have been mapped as they actually run, with the people who run them.
- ☐ Each process names the point at which a person reviews or overrides the output.
- ☐ When a process changes, someone prepares the people who work in it, and the preparation is part of the plan.
- ☐ A day is set on which the old manual route stops, and it has been agreed.
- ☐ The exceptions that people handle by hand today have been listed.

If you cannot tick most of these

Map one process as it runs, not as designed, before choosing where AI fits.

Sponsorship

A sponsor who funds a demonstration is not a sponsor who funds a live system.

- ☐ An executive has agreed to fund AI beyond the pilot stage, and the conditions for that funding are written down.
- ☐ The sponsor knows what a first AI project may deliver and what it may not.
- ☐ If a first project falls short, the organization has agreed in advance how it will review it.
- ☐ The sponsor takes part in the decision to go live, not only the decision to start.
- ☐ Funding covers running and changing the system after launch, not only building it.

If you cannot tick most of these

Agree the conditions for production funding before the pilot starts, so that success has somewhere to go.

Operations

The first sign of a failing AI system is usually a customer.

- ☐ Someone is assigned to watch how each live system behaves, and knows what to watch.
- ☐ There is a tested fallback for when the system is down or wrong, and a named person to call.
- ☐ The organization would know if a model or a vendor behind a live system changed.
- ☐ Output is sampled and checked against agreed measures, on a schedule.
- ☐ Switching a system off is a decision someone is allowed to take, and it has been practiced.

If you cannot tick most of these

Decide who watches, and what they watch, before launch.

Governance

Rules written after a failure are written in a hurry.

- ☐ Written rules say what AI may and may not be used for, and who may approve a use.
- ☐ The people approving a use case can see how it was tested and what it was tested on.
- ☐ A list of AI systems in use, including features built into bought products, is kept and has an owner.
- ☐ A route exists for reporting a serious error, and it has been rehearsed.
- ☐ The rules are known by the people who must follow them, and someone checks that they do.

If you cannot tick most of these

Write a short rule that people follow, then add detail where it is needed.

Evidence to collect

For each area, note where the proof lives. If nobody can say, treat the statement as not yet true.

Area	The proof you would show	Where it lives, and who holds it
Data	A source list with owners, and a sample checked against the live source.	
Ownership	The owner's name, their agreement, and the decision criteria.	
Process	The process maps as run, and the date the old route ends.	
Sponsorship	The written conditions for production funding.	
Operations	The monitoring plan and the tested fallback.	
Governance	The written rules, the list of systems and the incident route.	

Limits of this document

This document is a working aid written by XONIK. It is general practitioner guidance and not legal, financial, technical-audit or regulatory advice, and it carries no certification. It reflects your own reading of your own organization, so it is only as accurate as the answers you give. Nothing in it guarantees a result.