



The Human Pause: **AI** Decisions Resource Toolkit

FOR BOARDS, CFOs, AND
SENIOR LEADERSHIP TEAMS

nakedai.io



THE RESOURCES

Introduction to the AI Decisions Resource Toolkit

Welcome to the AI Decisions Resource Toolkit, designed to support our Human Pause framework. This toolkit offers a set of questions and insights that serve as a starting point for our discussions around AI in the business.

These resources help us consider ownership, functionality, and impact before integrating AI solutions. By addressing these vital topics, we aim to facilitate informed decision-making that aligns with our organisational goals and values.

Use this toolkit as a practical guide to navigate the complexities of AI together, ensuring our decisions remain transparent and collaborative.

- 03** The Decision Stack
- 05** The AI Integrations Matrix
- 09** Who Owns The AI Outcome
- 11** Where Does AI Actually Belong
- 13** The AI Spend Sanity Check
- 15** Define The Pause Trigger Before You Start
- 17** The Seven Question AI Decision Framework
- 19** TakeThe Human Pause Score

The Decision Stack

Most AI initiatives fail not because the technology is wrong, but because the decision to pursue them was never properly made.

The Decision Stack is a ten-layer framework for assessing AI decisions before capital, reputation, and execution capacity are committed. Each layer is a question that must be answered clearly before the one above it can be trusted. Gaps at the foundation compound upwards - which is why the organisations that struggle most with AI are rarely those that chose the wrong tool. They are those that approved the wrong decision.

VISUAL OVERVIEW

10

Regulatory Resilience

Would this decision withstand legal or regulatory scrutiny if challenged?

9

Board Defensibility

Can this be explained and justified to a board or investor without specialist knowledge?

8

Deployment Pathway

Is there a clear, governed route from pilot to operational deployment?

7

Organisational Readiness

Does the organisation have the capability and capacity to deliver and manage this?

6

Downside Definition

Have the specific consequences of failure been identified concretely?

5

Risk Ownership

Who owns the financial, legal, regulatory, and reputational risk if this fails?

4

Decision Ownership

Is one named individual personally accountable for this decision and its outcome?

3

Alternatives Assessment

Have non-AI alternatives been considered and ruled out?

2

Decision Structure

Is this decision formally documented, or does it exist only in conversation?

1

Business Outcome

What commercial result must this AI initiative produce?

FOUNDATION

Gaps at the foundation compound upwards. The stack must hold at every layer.

The Decision Stack

1

Layer 1: Business Outcome

What commercial result does this AI initiative need to produce? If this cannot be stated clearly, nothing else can be decided. A technology choice made before the outcome is defined is not a decision - it is a purchase.

2

Layer 2: Decision Structure

Has the decision been formally documented, or does it exist only in conversations and assumptions? A decision that cannot be written down has not yet been made.

3

Layer 3: Alternatives Assessment

Has AI been assumed as the answer, or have non-AI alternatives been considered and ruled out? Choosing AI without testing alternatives is not a decision - it is a preference dressed as a strategy.

4

Layer 4: Decision Ownership

Is there one named individual accountable for this decision and its outcome? Shared ownership is no ownership. If three people own a decision, no one does - and when something goes wrong, the post-mortem produces a committee rather than a lesson.

5

Layer 5: Risk Ownership

Who owns the financial, legal, regulatory, operational, and reputational risk if this goes wrong? Risk ownership and decision ownership are not the same thing - and confusing them is one of the most common and expensive governance failures in AI investment.

6

Layer 6: Downside Definition

Have the specific consequences of failure been identified - not in general terms, but concretely? "This could go wrong" is not a risk. A named consequence, with an owner and a cost attached to it, is.

7

Layer 7: Organisational Readiness

Does the organisation have the capability, data, leadership, and governance capacity to deliver and manage this? Most AI failures are not technology failures. They are readiness failures that were visible before the project started.

8

Layer 8: Deployment Pathway

Is there a clear, governed route from pilot to operational deployment, or is the pilot isolated? A pilot that cannot explain how it becomes an operational system is an experiment, not a decision. Experiments are valuable; they should not be funded as deployments.

9

Layer 9: Board Defensibility

Can this decision be explained and justified clearly to a board or investor without specialist knowledge? If the person accountable for the decision cannot explain it plainly, the decision is not yet clear enough to act on.

10

Layer 10: Regulatory Resilience

Could this decision withstand legal or regulatory scrutiny if challenged? This is not a question about compliance documents. It is a question about whether the decision, the evidence, and the accountability chain would hold up under examination.

The ten layers are sequential for a reason. A strong answer at Layer 9 cannot compensate for a missing answer at Layer 4. Gaps at the foundation do not stay contained - they surface later, at the point where they are most expensive to fix.

The Human Pause is a structured assessment of all ten layers before commitment is made. It is not a slow-down. It is the work that makes everything that follows faster, cleaner, and more defensible.

The AI Integrations Matrix.

A board-level map of where AI sits, what it touches, who owns it, and how much governance it needs. Use it to calibrate the size of the pause to the size of the decision.

THE MATRIX



WHAT THE MATRIX HELPS NAKEDAI DO

- 01 Locate the AI activity
- 02 Identify ownership
- 03 Judge proportional risk
- 04 Trigger The Human Pause

Where does your organisation sit on this matrix?

Take the Human Pause Score to find out where AI is, where it should be and what to do next



The AI Integrations Matrix

Map the decision before you size the governance

USE THIS WHEN

Use this framework when an AI proposal is on the table and the discussion has moved to platforms, suppliers or pilots before anyone has established what kind of decision this actually is.

WHY IT MATTERS

The same technology can be a minor convenience for one person or a system the organisation depends on. The interface may look identical in both cases. The decision is not the same, and neither is the governance it needs.

Two questions place any AI use on the matrix. How much does it influence a material decision, and how deeply is it integrated into the way the organisation works?

The Four Positions

1. Personal productivity

Low impact, ad hoc integration.

An individual uses AI to work faster and think better: drafting, summarising, research prompts. The output stays with them, and they are close enough to it to see when it is wrong.

Ownership sits with the individual, with manager guidance. Governance should be proportionate: acceptable-use guidance and data awareness.

2. Team workflow

High impact, ad hoc integration.

A department uses AI in repeatable workflows: legal review, HR screening, marketing production. The output no longer stays with one person. It travels into other people's work and informs what they do next.

Ownership moves to a function lead. This position needs policy, review and consistency, because one person's unchecked output becomes another person's starting assumption.

3. Operational process

Low impact, embedded integration.

AI becomes part of how the business runs: customer service, finance checks, procurement analysis. Each individual output may be minor. The volume is not.

Ownership sits with operations or the process owner. This position needs controls, an audit trail and escalation, because a small error repeats at scale before anyone sees the pattern.

4. Decision support

High impact, embedded integration.

AI influences material decisions: hiring, lending, pricing, mergers and acquisitions, compliance. These decisions have real consequences for real people.

Ownership requires an executive sponsor and board accountability. This position needs human oversight, explainability and defensibility. Oversight must be real, which means the reviewer needs time, context and the authority to reject the output.



The AI Integrations Matrix

Map the decision before you size the governance

Map two positions, not one

Most organisations map where a pilot sits today and stop there. The more useful question is where success will take it.

A pilot may use copied files, a small group of users and manual review. The intended operating model may connect to live data, serve several departments and shape decisions that matter. Treating both as the same use creates a false sense of continuity. Technical capability often scales smoothly. Ownership, review capacity and risk do not.

Write down two positions before a pilot begins. Where the AI sits during the pilot, and where it will sit if the pilot succeeds. Then describe what has to be true to move between them.

That movement is the decision.

AI you did not decide to buy

Most organisations map where a pilot sits today and stop there. The more useful question is where success will take it.

A pilot may use copied files, a small group of users and manual review. The intended operating model may connect to live data, serve several departments and shape decisions that matter. Treating both as the same use creates a false sense of continuity. Technical capability often scales smoothly. Ownership, review capacity and risk do not.

Write down two positions before a pilot begins. Where the AI sits during the pilot, and where it will sit if the pilot succeeds. Then describe what has to be true to move between them.

That movement is the decision.



The AI Integrations Matrix

Map the decision before you size the governance

The leadership check

Use these questions with your leadership team.

If an answer is vague, leave it blank and treat that as useful evidence.

CURRENT POSITION

Where does this AI use sit on the Matrix today?

INTENDED POSITION

Where will it sit if it succeeds?

CONDITIONS TO MOVE

What must be true before it moves from one position to the other?

OWNERSHIP

Who owns it now, and who owns it then?

EMBEDDED AI

What AI is already running inside software we pay for, and where does that sit?

Size the pause to the size of the decision.

A personal drafting tool does not need a board process. A system influencing material decisions does not become safe because it has the same chat box as the drafting tool.

The purpose is not to make every AI use look large. It is to stop a large use being governed as though it were small.

If you cannot agree where this sits, or where it is heading, take a Human Pause before approval.

Where does your organisation sit on this matrix?

Take the Human Pause Score to find out where AI is, where it should be and what to do next.



Who Owns The AI Outcome

A HUMAN PAUSE ownership check before AI spend becomes operational risk.

USE THIS WHEN

Use this before approving an AI pilot, platform, workflow change, or automation project where accountability is still vague.

WHY IT MATTERS

AI projects often have a budget owner, a delivery team, and a technical lead. That is not the same as having one named business owner who carries the outcome. If ownership is unclear at approval, it usually becomes even less clear when cost, adoption, trust, or delivery problems appear.

The Ownership Question

- If this project underperforms, disrupts work, costs more than expected, or damages trust, which leader owns the consequence?
- A useful AI decision names a person who owns a business outcome, not just the person managing the project.

3 Ownership Decisions to Make

- Name the outcome owner: who is accountable for the business result this AI project is meant to create?
- Name the adoption owner: who is responsible for whether people affected can actually use, trust, and work with the change?
- Name the pause authority who has permission to reduce scope, redesign, or stop the project if the evidence is not holding?

What good looks like

- The outcome is specific enough to explain in plain English.
- The owner has enough authority to change the plan, not just report on it.
- The leadership team agrees what evidence would prove the project is working.
- The decision can be defended financially, operationally, and reputationally.



Who Owns The AI Outcome

A HUMAN PAUSE ownership check before AI spend becomes operational risk.

Working Page

Use this page with your leadership team. If the answer is vague, leave it blank and treat that as useful evidence.

CHECK	QUESTION	ANSWER / EVIDENCE
BUSINESS OUTCOME	What measurable result is this AI project expected to improve?	
NAMED OWNER	Who owns the result if the project underperforms?	
PEOPLE/ ADOPTION IMPACT	Who will use, check, trust, or be affected by the change?	
PAUSE AUTHORITY	Who can reduce scope, redesign, or stop the project?	
EVIDENCE TO REVIEW	What will we measure, and when will we review it?	

If you cannot name the owner, the decision is not ready. Take a Human Pause before approval.

RESOURCE 2



Where Does AI Actually Belong

Four work decisions before you automate, assist, augment, or keep human

USE THIS WHEN

Use this when a team is asking where AI could be added, but the work has not yet been mapped task by task.

WHY IT MATTERS

The strongest AI strategy is not to put “AI everywhere”. It is knowing where AI belongs, where it does not, and watching you stay human because the organisation still carries the accountability.

The Four Work Decisions

- **Automate:** repetitive rules-based checkable work where the output can be verified
- **Assist:** cognitive load work, such as summarising, searching, preparing, organising, or comparing information
- **Augment:** skilled judgement where AI can structure the evidence, but a person still decides and owns the outcome.
- **Keep human:** the work involving trust, ethics, relationships, exceptions, reputation, well-being, or final accountability.

The better question

- **Do not start with:** Where can we put AI?
- **Start with:** Where does AI actually belong in this work?
- A job rarely changes in one neat block. AI usually changes tasks, decisions, handoffs, evidence, and accountability.

When to be careful

- Do not add AI to a process simply because AI is available.
- Do not replace a working process if AI would add complexity without improving the outcome.
- Do not treat human judgement as an inefficiency when it is actually the control the business needs.



Where Does AI Actually Belong

Four work decisions before you automate, assist, augment, or keep human

Working Page

Use this page with your leadership team. If the answer is vague, leave it blank and treat that as useful evidence.

CHECK	QUESTION	ANSWER / EVIDENCE
TASK OR DECISION	What specific part of the work are we looking at?	
CURRENT PROBLEM	What is slow, risky, costly, inconsistent, or painful today?	
AI ROLE	Automate, Assist, Augment, or Keep Human?	
HUMAN ACCOUNTABILITY	Who checks, decides, explains, or owns the outcome?	
DO NOT CHANGE IF	What would make us leave this process alone for now?	
OWNERSHIP	Who owns outcome, adoption, cost, and risk?	

The team cannot map the work cleanly. The project is not ready for approval. Take a Human Pause first.



RESOURCE 3

The AI Spend Sanity Check

Four questions before AI spend becomes speculative

USE THIS WHEN

Use this before signing off spend where the idea looks impressive, but the problem, workflow, data, or value is still not clear enough.

WHY IT MATTERS

AI spend becomes speculative when it solves a problem the business has not properly named. A strong business case should survive plain English questions from the CEO, COO, and CFO before money, people, or operational change are committed.

Question one: What problem are we solving?

- Avoid vague goals such as "improve efficiency" or "increase productivity." Name the actual business problem.
- Ask whether the problem is expensive, frequent, risky, or strategically important enough to justify an AI intervention.

Question 2: Have we fixed the workflow first?

- AI will not rescue a process that no one understands.
- If the workflow is inconsistent, unclear, or full of avoidable handoffs, simplify that first.

Question 3: Is the data ready for this use?

- Check whether the data is accurate, accessible, permissioned, and appropriate for the decision being made.
- Bad data does not become safe because a new system is placed on top of it.

Question four: What value will this create?

- Name the financial, operational, customer, or risk outcome you expect to improve.
- If the value cannot be measured or described clearly, approval should pause until the business case is stronger.



RESOURCE 3

The AI Spend Sanity Check

Four questions before AI spend becomes speculative

Working Page

Use this page with your leadership team. If the answer is vague, leave it blank and treat that as useful evidence.

CHECK	QUESTION	ANSWER / EVIDENCE
PROBLEM STATEMENT	What specific business problem are we solving?	
WORKFLOW STATUS	Is the process clear enough to improve, or are we automating confusion?	
EXPECTED VALUE	What value will be created, and how will we know?	
SIMPLER ALTERNATIVE	Is there a simpler process or decision fix we should try first?	

If you cannot name the owner, the decision is not ready. Take a Human Pause before approval.



Define The Pause Trigger Before You Start

A simple way to stop AI projects drifting into cost, complexity, and low adoption

USE THIS WHEN

Use this before launching a pilot or programme where everyone is enthusiastic, but no one has defined what evidence would make the team stop, reduce scope, or redesign.

WHY IT MATTERS

Broad AI programmes can drift because the first decision is exciting and the later evidence is messy. A pause trigger protects the organisation by a green, before launch, what would count as a warning sign.

Set triggers before the project starts

- **Adoption:** if the people affected are not using or trusting the change by a defined point.
- **Cost:** if implementation, training, change, or support costs materially change the business case.
- **Quality:** if outputs require too much rework, checking, or exception handling.
- **Workflow:** if the project adds complexity rather than reducing friction.
- **Trust or reputation:** If customers, commerce, staff, or leaders cannot explain or defend the use of AI

What is a pause trigger?

- A pause trigger is the evidence that tells the leadership team to stop, reduce scope, redesign, or choose a smaller intervention.
- It is not a failure label. **It is a decision control.**

The discipline

- Start narrow. Solve one name problem well.
- Measure it honestly.
- Decide in advance what evidence would make you pause rather than drift.



Define The Pause Trigger Before You Start

A simple way to stop AI projects drifting into cost, complexity, and low adoption

Working Page

Use this page with your leadership team. If the answer is vague, leave it blank and treat that as useful evidence.

CHECK	QUESTION	ANSWER / EVIDENCE
PROJECT SCOPE	What problem is this AI project meant to solve first?	
SUCCESS EVIDENCE	Would prove the project is working?	
PAUSE TRIGGER	What evidence would make us pause, reduce scope, or redesign?	
REVIEW DATE	When will the leadership team review the evidence?	
DECISION OWNER	Who has authority to make the pause or redesign decision?	
OWNERSHIP	Who owns outcome, adoption, cost, and risk	

If there is no pause trigger, there is no real control. Take a Human Pause before the project becomes hard to unwind.



The Seven Question AI Decision Framework

A practical HUMAN PAUSE checklist before approving an AI project

USE THIS WHEN

Use this in the leadership meeting before approving AI spend, pilots, workflow changes or new operational use of AI.

WHY IT MATTERS

A confident idea is not the same as a decision-ready project. These seven questions help leadership teams test whether the project is clear, owned, proportionate, and defensible before commitment.

The Seven Questions

- **Problem:** what specific business problem are we trying to solve?
- **Fit:** does AI actually belong in this part of the work?
- **Work type:** should this be automated, assisted, augmented or kept human?
- **Simplicity:** is there a simpler process, workflow or decision fix before adding AI?
- **Value:** will this create measurable value, or just add complexity?
- **Ownership:** who owns the outcome, adoption, cost and risk?
- **Pause trigger:** what evidence would make us stop, reduce scope or choose a smaller intervention?

How to use it

- Put the seven questions on the agenda before approval, not after launch.
- Ask the CEO, COO and CFO to answer from their own responsibilities: strategy, operations and financial discipline.
- If the answers are vague, pause. The work is not to block AI; it is to make the decision stronger.

The Decision Rule

- If you can answer the seven questions clearly, the project may be ready to move forward with proper ownership and review.
- If you cannot, the next step is not more enthusiasm. The next step is clarity.



The Seven Question AI Decision Framework

A practical HUMAN PAUSE checklist before approving an AI project

Working Page

Use this page with your leadership team. If the answer is vague, leave it blank and treat that as useful evidence.

CHECK	QUESTION	ANSWER / EVIDENCE
PROBLEM	What specific problem are we solving?	
FIT	Why does AI belong here?	
WORK TYPE	Automate, assist, augment, or keep human?	
SIMPLICITY	What simpler fix have we considered	
VALUE	What measurable value will this create?	
OWNERSHIP	Who owns outcome, adoption, cost, and risk	
PAUSE TRIGGER	What evidence would make us stop, reduce, scope, or redesign?	

If your leadership team cannot answer all seven, take a Human Pause before approval.



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