

A LEGAL-SECTOR BRIEFING



NAKEDAI

# AI IN THE BACKGROUND

Liability, governance, litigation and  
the hidden risk of unowned AI  
decisions.

[nakedai.io](https://nakedai.io)



## WHAT THIS IS

# A legal-sector briefing for serious AI decisions

## For boards, general counsel, law firms and senior advisers

AI is no longer only something a person opens, prompts and checks. It is moving into workflows, board papers, diligence, litigation, compliance, dispute resolution, communications, cyber response, client advice and professional judgement.

That changes the risk. The question is no longer simply whether the tool works. The question is whether the organisation can explain what AI was allowed to touch, who owned the decision, what safeguards existed, what human judgement remained, and whether the outcome could be defended later.

This briefing is written for the moment before AI becomes embedded, relied upon or difficult to reverse.

**AI is becoming embedded in workflows faster than organisations can explain, govern or defend the decisions around it.**



## INSIDE

# The legal questions behind the AI question

- 04** What lawyers are already worried about
- 05** The litigation record starts before the litigation
- 06** Evidence, fabrication and trust
- 07** Reliance, negligence and professional judgement
- 08** Client pressure, in-house teams and legal economics
- 09** Authority, Agentic AI, liability and regulatory exposure
- 10** Disputes, arbitration, cyber and crisis
- 11** Due diligence and transactions
- 12** The AI Integrations Matrix
- 13** The Human Pause
- 14** The Human Pause - The Question
- 15** What the Human Pause is.
- 16** How we work
- 17** Methodology and The Decision Stack
- 18** Where NakedAI fits
- 19** The next step



## THE PROBLEM

# What lawyers are already worried about

The legal market is not only asking whether AI can make work faster. It is asking what happens when AI begins to affect evidence, reliance, professional duty, client expectations, public scrutiny and the economics of legal work.

- **Evidence:** Can we prove what is real, synthetic, altered or AI-generated?
- **Reliance:** When is it reasonable to rely on AI, and what judgement must still be exercised?
- **Professional duty:** Does AI change the standard expected of lawyers, advisers, directors or experts?
- **Privilege and confidentiality:** What happens when prompts, outputs or workflows expose material that should have remained protected?
- **Client pressure:** How do firms deliver efficiency without weakening quality, supervision or accountability?
- **In-house accountability:** Are general counsel being asked to manage AI risk after the business has already moved?
- **When AI starts to act:** Who authorised the AI to act, trigger, escalate or decide?
- **Class actions and liability:** Will weak AI governance become evidence in future claims?
- **Disputes and crisis:** What happens when AI shapes evidence, procedure, issue selection or settlement dynamics?
- **Cyber and fraud:** How does the organisation respond when AI is used against it?



## THE LITIGATION RECORD

# The litigation record starts before the litigation

By the time an AI dispute reaches lawyers, the record may already exist.

- **Board papers.**
- **Procurement notes.**
- **Vendor assessments.**
- **Risk registers.**
- **Testing records.**
- **Human review processes.**
- **Matter protocols.**
- **Privilege decisions.**
- **Escalation choices.**
- **Warnings unheeded.**

**The legal issue may be the output. The real issue is often the decision that allowed the output to matter.**



## EVIDENCE

# When evidence becomes uncertain

AI changes the evidential landscape. Deepfakes, synthetic documents, AI-generated communications, manipulated images, altered audio, automated summaries and hallucinated references all create a harder question for legal teams: can this be trusted?

That question is not only technical. It is legal, operational and reputational. Legal teams will need to ask how material was created, preserved, verified, challenged and relied upon.

As AI-generated material becomes easier to create, organisations will need more than detection tools. They will need decision discipline around evidence, verification and reliance.

**The Legal question:** not simply whether evidence exists, but whether its integrity can be shown, explained and defended.



## RELIANCE

# The professional risk is not simply that AI may be wrong

It is that a professional may rely on AI without being able to explain why that reliance was reasonable.

The risk may also cut the other way. As AI becomes part of ordinary practice, a professional may be challenged for failing to use it where its use would have been reasonable, efficient or expected.

The defensible position is neither "use AI" nor "do not use AI". The defensible position is clarity: being able to explain and defend the decision either way.

### **A defensible AI decision can show:**

- What decision was being made.
- What AI was allowed to influence.
- What limits were understood.
- What human judgement was applied.
- What evidence was kept.
- Why the choice was reasonable.



## CLIENT PRESSURE AND PROFESSIONAL DUTY

# Efficiency is not the same as judgement

Law firms are being pulled in two directions. Clients expect AI to make work faster, cheaper and more efficient. At the same time, they expect professional standards not to fall.

In-house legal teams sit where the pressure gathers. The business wants speed. The board wants assurance. Vendors promise transformation. Regulators expect oversight. The organisation wants lower cost.

Legal is brought in too late if it is only asked to approve what the business has already decided.. AI governance should begin when the organisation first asks: “what decision are we making?”

**The question for firms is not only how AI changes delivery. It is how they preserve trust while delivery changes.**



## AUTHORITY AND AGENTIC AI

# When AI starts to act

Agentic AI changes the question again. A system that produces an answer is one thing. A system that takes steps, triggers workflows, escalates issues, instructs tools, or influences connected systems is another.

The governance question is no longer only: “who checks the output?” It becomes: “who authorised the action?” AI-related claims may arise from:

- defective systems
- discriminatory outputs
- misleading assurance
- poor human review
- weak documentation
- over-reliance
- bad procurement
- failure to comply with emerging regulation

**The question is not only what the AI did. It is who allowed it to matter.**



## DISPUTES, CRISIS AND PUBLIC SCRUTINY

# AI is changing both the dispute and the process

It is not only changing the subject matter of disputes. It is changing dispute resolution itself: case management, document review, issue identification, chronologies, evidence analysis, drafting, settlement dynamics, and procedural efficiency.

Used well, AI may make disputes faster and more focused. Used carelessly, it may create new arguments about fairness, evidence, transparency, confidentiality and reliance.

AI governance is also about how the organisation responds when AI is used against it: deepfake fraud, synthetic identity, AI-assisted phishing, malicious GPTs, ransomware, false evidence and impersonation.

AI failures rarely remain purely technical. A disputed output, biased decision, manipulated image, unexplained system or poorly governed workflow can move quickly from internal issue to public narrative.



## DUE DILIGENCE AND TRANSACTIONS

# Speed is not certainty

In transactions, AI can accelerate red-flag review and issue identification. That does not make the output certain.

If AI-supported diligence affects valuation, risk allocation, warranty negotiation, disclosure, deal confidence or board approval, the decision trail matters.

AI may support diligence. It should not silently replace the judgement that gives diligence its value.

### Questions that matter:

- **What was AI asked to do?**
- **What data did it review?**
- **What did it miss or overstate?**
- **Who validated the output?**
- **What judgement remained with the lawyer?**
- **What was escalated, ignored or assumed?**



## THE AI INTEGRATIONS MATRIX

# Where AI sits

### What it influences, and whether the decision can be defended.

The Matrix does not ask whether AI is good or bad. It asks where AI has entered the organisation, what it is allowed to influence, who owns the outcome, and what evidence would show the decision was reasonable if challenged later.

#### Personal productivity

**What it looks like:** Individual use: drafting, summarising, research support, internal notes.

**Legal concern:** Confidentiality, privilege, hallucination, over-reliance.

**Decision question:** Is the individual using AI safely, or has informal use already created professional or confidentiality risk?

#### Operational process

**What it looks like:** AI embedded in how the organisation works: intake, triage, claims, compliance, cyber response, procurement.

**Legal concern:** Liability, data protection, bias, resilience, escalation, accountability.

**Decision question:** Has AI become part of how the organisation operates before anyone owns the consequences?

#### Team workflow

**What it looks like:** Shared use inside a team: document review, diligence, knowledge management, case preparation.

**Legal concern:** Supervision, consistency, quality control, matter protocols.

**Decision question:** Is AI use being managed as a controlled workflow, or has it become a habit without ownership?

#### Decision support

**What it looks like:** AI influencing advice, evidence, settlement, valuation, board papers, litigation strategy or regulated decisions.

**Legal concern:** Reliance, causation, professional negligence, directors' duties, litigation exposure.

**Decision question:** Can the organisation explain the role AI played and why relying on it was reasonable?



## THE HUMAN PAUSE

# The moment before commitment

The Human Pause is not a delay mechanism. It is a decision discipline: the moment before procurement, rollout, scale, reliance or board approval hardens into exposure.

It asks what decision is really being made, what business outcome is being pursued, who owns the decision, what risk is being accepted, what alternatives have been rejected, what governance is required, and what must remain human.

The Decision Stack tests whether the decision is clearly defined, owned, governed and defensible across ten layers.

### WHAT WE TEST

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



THE PROBLEM THE HUMAN PAUSE SOLVES

# Most AI failures do not start with model capability

They start earlier. Unclear ownership. Weak business cases. Governance arriving as paperwork after substantive 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. But it is not until the failure reaches a lawyer's desk - in a dispute, disclosure, or complaint - that the real problem emerges: the record shows no one was thinking clearly about what was being approved.

Every week, leadership teams approve AI initiatives. Vendors are selected. Budgets are signed off. Projects are launched. Almost none of those teams asked the right questions first. When they should have asked what decision was being made, who owned it, what could go wrong and whether the organisation could genuinely govern it, they asked instead whether the tool worked. By then, the defensible moment has passed.

**The gap is not in the model. The gap is in the record. And the record is what will be examined later.**



## THE PROBLEM THE HUMAN PAUSE SOLVES

**Not: “Can we implement this AI?”  
That comes later, and only if the  
answer to the earlier question is yes.**

The earlier question is whether the decision itself is defensible. Whether the organisation understands what AI will control or influence, who owns the outcome, what the exposure is if it fails, and whether governance and accountability can be established before commitment.

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:

1. Ownership of the decision.
2. Clarity of the business case.
3. Readiness of the organisation to govern what is being approved.

The model can be excellent whilst the deployment still fails because of these.

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



## HOW WE WORK

# 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 judgement.**



## METHODOLOGY

# The Decision Stack.

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

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

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



## WHERE NAKEDAI FITS

# Before legal exposure crystallises

NakedAI sits before implementation, procurement, rollout, governance failure, litigation, professional negligence, evidential dispute and reputational crisis.

We are not an AI tool vendor. We do not begin with the software. We begin with the decision.

We help boards, leadership teams and their advisers understand where AI is entering the organisation, what it touches, who owns it, what governance is required, and whether the decision can be defended.

## How law firms can use this

- Private client roundtables on AI in the background
- Board-level AI decision briefings for regulated or high-risk clients
- AI diligence readiness sessions for corporate, M&A and PE teams
- Pre-procurement decision reviews before clients buy or embed tools
- AI risk and ownership workshops where legal advice meets governance reality

**The law firm remains a legal adviser. NakedAI brings the decision architecture.**



THE NEXT STEP - TAKE THE HUMAN PAUSE SCORE

# If AI is already being discussed, the decision has begun

Before it becomes expensive, embedded or difficult to defend, pause.

The Human Pause Score is a structured ten-minute diagnostic. It assesses your AI decision across the ten layers of the Decision Stack and returns a 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 a decision that  
is not ready  
to proceed.**

---

**CONTACT**

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

LONDON · EDINBURGH · TORONTO · AUCKLAND