

Five Questions to Ask Before Approving Any AI Initiative

A decision framework for non-technical managers · From the VERDICT methodology

Most AI decisions in organisations are made on instinct, vendor pressure, or avoidance. None of these hold up when leadership asks why. These five questions, drawn from the VERDICT Framework, give any manager a structured starting point for evaluating an AI proposal. They do not replace a full evaluation. They tell you whether one is needed.

01 What decision is this AI actually making, and who is currently accountable for it?

Before evaluating any AI tool, name the specific decision or output it will produce. Not "it will help with recruitment" but "it will shortlist candidates based on CV scoring." Then ask: who is accountable for that decision today? If the AI tool produces a wrong output, the accountability does not transfer to the software. It stays with the manager. If nobody can name who is accountable, the proposal is not ready to evaluate.

02 Is this process rule-based enough to automate, or more relational than it appears?

Some processes look systematic on the surface but depend on judgment, relationships, or contextual knowledge that cannot be codified. Take a matching process: pairing professionals to projects, for example. It may appear rule-based, but the best outcomes often depend on knowledge that lives in people's heads, not in databases. If the process depends on judgment that nobody can articulate as rules, AI augmentation may be appropriate. Full automation is not.

03 What happens when it gets it wrong?

Every AI system produces errors. The question is not whether it will be wrong, but what the consequences are when it is. A wrong summary of a meeting is low-stakes. A wrong clause extracted from a contract is not. Before approving any AI initiative, define the failure mode: what does a wrong output look like, who would notice, how quickly, and what is the cost of not catching it in time?

04 What data is this tool consuming, where does it go, and who else can see it?

AI tools connected to organisational systems (HRMS, procurement platforms, CRM) process data that may include employee information, commercial terms, or client details. Before any pilot: confirm what data the tool ingests, whether it leaves your environment, whether it is used to train the model, and whether your data protection obligations under GDPR, DPDP Act 2023, or contractual confidentiality are met. If the vendor cannot answer these questions clearly, that is your answer.

05 What would we need to see in a 30-day pilot before committing?

The strongest position is never a simple yes or no. It is "pilot, with conditions." Define in advance what success looks like in 30 days: which metrics, measured how, reviewed by whom. A vendor who resists a structured pilot, or who cannot define measurable success criteria, is selling confidence, not capability. Your organisation deserves the latter.

These five questions are drawn from the VERDICT Framework, a structured decision methodology for evaluating AI adoption, developed by The Work Architect. The full framework is delivered through a two-day capability programme designed for managers who are accountable for AI decisions in their organisations.

Start a conversation: anshuman@theworkarchitect.com · theworkarchitect.com