



NAKEDAI — PILLAR ARTICLE

# The AI Decision Stack

*Ten layers. Each one is the foundation for the one above it. Miss one and the whole thing is at risk.*

Most AI initiatives fail not because the technology was wrong, but because the decision to pursue them was never properly made. The AI Decision Stack is the framework we use at NakedAI to assess whether that decision is sound — across every dimension that matters — before any capital is committed.

There are ten layers. They are sequential for a reason: a strong answer at Layer 9 cannot compensate for a missing answer at Layer 4. The stack must hold at every level.

## Layer 1: Business Outcome

What commercial result must this AI initiative produce? This is the foundation. If the answer is vague — "improved efficiency", "better customer experience", "staying competitive" — the decision is not yet ready to be made. A technology choice made before the outcome is defined is not a decision; it is a purchase. Everything else in the stack depends on a clear, specific answer here.

## Layer 2: Decision Structure

Has this decision been formally documented, or does it exist only in conversations and assumptions? Decisions that cannot be written down have not yet been made. A documented decision creates accountability; an undocumented one creates ambiguity. At this layer we establish whether the decision has a formal structure — or whether we are still in the pre-decision phase.

### **Layer 3: Alternatives Assessment**

Have non-AI alternatives been genuinely considered and ruled out? AI is not always the right answer. Sometimes a process change, a hire, or an existing tool would produce the same result at a fraction of the cost and risk. Skipping this layer means the decision was made in favour of AI before the problem was fully understood.

### **Layer 4: Decision Ownership**

Is one named individual personally accountable for this decision and its outcome? Not a committee. Not a department. One person. Shared accountability is no accountability. When AI initiatives fail, the absence of a single accountable owner is almost always a contributing factor. This layer establishes who that person is — and confirms they have accepted the responsibility.

### **Layer 5: Risk Ownership**

Who owns the financial, legal, regulatory, and reputational risk if this fails? Decision ownership and risk ownership are not always the same person. A CFO may own the budget decision while the legal team owns the regulatory exposure. This layer maps where the risk sits — and whether the people who own it are aware of and have accepted that responsibility.

### **Layer 6: Downside Definition**

Have the specific consequences of failure been identified concretely? Not "the project might not deliver" — but what, specifically, would happen: financial exposure, regulatory action, reputational damage, operational disruption. Organisations that cannot answer this question clearly have not finished assessing the decision. Defining the downside is not pessimism; it is due diligence.

### **Layer 7: Organisational Readiness**

Does this organisation have the capability and capacity to deliver and manage what it is committing to? Many AI initiatives fail not because the technology was wrong but because the organisation was not ready to absorb it: the data was not clean, the teams were not trained, the processes were not adapted, and the leadership attention was not sustained. Readiness must be assessed honestly — not optimistically.

## **Layer 8: Deployment Pathway**

Is there a clear, governed route from pilot to operational deployment? A pilot that succeeds but cannot scale is not a success — it is a delayed failure. This layer assesses whether the path from proof of concept to live operation has been mapped: who governs the transition, what the success criteria are, and what happens at each stage if things do not go to plan.

## **Layer 9: Board Defensibility**

Can this decision be explained and justified to a board or investor without specialist knowledge? If the only people who can explain why this decision makes sense are the people who proposed it, the decision has not been made at the right level. Board defensibility means the logic is clear, the risks are understood, and a reasonable non-specialist could assess the decision on its merits.

## **Layer 10: Regulatory Resilience**

Would this decision withstand legal or regulatory scrutiny if challenged? The regulatory environment around AI is changing rapidly. What is permissible today may not be permissible in twelve months. This layer assesses whether the decision has been tested against current and anticipated regulatory requirements — and whether there is a governance structure in place to respond if those requirements change.

## **The stack is the standard**

These ten layers are not a checklist to be completed and filed. They are the minimum standard for a defensible AI decision. An organisation that can answer each question clearly, with evidence, is in a strong position to proceed. An organisation that cannot is not yet ready — regardless of how compelling the technology looks or how persuasive the vendor is.

The Human Pause is our structured assessment of all ten layers. It produces a clear score, a gap analysis, and a proceed / pause / redesign / stop recommendation — before any commitment is made.

---

## **Find out where your AI decision stands**

The Human Pause Score assesses your decision across all ten layers of the AI Decision Stack and gives you an immediate result — including your recommended next step.

**Take the Human Pause Score at [nakedai.io/human-pause-score](https://nakedai.io/human-pause-score)**