



// ORTENT WHITEPAPER • OCTOBER 2026

Nine questions, five stages.

How a board oversees AI.

Most boards get an aggregate answer when they ask where the company is on AI. This paper sets out the nine questions a board should ask every quarter, a five-stage scale that separates saying from doing, and what governed looks like for each question.

Companion to the AI Board Diagnostic at ortent.co/tools/ai-board-diagnostic and the field guide AI and the Board.

Why "where are we on AI?" gets the wrong answer

Every chair asks the CEO the same question: where are we on AI? Most CEOs answer it in aggregate. Copilots in marketing, a forecasting model in sales, an agent in pilot in customer success. The board nods, and nobody in the room can tell whether any of it is working.

AI is not adopted by a company. It is adopted inside individual workflows, and each workflow carries its own risk and its own case for AI. An aggregate answer hides exactly the detail the chair needs: which workflows have changed, what that cost, and what it earned.

The aggregate answer also tilts towards cost savings, because savings are easy to count. That is real money, and boards should take it. But a competitor can buy the same tools next quarter and take the same cost out. The question that tells a board whether AI is changing the company's position is a different one: did it make us any money we could not have made before?

The public record points the same way. Ortent scored 120 software, healthcare and life sciences companies from their own filings in August 2026. 92 report AI revenue that is live. Only 4 say how they fund new AI bets. Companies talk about AI revenue and go quiet on what it cost and how the next bets get paid for. In the NHS Board Radar's October 2026 refresh, AI appears in the board papers of 134 of 143 NHS organisations, yet no board pack yet shows AI producing a measured benefit sustained over more than one period.

In short. A board does not need to understand models to oversee AI. It needs a short, fixed set of questions, asked every quarter in the same order, and an honest scale for how far each answer has got. This paper gives both. The AI Board Diagnostic at ortent.co/tools/ai-board-diagnostic scores a board on them in about ten minutes.

One scale for every question

Each of the nine questions is scored on the same five stages, so the scale reads the same way down the page. The stages separate saying from doing.

0	1	2	3	4
Not discussed	Discussed	Owned	Measured	Governed
The board has not asked the question.	It has come up at the board, but nobody owns it and nothing is written down.	A named executive owns it and the work has started. The artefact is partial or out of date.	The artefact exists and reaches the board on a set cadence, with numbers.	Measured over more than one period, the board has taken a decision on it, and the company could show it to a buyer next week.

The rule that keeps it honest

A policy or a plan on its own never scores above Discussed. A policy says what should happen, not what does. Owned needs a named person who is accountable. Measured needs the artefact itself, not its title: if the answer to a board question is the name of a document, the follow-up is simple. Show me the artefact.

Governed is a high bar on purpose. It needs more than one period of history, a decision the board has taken on the evidence (to fund, stop or fix something), and the confidence that the company could hand it to a buyer next week. Most boards will have few questions at Governed in the first year. That is expected. What matters is the direction of travel from one quarter to the next.

Questions that do not apply

Two questions only apply where AI is in the product: the orchestration layer and data rights. The question on reconstructing a decision applies where AI is used in the product or the company's own operations. A board scores the rest and reads its result out of the questions that apply.

// 01 The map

Which workflows use AI today, at what risk, which are agent-ready, augment or human-only, and which one moves this quarter?

AI is not adopted by a company. It is adopted inside individual workflows, and each one carries its own risk and its own case for AI. An aggregate answer hides exactly what the chair needs to see.

OWNED	A named executive is building the operating model map, and the core workflows are drawn.
MEASURED	The map, with a risk score and an AI class on every workflow, is in the board pack every quarter.
GOVERNED	A current operating model with a risk score and an AI class on every workflow, reviewed by the board each quarter, with a named workflow moving this quarter.

ORTENT ESSAY BEHIND IT: AI READINESS IS A WORKFLOW-BY-WORKFLOW QUESTION

// 02 New money

Does AI make money the company could not have made before?

Savings are easy to count, so AI budgets drift towards them. A competitor can buy the same tools next quarter and take the same cost out. New revenue is the number that shows whether AI is changing the company's position.

OWNED	The CFO or CGO owns separating AI revenue from AI savings, and the definitions are agreed.
MEASURED	Both numbers appear separately in the board pack every quarter.
GOVERNED	New AI revenue is reported to the board separately from AI cost savings, every quarter, with at least two periods of history.

ORTENT ESSAY BEHIND IT: THE EFFICIENCY TRAP

// 03 Funding new bets

How are new AI bets funded, and who owns them?

Early bets on new AI revenue will not clear the return bar set for efficiency. Without a ring-fenced budget and a written stop point, they either never start or never stop.

OWNED	One executive owns new AI bets, and a separate budget line is agreed in principle.
MEASURED	The ring-fenced budget exists, each bet has a written stop point, and the board sees progress every quarter.
GOVERNED	A ring-fenced budget outside the efficiency budget, one named owner, a quarterly review and a written stop point for each bet. At least one bet has been stopped or moved up a stage on that evidence.

ORTENT ESSAY BEHIND IT: REINVENTORS

// 04 The inventory

Can the company produce a list of every AI tool on demand?

AI spend tends to grow department by department, with nobody owning how it fits together. Without a list the board cannot see the risks it is taking. The board does not need to audit the stack. It needs to know the audit can be produced quickly.

OWNED	One executive owns the AI tool inventory, and a first list exists, even if it is incomplete.
MEASURED	The list covers every part of the business, with owner, data seen and last review date, and the board sees it on a set cadence.
GOVERNED	Every AI tool in every part of the business, with its owner, the data it can see and the date it was last reviewed, produced within a day when asked and checked at least twice a year.

ORTENT ESSAY BEHIND IT: THE IT STACK NOBODY AUDITS

// 05 Reconstructing a decision

Could a non-engineer reconstruct any decision the system made in the last twelve months?

If nobody outside engineering can walk through a past decision, the company does not yet have an AI product it can defend in front of a regulator or a customer.

OWNED	An executive owns the audit trail, and logging of AI decisions has started.
MEASURED	Calls, inputs, outputs and versions are logged where the company controls them, and the board sees the coverage on a set cadence.
GOVERNED	Every call, input, output, model version and prompt version is logged where the company controls it, and a replay has been run for the board or an auditor.

ORTENT ESSAY BEHIND IT: DON'T BUILD THE LLM. INTEGRATE THE COMPONENTS.

// 06 The orchestration layer

Who owns the layer that decides which model or component does what?

If the honest answer is the model vendor, the company's advantage depends on that vendor's roadmap, pricing and licence terms.

OWNED	An executive owns the architecture decision and can say which tasks run on which components.
MEASURED	The map of tasks to components, and the dependence on each vendor, reach the board on a set cadence.
GOVERNED	The company owns the orchestration and the audit trail, can name which tasks run on which components, and has a tested plan for changing model vendor.

ORTENT ESSAY BEHIND IT: DON'T BUILD THE LLM. INTEGRATE THE COMPONENTS.

// 07 Data rights

Can the company evidence its rights to the data its AI was trained on and uses?

Questions about data rights tend to surface late and cost most, usually in diligence at exit. A written policy does not answer them. The evidence does.

OWNED	An executive owns data rights for AI features, and the records are being built.
MEASURED	Every AI feature has a data record (sources, rights, retention), and the board sees its status on a set cadence.
GOVERNED	A data record for each AI feature (sources, licences or consents, retention), kept current and reviewed before each new feature ships.

ORTENT ESSAY BEHIND IT: ASK FOR PROOF, NOT POLICY

// 08 AI claims

Is there evidence behind every AI claim the company makes?

Every categorical claim on a product (autonomous, AI-powered, an accuracy figure) needs a file behind it. In a sale process, a buyer will ask for it.

OWNED	An executive owns AI claims, and a list of the claims made in public and in sales material exists.
MEASURED	Each claim has a substantiation file, and the board sees how many claims have evidence behind them on a set cadence.
GOVERNED	A substantiation file for every AI claim on the website, in proposals and in the pitch deck (autonomous, AI-powered, any accuracy figure), checked before each release and before a raise or sale.

ORTENT ESSAY BEHIND IT: CONSISTENCY BEATS POLISH. SPEED IS A SIGNAL.

// 09 Regulation and standards

Does the board know which AI rules apply to the company, and which framework it runs its AI work to?

The rules are moving. The EU's Digital Omnibus on AI, in force since July 2026, pushed the AI Act's high-risk obligations back to December 2027, and to August 2028 for AI built into regulated products, while most of the Act's transparency duties were not delayed. A board that has not written down which rules apply, and from when, tends to learn the dates from a customer, an insurer or a buyer. A recognised framework, such as ISO/IEC 42001 (the certifiable AI management system standard) or the NIST AI Risk Management Framework, gives the work a shape a buyer can check.

OWNED	An executive owns the regulatory view, and a first list of the rules that may apply exists.
MEASURED	The list, with dates and the company's position against each rule, and progress against a chosen framework, reaches the board on a set cadence.
GOVERNED	A written view of which AI rules apply and from when (for example the EU AI Act, where the company sells into or operates in the EU), mapped to a recognised framework such as ISO/IEC 42001 or the NIST AI RMF, and reviewed by the board whenever the rules change.

ORTENT ESSAY BEHIND IT: ASK FOR PROOF, NOT POLICY

Not legal advice. This question asks whether the board has a written view of the rules and keeps it current, not whether that view is right. Where the company is unsure which rules apply, take legal advice.

One board, scored on all nine

The company below is fictional and illustrative. It is a growth-stage B2B software company with AI features in its product, a PE investor on the board, and a sale process expected in two to three years. The chair and the CEO filled in the diagnostic separately and agreed the stages below.

QUESTION	STAGE REACHED
1. The map	Owned
2. New money	Discussed
3. Funding new bets	Not discussed
4. The inventory	Discussed
5. Reconstructing a decision	Owned
6. The orchestration layer	Measured
7. Data rights	Discussed
8. AI claims	Discussed
9. Regulation and standards	Discussed

The read. 1 of 9 questions is measured or governed. 8 are still at owned or below. Six questions sit below Owned, which points to the operating model rather than one missing owner. That is the shape the 90-day Diagnostic is built for.

Fix first: funding new bets. It is the only question the board has never asked. There is no budget for new AI revenue outside the efficiency budget, so every bet is judged on the same return as a cost saving, and none has started.

The 90-day plan. The two lowest questions move up first. Ties are broken in the order a buyer asks in diligence, so AI claims comes before data rights, regulation, the inventory and new money.

- **Funding new bets, to Owned.** The CEO names one executive to own new AI bets. The board agrees a small ring-fenced budget line in principle and sees the first two candidate bets, each with a written stop point, at its next meeting.
- **AI claims, to Owned.** The CMO lists every AI claim on the website, in proposals and in the pitch deck. The board sees the list next quarter, with the claims that have no evidence behind them marked.

Run again next quarter, the same board should see two questions move up a stage and none move down.

Same questions, same order, every quarter

The value is in the repetition. Ask the nine questions in the same order every quarter, the way a board works through a risk map, and track how the answers move, not just what they say.

The persistence rule

One quarter without a good answer is a process problem. Two quarters is an ownership problem. Three quarters means the leadership has not taken it on. A question that sits at Discussed for three meetings is telling the board something about the executive team, not about AI.

Two people, filled in separately

The diagnostic works best when two people fill it in separately: the chair and the CEO, or a NED and the CFO. Where one answer is lower than the other is where the board conversation should start. The gap usually means one side has seen an artefact the other has not, or one side is counting a plan as progress.

What goes in the board pack

One page: the nine questions, the stage each has reached, the stage last quarter, and the one question being moved this quarter with its owner. The artefacts behind Measured and Governed sit in the appendix, not in the main pack.

Where this comes from. The nine questions come from Ortent's field guide AI and the Board and the essays behind it. The five stages borrow the logic of the NHS Board Radar, which separates what a board says about a priority from what it has delivered, and the evidence rules of the Diligence Radar, where a policy on its own is never proof.

From the read to the work

The diagnostic names the question to fix first. What fits depends on how many questions sit below Owned. Ortent's AI work runs through the same five services as everything else. Engagements start with a 30-minute call at ortent.co/contact.

Sprint

One or two questions below Owned, each with a clear fix. One board-ready written answer in two weeks.

90-day Diagnostic

Three or more below Owned. A full commercial and operating read against evidence, including the operating model and the AI overlay.

Advisory or NED

Everything at Owned or above. Quarterly challenge on the nine questions from someone who has sat on both sides of the table.

About the author

Andrew Wyatt is Principal of Ortent Advisory. He is a four-time exit operator across enterprise software and telecoms, most recently Chief Growth Officer at Sapio Sciences. He has over 30 years operating, with board roles across PE and VC-backed businesses since 2006.

Take it further

The AI Board Diagnostic: ortent.co/tools/ai-board-diagnostic

The Claude prompt, for the version scored from real evidence: ortent.co/tools/ai-board-diagnostic/prompt

The field guide AI and the Board: ortent.co/insights/ai-and-the-board

What 120 companies disclose about AI: ortent.co/tools/ai-value-gap/outside-in

This paper is general guidance for boards. It is not legal, regulatory or financial advice. © 2026 Ortent Advisory Ltd. Ortent is a registered trade mark of Ortent Advisory Ltd.