



NAKEDAI

The Human Pause.

A guide to strategic AI clarity
and governance for leaders.

FOR BOARDS, CFOs, AND
SENIOR LEADERSHIP TEAMS

nakedai.io



WHAT THIS IS

A short guide for considered AI decisions.

This booklet sets out who we are, what we do, and how we work. It is written for the moment between recognising that AI must be addressed and committing capital to a particular answer.

If you read it through, it will take fifteen minutes. If you only have five, the assessment on page fifteen is the most useful place to start.

- 03** The Problem
- 04** The Question
- 05** Who we are
- 06** Malcolm Maxwell
- 07** Marla Ubhi
- 08** Saakshi Saggu
- 09** In practice
- 10** The Human Pause
- 11** The Decision Stack
- 12** Four valid outcomes
- 13** The AI Decision Pack
- 14** What follows
- 15** TakeThe Human Pause Score



THE PROBLEM

Most AI failures do not start with model capability.

They start earlier. Unclear ownership. Weak business cases. Governance arriving as paperwork after the substantial decisions have already been made.

By the time a board sees a failure forming, the capital is committed, the vendor is engaged, and the damage is already taking shape.

Every week, leadership teams approve AI initiatives. Vendors are selected. Budgets are signed off. Projects are launched. Almost none of them have asked the right question first.



Not 'which AI tool should we buy?' **That question comes later, and only if the earlier one is yes.**

The earlier question is whether the decision itself is sound. Whether the organisation knows what it is trying to achieve, who owns the outcome, what happens if it fails, and whether it is genuinely ready to govern what it is about to commit to.

There is a moment, just before commitment, when every option remains open. Once that moment passes, the cost of correcting course rises sharply.

Three things tend to fail first, none of them about technology: ownership of the decision, clarity of the business case, and readiness of the organisation to govern what is being approved. The model can be excellent and the deployment still fail because of these.

**The Human Pause exists for
the moment before that.**



We work at the point where AI decisions are formed.

Our role is not to make work easier. It is to ensure decisions are better.

We have seen that efficiency improves output. It does not improve judgement. AI can accelerate activity without improving the quality of the decision behind it.

We use AI to introduce challenge, not comfort: to surface risk, to prompt better thinking, to test whether a decision holds. We do not treat AI as a replacement for judgement. We treat it as a tool that strengthens it.

We operate as a team of nine.

Only three are human. All accountability sits with us.



Malcolm Maxwell.

CTO / Chief Bot Whisperer

AI architecture and systems



Malcolm designs AI systems that remain accountable, secure, and defensible over time.

His work is grounded in the operational realities of running AI inside an organisation: where the data lives, how the model behaves under stress, what a human reviewer actually sees, and what evidence remains when something has to be explained to a regulator.

His instinct is to ask what the system will look like a year after launch, not just whether it works on day one. Too many AI deployments make it through go-live and then drift; his job is to design out that drift before it begins.

Considered, structural, sceptical of comfort answers.



Marla Ubhi.

CEO / Chief People Whisperer

Decision and operational leadership



Marla works at the point where decisions carry financial and operational consequence.

She translates between the language of the board and the language of the engineering team, ensuring that what gets approved at the top is what actually gets built, governed, and defended below.

In practice, much of her work is translation. The board speaks in commercial terms; the engineering team speaks in system terms; the gap between the two is where most AI initiatives lose their footing. She closes that gap until everyone is making the same decision.

Direct, commercial, allergic to ambiguity.



Saakshi Saggu.

Associate / Chief Knowledge Keeper

Knowledge and decision support



Saakshi ensures decision clarity is captured, structured, and retained across the work.

Her role is to make sure nothing important is held only in the memory of a single person. The decision record, the rationale, the dissent, the open question that was never fully answered: she keeps the trail intact.

In fast-moving engagements, the things that matter most are the things easiest to forget: why an approach was rejected, who raised an objection later proved right, what evidence was relied on. Her work is to ensure that record is intact when the organisation needs it.

Methodical, calm, attentive to what others miss.



PROOF

In practice.

“What Malcolm delivered was twofold: he validated the core approach we'd taken, allowing us to rule out architectural paths we'd been considering, and identified targeted improvements with clear, practical guidance. The fees paid for themselves within a week. They were a pleasure to work with: thoughtful, engaged, and genuinely collaborative.”

Jon Elek

Founder and CEO, Althea Research Ltd



THE HUMAN PAUSE

The diagnostic.

The Human Pause is the first thing we do, and for many decisions it is also the only thing required.

Few leadership teams can clear a week of senior calendar at once, so we do not ask them to. Most engagements run as focused sessions, an hour or two at a time, paced around your team's availability over two to four weeks. Where a decision is urgent and the team can clear a day, the same work compresses into a single sitting. The substantive analysis sits with us between sessions, so the time you commit is the time that counts.

Preparation is minimal and targeted. The work is grounded in how your organisation actually operates, not in generic use cases or theoretical opportunity.

WHAT WE TEST

Whether the decision is clearly defined, owned, governed, and defensible against the ten layers of the AI Decision Stack.



METHODOLOGY

The Decision Stack.

Ten layers. Each is the foundation for the one above it. The stack must hold at every level

10	Regulatory Resilience
09	Board Defensibility
08	Deployment Pathway
07	Organisational Readiness
06	Downside Definition
05	Risk Ownership
04	Decision Ownership
03	Alternatives Assessment
02	Decision Structure
01	Business Outcome

A strong answer at Layer 9 cannot compensate for a missing answer at Layer 4.



WHERE IT LEADS

Four valid outcomes.

The Pause does not exist to produce a yes. It exists to produce the right answer, which sometimes is yes, and sometimes is something else. A meaningful proportion of the work we do ends in something other than Proceed. That is the value.

Proceed.

The decision is clear, ownership is in place, the business case is sound, governance is real. Move.

Pause.

The decision could be sound, but specific layers are not yet ready. Address those, then revisit.

Redesign.

The underlying problem is real, but the AI approach being considered is wrong for it. Reframe.

Stop.

AI is not the right answer to the problem. The expensive answer is to commit anyway. The cheaper answer is to stop now.



THE DELIVERABLE

The AI Decision Pack.

A board-grade written record built from eight artefacts. Designed to give whoever holds the budget the evidence they need to make a defensible call.

01	AI Decision Clarity Score	02	Decision-Owner Map
03	Business-Outcome Map	04	AI Spend Risk Register
05	Governance & Accountability Gap Analysis	06	Vendor & Readiness Challenge Sheet
07	Proceed / Pause / Redesign / Stop Recommendation	08	90-day Decision Roadmap



WHAT FOLLOWS

Where the diagnostic leads to action.

Where the recommendation is to proceed, two further engagements may follow. Each is scoped on its own merits.

STAGE 02

AI Decision Architecture Sprint

Translates the agreed decision into system design and the governance architecture required to deliver it. Boundaries, controls, accountability.

STAGE 03

AI Oversight / Decision Board

Decision rights, escalation pathways, ongoing oversight. Operates as second opinion or retained involvement, ensuring decisions remain valid as systems evolve.

OR IT ENDS HERE

No further engagement needed.

Where the recommendation is pause, redesign, or stop, the Decision Pack is the deliverable. You leave with a board-grade record and the evidence to defend the choice.



THE DIAGNOSTIC

Take the Human Pause Score.

A structured ten-minute diagnostic. It assesses your decision across all ten layers of the AI Decision Stack and returns an immediate, specific result with your recommended next step.



Or visit:

nakedai.io

[/human-pause-score](https://nakedai.io/human-pause-score)



**A decision that
cannot be explained,
owned, and defended,
is not ready
to proceed.**

CONTACT

Thomas Ford House, 23–24 Smithfield Street, London EC1A 9LF
+44 (0) 7769 530 558 · hello@nakedai.io · nakedai.io

LONDON · EDINBURGH · TORONTO · AUCKLAND