The questions a board cannot delegate

A board can delegate the building of an AI system. It cannot delegate the duty to know what that system decides, on whose data, under whose law, and who answers when it is wrong.

The deadline moved. The duty did not.

Regulation (EU) 2026/1744, the Digital Omnibus on AI, came into force on 27 July 2026. It deferred the centre of the EU AI Act. Obligations for stand alone high risk systems now bite from 2 December 2027, and for high risk systems embedded in regulated products from 2 August 2028. Article 4, the AI literacy duty, was softened from ensuring a sufficient level of literacy among staff to taking measures that support it. The transparency duties in Article 50 were not deferred and have applied since 2 August 2026, to a penalty ceiling of fifteen million euro or three per cent of worldwide annual turnover.[1]

Boards have read the deferral as breathing space. Some of it is. A deferral changes when a regulator can act; it does not change when a shareholder, an insurer, an auditor, a claimant or a customer can.

Directors in the United Kingdom owe the duty under section 174 of the Companies Act 2006 to exercise reasonable care, skill and diligence, measured in part against what a reasonably diligent person with their knowledge and experience would do.[2] That duty has no commencement date to wait for, and it rises as the board's own knowledge rises. Every briefing a board accepts on AI lifts the standard it will later be judged against.

There is no cross sector AI statute in the United Kingdom. No such bill appeared in the May 2026 King's Speech, and ministers have preferred targeted intervention to a single encompassing act.[3] Regulators work with the powers they already hold. So AI exposure reaches a board dressed as something familiar: a data protection matter, a consumer fairness matter, a disclosure matter, a procurement dispute, a control failure. There is no AI compliance box to tick, and no AI defence to plead.

Your controls declaration now covers the model

Provision 29 of the UK Corporate Governance Code 2024 applies for financial years beginning on or after 1 January 2026. The board must declare the effectiveness of its material internal controls, across reporting, compliance and operational controls, and explain the basis on which it reached that view.[4] The first of those declarations reach published annual reports during 2027.

Set that against what has been automated quietly in the past eighteen months. Supplier onboarding checks. Contract review. Candidate screening. Payment anomaly detection. Complaint triage. Draft management information that then travels upward through the reporting chain. Where a model now sits inside a control the board has called material, the declaration covers the model, its supplier, its version and its failure mode.

The question to put to the executive is not whether the organisation uses AI, but which of the controls it is about to declare on depend on a system nobody in the room can explain.

WHAT MOVED

2 Dec 2027
DEFERRED

Stand alone high risk obligations under the EU AI Act, reset by Regulation (EU) 2026/1744.

27 Jul 2026
IN FORCE

The Digital Omnibus on AI amended the Act and moved the timetable, not the duties of directors.

€15m
CEILING

Or three per cent of worldwide turnover, for breach of the Article 50 transparency duties now in force.

WHAT DID NOT MOVE

1 Jan 2026
CONTROLS

Financial years from this date fall under Provision 29. The declaration on material controls follows in 2027.

s.174
DUTY

Reasonable care, skill and diligence. The standard rises with what the board has been told.

2026 DERIVATIVE SUITS

1 **Adobe**
Filed 24 April 2026. Training data provenance and public claims of commercial safety.

2 **Microsoft**
Filed 30 June 2026, with further complaints in August and September.

3 **Nvidia**
Filed 31 July 2026. The theory spread to a third circuit within a hundred days.

Liability is arriving from three directions

Shareholders. Three derivative suits filed in the United States during 2026 put board approval of AI strategy directly in issue. Adobe's chief executive and directors were sued on 24 April 2026; Microsoft's officers on 30 June 2026, with further complaints in August and September; Nvidia's on 31 July 2026.[5] Each concerned the provenance of training data. None required the claimant to prove that infringement had occurred. The allegation was that the leadership understood the provenance risk and proceeded.

The mechanism travels even where the cause of action does not. The standard applied in those complaints came from the companies themselves: codes of conduct, responsible AI commitments, investor disclosures, marketing claims. A published pledge becomes an obligation the board must be able to evidence. Three commitments a board can prove are worth more than a page of principles nobody can audit.

Insurers. The market has begun to price board conduct on AI directly. Through 2025 and 2026 insurers replaced silence with explicit wording. ISO introduced a generative AI exclusion into commercial general liability in January 2026. Absolute AI exclusions have appeared in directors and officers, employment practices and fiduciary lines, drafted to exclude claims arising out of any use, development or deployment of artificial intelligence. Technology errors and omissions cover is narrowing. Cyber has held up best so far.[6]

An absolute AI exclusion on a directors and officers policy is a governance fact, not a procurement detail. Somebody should be able to tell the board what the current endorsement says, and who accepted it.

The board itself. KPMG and the INSEAD Corporate Governance Centre published five AI board governance principles on 14 April 2026, alongside a finding that close to three quarters of boards are seen as having only moderate or limited AI expertise.[7] Perception is the operative word, because perception is what an underwriter, a regulator or a claimant's counsel works from. A board that cannot interrogate a model risk paper will approve it.

PLACEMENT

ToR

AMEND

A duty that appears in no committee’s terms of reference is assigned to nobody.

The committee matters less than the record

The first structural question most boards ask is where AI oversight belongs. The audit committee already owns controls and disclosure. The risk committee already owns appetite. A dedicated technology or AI committee concentrates scarce expertise, and can also quarantine the subject from the directors who most need to follow it. All three arrangements work, and all three fail in the same way.

Placement matters less than three things: a standing item with a named executive owner; a reporting line that surfaces failures rather than adoption; and terms of reference amended so that the duty is written down. Where AI touches a material control or a public disclosure, the audit committee cannot be routed around, whatever else is set up alongside it. An oversight duty that appears in no committee's terms of reference is not assigned to anyone.

THE MEASURE

3 in 4

EXPERTISE

Boards seen as having only moderate or limited AI expertise. KPMG and INSEAD, April 2026.

Seven things cannot be handed down

Management owns models, suppliers, pipelines, evaluation and deployment. Boards that try to own those usually arrive late and add little.

- **The appetite:** which decisions a model may take alone, and which always require a person who can be named.
- **The inventory:** whether one register of AI in use exists, whether it reconciles with what procurement and the expense ledger show, and who

signs it.

- **The declaration:** which material controls now depend on a model, and what evidence supports the statement of effectiveness.
- **The disclosures:** whether public commitments on AI match practice, in the annual report, on the website, in tender responses and in customer contracts.
- **The escalation route:** how a model failure reaches the board, how fast, and through whom.
- **The competence:** whether the board can challenge what it is shown, and what is being done where it cannot.
- **The record:** what the minutes show of the evidence relied on, the alternatives considered and the dissent expressed.

Delegation moves the work. It does not move the duty.

THE TEST

90

DAYS

One register, one page of appetite, one page of standing report. Length is not assurance.

Ninety days of useful work

1. Commission one register, with a single named owner, covering procured tools, features switched on inside existing suppliers, and staff use of consumer accounts. Date it. A register nobody signs is a list.
2. Approve one page of appetite. Name the decisions a model may not take alone. Record it in the minutes rather than in a policy pack.
3. Reconcile disclosure with practice. Withdraw any public claim about AI that cannot be evidenced, before someone else tests it.
4. Ask the broker for every AI endorsement and exclusion across directors and officers, technology errors and omissions, cyber and crime cover. Take the answer in writing, with renewal dates.
5. Put AI dependency into the Provision 29 controls mapping now, while there is time to remediate rather than explain.
6. Reduce the standing report to one page: what changed, what failed, what was withdrawn, what awaits approval. Adoption rates and usage volumes are not oversight.
7. Set the literacy measure. Article 4 asks for measures that support AI literacy among staff, and a board that cannot describe its own is inviting a finding.

The clock runs to December 2027

None of this needs a new statute, a new committee or a maturity model bought from a consultancy. It needs a board that asks for evidence and notices when the evidence does not arrive. The deferral runs to December 2027. That is enough time to build a register, a declaration that holds and a record that shows a board exercising judgement. It is also enough time to arrive at the deadline with a policy document and nothing behind it. Neither outcome waits on a regulator. Both are chosen now, or by default.

Where this series goes next

This paper concerns the duty. The four that follow take the questions it raises one at a time. Paper 02 examines where data lives and who can lawfully compel access to it, since residency and sovereignty are not the same thing. Paper 03 takes the supplier a board cannot exit, and what substitutability means when the model is the product. Paper 04 sets out how a board gets evidence instead of assurance theatre, and what an assurance report on AI should contain. Paper 05 makes the commercial case for sovereignty as a strategy rather than a cost.

Twelve questions, and the answers that should worry you

Ask the executive	Worry if the answer is
How many AI systems are in use across the group, and who signed the register?	"We are pulling a list together." A register without an owner and a date is a list.
Which of our material internal controls now depend on a model?	"None, as far as I know." Absence of awareness is not absence of dependency.
What can a model decide without a named person approving it?	"It only makes recommendations." Then ask how often a recommendation has been rejected, and by whom.
Which public commitments have we made about AI, and what evidences each one?	"It is all in our responsible AI statement." The statement is the claim. The evidence is separate.
Where do the training data and the retrieval sources come from?	"The vendor takes care of that." Licensing a model does not license what it reaches at query time.
What do our insurance policies now say about AI, and which exclusions apply?	"We are fully covered." Ask for the endorsement wording and the renewal date.
What happened the last time a model got something materially wrong?	"Nothing has gone wrong." Either little is deployed, or little is being detected.
Who can lawfully compel access to the data our AI processes?	"It is hosted in the European Union." Residency is geography. Compulsion follows jurisdiction.
If this supplier doubled the price or changed the model version, what would we do?	"We would move to another provider." Ask for the exit test, who ran it, and when.
How would the board learn of an AI failure, and within what time?	"Through the usual channels." Name the trigger, the route and the hours.
What are we declining to do with AI because we cannot yet govern it?	"Nothing is off limits." A board with no declared limit has enthusiasm, not appetite.
Which of us can challenge a model risk paper on its merits?	Silence, or a glance toward one director. That is a board development plan, not a comfort.

NOTES AND SOURCES

1 Regulation (EU) 2026/1744, the Digital Omnibus on AI, in force 27 July 2026. White & Case, *EU AI Omnibus enters into force, amending the AI Act*, July 2026; Gibson Dunn, *EU AI Act Omnibus agreement: postponed high risk deadlines and other key changes*, May 2026. Article 50 duties and penalty ceilings: Jones Walker, *Yes, August 2 still matters*, August 2026.

2 Companies Act 2006, section 174.

3 House of Lords Library, *AI regulation in the UK: debate on the need for cross sector legislation*, 2026.

4 Financial Reporting Council, *UK Corporate Governance Code 2024*, Provision 29, applying to financial years beginning on or after 1 January 2026; FRC Provision 29 information sheet, 29 January 2026.

5 SEIU Pension Plan Master Trust v Narayen (N.D. Cal., filed 24 April 2026); Anderson v Nadella (W.D. Wash., filed 30 June 2026); Berliner v Huang (N.D. Ill., filed 31 July 2026). Ropes & Gray, *AI in the boardroom, shareholders in the courtroom*, September 2026.

6 Fenwick, *The end of silent AI: emerging AI exclusions, coverage fragmentation and practical implications*, 2026.

7 KPMG International and the INSEAD Corporate Governance Centre, *AI Board Governance Principles*, 14 April 2026.

AltLibre Papers, board series, paper 01 of five. Adrian Hollister, AltLibre, 22 August 2026. altlibre.com