# The Critical Friend's Brief

## *AI at Board Level*

A briefing for Boards and NEDs on AI governance in mid-2026

*AI can be technical.*
*Its adoption is not.*
*This briefing is for the people who have to govern it.*

# A note before you start

For over 30 years I've been called in to diagnose issues with, then rescue, projects across banks, building societies, gaming companies, retail & travel. The pattern repeats: by the time someone calls me, the Board already knew something was wrong but didn't have the language to say so out loud.

AI is the next version of this conversation.

I'm writing this briefing because there are too many Board meetings where AI gets a ten-minute slot, a slide deck full of buzzwords, and a polite nod from the executive. Nobody pushes back, because nobody feels qualified to push back. That isn't governance. That's theatre.

This **is not** a "what is AI?" primer. You don't need to know in detail how a neural network works to govern one, any more than you need to know how a combustion engine works to chair a logistics business.
This **is** the set of questions I think every Board should be asking, and the red flags I think every NED should be watching for. That all said, I strongly believe you need foundational knowledge of AI to engage in these discussions – if you're still stuck on this point then let's chat.

This briefing is direct, unhedged and occasionally uncomfortable. That's deliberate. AI governance done politely is AI governance done badly. If you finish this briefing with more questions than answers, it's done its job.

*Calm and confident,*

**Geoff Porritt** - Director and Principal Consultant, ENOS Limited

***If your Board had no structured AI education in the last twelve months, the questions in this pack will be ducked successfully by your Exec***

# Where does AI sit?

Every generation of senior leaders has had to govern a technology shift they didn't fully understand; Y2K, dotcom, video conferencing and cloud. Now AI, and on its heels, Agentic AI. Some shifts were overhyped and faded. Some were quietly transformative. Some are still unfolding, and we won't know for a decade how they land.

*Here's the uncomfortable truth: your job isn't to predict which one AI turns out to be. Your job is to govern it **AS IF** it could be any of them.*

| Y2K | Dotcom | Video Conferencing | AI / Agentic AI |
|---|---|---|---|
| *Overhyped, briefly disruptive, faded* | *Crashed quickly, then quietly rebuilt the economy* | *Slow burn, then suddenly indispensable* | *Still unfolding - your call to make* |

Sat in 2026, the most likely outcome is somewhere between Dotcom and Video Conferencing - a noisy first wave, a correction, and then a quieter integration into how organisations actually run. But "most likely" is not "certain", and Boards don't get to govern based on "most likely".

One prominent AI vendor refused wide-scope US government use cases on ethical grounds; other model providers said yes to the same brief. That isn't a technology question. **That's a values question - and values questions are Board questions.**

For example, two boards of similar size could legitimately reach different conclusions on whether to use AI in credit checks, and both can be governed; the failure is not having had the conversation.

# What's the key difference about this technology?

*If you take just **ONE** thing away from this briefing, please take this.*

## Google knows facts.

When you type a question into a search engine, it goes and fetches information that exists somewhere in the world and shows you the bit it thinks is most relevant. It might be wrong about relevance, but the information itself is real. It came from a source.

Dry, transactional, factual.

## AI predicts the next word.

When you type a question into ChatGPT, Claude, Copilot, Gemini or DeepSeek, it does not go and fetch anything. It predicts, one word at a time, the most statistically likely next word given everything that came before it.

Rich, verbose, conversational - and crucially, **IT DOES NOT KNOW** whether what it's saying is true.

## A worked example, and genuine true story.

One Sunday I asked an AI: "What day is tomorrow?" The answer came back: "Tomorrow is Tuesday, because today is Monday". Why? It turned out the AI had assumed Sunday was a working day for me and, as I'd done nothing the day before, "today" must be Monday. In the background it had inferred my regular working pattern, without me telling it. Then it had ***predicted***.

Every word the AI produced is grammatically and contextually correct. The reasoning sounds confident. And yet it is complete nonsense. Now imagine that same confident, plausible, invented reasoning applied to a credit decision. A clinical recommendation. A sanctions screening result.

***This is the governance challenge. In one paragraph.***

# AI usage - The single biggest cost-overrun risk

Many Boards are used to software contracts that look like this: pay a licence fee, pay a per-user subscription, that's roughly the year's spend. Predictable. Budgetable. Boring. AI doesn't work that way. If your finance team is provisioning for AI tools using the old model, you already have a problem you don't yet know about.

## Think of it like a club.

- **The subscription is the door charge.** Pay it, you get in
- **The AI-use is the drinks.** Every question you ask, every document you upload, every response the AI generates - each is a round at the bar. Some rounds are cheap. Some are eye-wateringly expensive
- **Now put an Agent in the room.** An Agent is allowed to keep ordering rounds, without checking with you first, to "get the job done". If it loops, gets confused, or chooses an inefficient path, it'll keep ordering until you eventually notice
- ***NEVER leave the Agent with your credit card....!***

*AI costs scale with USAGE, not with users. A single power user running complex agent workflows can rack up more spend in a week than a hundred light users do in a year.*

**Three things to confirm with your executive:**

1. **Who has set the cost ceilings, and where?** Per-user, per-agent, per-month - all three should have limits.
2. **What's the monitoring cadence?** Monthly is too slow. Daily is realistic. Real-time is best for high-volume use cases.
3. **What's the escalation trigger?** At what spend level does someone get woken up?

**WHAT IS AGENT/AGENTIC AI?**

'Traditional' **AI** is a human user prompt that creates AI output e.g. ChatGPT, CoPilot, etc (there are dozens)
**Agent(ic) AI** is where AI Agents are given a business goal (e.g. run a marketing campaign for launching product X), the tools (e.g. other software, purchasing power) and left to create and execute the plan themselves.

# Whole new classes of Risk - not just more of the old ones

AI doesn't just add to the risks you already manage; it introduces whole new CLASSES of risk your current framework may not be built to catch. Your job on this page isn't to work every risk to the ground - it's to satisfy yourselves that each class below has an owner, a control and a test, because the ones you haven't named are the ones that find you first.

## ACCURACY & HALLUCINATION

*The model sounds certain about its output. That tells you nothing about whether it is right.*

1. Where do AI outputs feed decisions with no human in the loop?
2. What's our tolerance for a confident, plausible, wrong answer - by use case?
3. How would we catch a hallucination before a customer or regulator did?

## COMPETITIVE vs PEERS

*Are we behind or ahead, and on which dimensions?*

1. Is AI bringing down costs for our peers? How do we know?
2. Customer expectation of AI-grade service; what opinions are we asking, and of whom? What is other market research saying?
3. Capability and talent gap vs peers; are we losing or gaining talent due to our use of AI?

## REGULATORY

*What applies? What's coming? What are we ready for?*

1. EU AI Act - duties phased to 2026
2. ISO/IEC 42001
3. UK GDPR Article 22 and ICO guidance
4. Sector-specific obligations

## SECURITY

*Have each of these been identified as distinct risk classes, and have each been tested?*

1. Malicious prompt injection
2. Data leakage & training-data provenance
3. Model integrity & tampering
4. PROD vs non-PROD separation

## BIAS, FAIRNESS & DISCRIMINATION

*Could the model treat people unequally - and could we prove it didn't?*

1. Decisions made about individuals (HR, credit, claims)
2. Proxy discrimination hidden in the model or training data
3. Explainability - can we defend a single decision to the person impacted?

## CONCENTRATION & LIABILITY

*If the AI gets it wrong, or is unavailable, who carries the cost - and is anyone insuring it?*

1. Is there an over-reliance on one model provider or model?
2. Legal position when a wrong AI decision causes loss - our's, the vendors or unclear?
3. Is anyone insuring this? Or are we carrying full exposure?

***If your Exec can't show your Board where each of these is owned, mapped and tested, your risk framework hasn't caught up.***

# Q1 What is our stance on AI?



*If your Board hasn't formally agreed an answer to this, you don't have an AI strategy; you have a collection of accidents waiting to be called a strategy in hindsight. Every other question in this pack assumes you have answered this one. Many Boards haven't.*

## What you might hear from your Exec

*"We're keeping a close eye on the space."*

*"We're already using it across the business."*

*"It's covered in our digital transformation programme."*

*"We're trialling Copilot at the moment."*

## What you ought to hear

**A clear, single-sentence position** – are you a leader? fast follower? careful adopter? or deliberate deferrer? - with a stated reason.

"Fast follower because our regulator hasn't finalised its position" is defensible. So is "defer until we've fixed our data foundations - we'd just be putting GIGO through a faster pipe".

Both are stances. Both can be governed. "We're watching the space" is neither.

**THE MOST DANGEROUS ANSWER:** *"Different parts of the business are doing different things."*

**Translation:** You have shadow AI. Individuals and teams using tools you don't have contracts with, sharing data you haven't risk-assessed, and creating audit problems that haven't yet surfaced. Your first job isn't to clamp down. It's to **know**. You cannot govern what you cannot see.

**Follow-up to put to your Exec, if that first answer felt thin:**

*"Where is our stance on AI documented? Who signed it off? When was it last reviewed - and against what criteria?"*

***If those three get a confident answer, you have a stance. If they don't, you have homework.***

# Q2 Where are our AI risks, issues and opportunities?



*The most common driver of Board panic on AI in 2026 is the feeling of being left behind. That feeling has its place, but it isn't the whole question. The Board's job here is broader: where is the organisation EXPOSED, and where could it WIN? Treat AI as you would any other agenda item - risks AND opportunities, evidenced and on the same page.*

## What you might hear from your Exec

*"Our competitors are way ahead of us."*

*"We're at the same stage as everyone else in our sector."*

*"We're keeping a close eye on the regulatory position."*

*"Cyber is covered by our existing controls."*

## What you ought to hear

**For each of the risk classes set out earlier** - the Board should hear evidence of where it is owned, where it is mapped, how it has been tested and which sub-committee carries it

**Balancing Competing Goals** - The accepted, documented trade-off position between the competing forces of AI and ESG

**Strategic Opportunity** - Futures we might close off by not investing. Any wins by moving now?

**THE MOST DANGEROUS ANSWER:** *"AI is covered by our existing risk framework."*

**Translation:** Nobody has actually mapped AI risk separately. AI introduces attack vectors, failure modes and regulatory obligations that pre-AI frameworks were never designed for. "It's covered" without showing the work is the risk equivalent of "we have antivirus, we're fine". Push back.

**Follow-up to put to your Exec, if that first answer felt thin:**

*"Walk us through the evidence on your top three risks - and tell us which one of them keeps you awake at night."*

***Being exposed is survivable. Not knowing you are exposed is not.***

# Q3 Do we have a plan, and who owns it?



*Plans get nodded through at Board meetings all the time. AI plans rarely survive contact with the budget, the data team, or the regulator. Two questions worth pressing: "Is this a plan or a wish list?" and "If it fails, whose head is on the block?"*

## What you might hear from your Exec

*"Our CTO is leading it."*

*"We've engaged [a big consultancy] to develop the roadmap."*

*"We're starting with use cases and building from there."*

*"It's part of our digital transformation programme."*

## What you ought to hear

**A named Executive owner** - with AI in their objectives and pay tied to it

**A deliberate choice between consultants and internal** - fast/costly vs slow/owned

**Top-down or bottom-up** - either can be right. Bottom-up done deliberately (training, ring-fenced time, harvest owner, kill-criteria) is strong; by accident it will drift

**A capacity plan** - to free people up to actually ***do*** AI - nobody does it well in their spare time

**THE MOST DANGEROUS ANSWER:** *"It's a steering group across the leadership team."*

**Translation:** Nobody owns it. AI by committee means nobody can be sacked when it doesn't deliver, which means nobody loses sleep, which means it doesn't deliver. Shared accountability is a polite phrase for none. For regulated entities, name the SMF and document the control structure underneath

**Follow-up to put to your Exec, if that first answer felt thin:**

*"If this AI programme fails to land in eighteen months, whose objectives will reflect that, and what will the consequence be?"*

***A plan without an accountable owner is a press release. Read it as such.***

# Q4 What are the business goals, and how do we measure them?



*This is where most AI initiatives sleepwalk. "Improve productivity", "stay competitive", "explore the technology" - all valid-sounding, none measurable. Six months in, nobody can say whether it is working, so the budget gets renewed by default. That isn't investment. That's an annuity to your model provider.*

## What you might hear from your Exec

*"We're aiming for a 20% productivity uplift."*

*"We're testing customer-facing use cases."*

*"We're exploring where AI can add value."*

*"We're starting internal, to learn before going external."*

## What you ought to hear

**Goals stated in business outcome terms** - revenue, cost, risk, customer satisfaction, time-to-market. Tied to a stated strategic intent - which might be "build capability" as well as "increase revenue"

**A baseline measurement** - captured ***before*** deployment. Without it, all productivity claims after the fact are untestable

**A measurement cadence** - with a named owner and a reporting line to the Board

**A success threshold** - that, if missed, triggers a Board-level review and not a quiet roll-over

**THE MOST DANGEROUS ANSWER:** *"It's hard to measure, because it's transformational."*

**Translation:** We do not want to be measured. Anything labelled "transformational" that can't be measured is a budget hiding behind language. Push back. Politely, but firmly.

**Follow-up to put to your Exec, if that first answer felt thin:**

*"What baseline did we capture before this went live, and what does the variance look like today?"*

***If you can't measure it, you can't govern it. If you don't govern it, you'll end up explaining why it failed.***

# Q5 What is it costing us, and what is the ROI test?



*The Board-level version of the cost conversation. Not "did we set limits?" (that's Exec hygiene). Rather: "Are we getting value for what we're spending, and at what point does that change?"*

## What you might hear from your Exec

*"It's a small spend relative to our IT budget."*

*"We're tracking against the business case."*

*"The ROI is in productivity savings."*

*"It's strategic - the ROI is longer-term."*

## What you ought to hear

**Total AI spend** - across the organisation - including shadow usage and AI features embedded in tools you've already bought

**A stated payback window** - with named milestones

**A floor** - below which the investment is killed

**A ceiling** - above which the Board is consulted before more is spent

**THE MOST DANGEROUS ANSWER:** *"Model providers' prices are coming down - our cost per token will fall."*

**Translation:** Probably true at the unit level. Utterly irrelevant if usage is growing faster than unit cost is falling - which it usually is. Your bill is going up, not down. This is the cost equivalent of "I save money every time I buy things in the sales".

**Follow-up to put to your Exec, if that first answer felt thin:**

*"What did we spend on AI last quarter, including shadow usage and embedded features - and what is it forecast to be next quarter?"*

***AI is the first technology spend where the bill arrives faster than the value. Set the floor and the ceiling NOW.***

# Q6 Who does this affect, and how are we preparing them?



*The technology question dressed up as a people question - or the other way around. Either way, this is the one most likely to come back to bite the Board if it has been ducked. AI changes who does what, how they do it, and who they do it with. Get this wrong and you lose the people whose tacit knowledge made the business work.*

## What you might hear from your Exec

*"We're rolling Copilot training out across the company."*

*"We're starting with a power-user group, then scaling."*

*"HR are leading on the people side."*

*"It's a personal-development conversation - self-paced learning."*

## What you ought to hear

**Who is in scope** - whole company, or a constrained group, with a stated reason

**Hires and leavers** - are we attracting and retaining the people we need, with evidence in the talent data?

**Critical-knowledge risk** - what walks out the door if specific people leave

**Which jobs will materially change in 2 years** – and a plan for the people in them

**Tone** - is AI being framed by leadership as something done ***with*** the workforce, or done ***to*** it? Have unions been kept in the loop?

**THE MOST DANGEROUS ANSWER:** *"We don't want to scare anyone, so we're keeping the messaging high-level."*

**Translation:** We know it will affect jobs, and we are hoping it will be done before they notice. People are not stupid. They notice. Trust evaporates fast and is very hard to win back; ask anyone who managed an outsourcing programme in the 2010s.

**Follow-up to put to your Exec, if that first answer felt thin:**

*"Which roles in this organisation will look materially different in twenty-four months, and what is our plan for the people doing them today?"*

***AI strategy IS workforce strategy. Treat it any other way and you are creating a problem you will own.***

# Score yourselves. Honestly.

*Tick where you can confidently say "yes, with evidence". Total your ticks and read the scoring panel.*

## Q1. Stance

- ☐ Written, Board-approved AI stance with a reason?
- ☐ Reviewed in the last twelve months?
- ☐ Every Board member can articulate it in one sentence?

/3

## Q2. Risks, issues and opportunities

- ☐ Evidence of competitive position - cost, customer, capability?
- ☐ Regulatory obligations mapped - EU AI Act, ICO, sector?
- ☐ AI security as distinct risk class?
- ☐ Documented delegation & consequences of AI decision-making?
- ☐ Is the trade-off position of AI vs ESG made clear?
- ☐ Future capabilities you might be closing off captured?

/6

## Q3. Plan and ownership

- ☐ Named Executive accountable for AI delivery?
- ☐ Their pay and objectives tied to AI outcomes?
- ☐ Deliberate choice between external and internal build?
- ☐ Tool-fit reappraisal cadence in place?
- ☐ Capacity-release - people freed up to do the work?

/5

## Q4. Business goal and measurement

- ☐ Goals stated in business-outcome terms?
- ☐ Baseline measurement captured BEFORE commencement?
- ☐ Measurement cadence with a named owner?
- ☐ Success threshold that triggers a Board review if missed?

/4

## Q5. Cost and ROI

- ☐ Total AI spend known, including shadow and embedded?
- ☐ Stated payback window with milestones?
- ☐ Floor below which the investment is killed?
- ☐ Ceiling above which Board consultation is required?

/4

## Q6. People and upskilling

- ☐ Talent impact - hires and leavers - measured in last 12 months?
- ☐ Critical-knowledge risk by role assessed?
- ☐ Whose roles will materially change in twenty-four months?
- ☐ Honest communication with the workforce about direction of travel?

/4

**Count your Yes answers (out of 26):**

**20-26** doing the job, work on closing any gaps **13-19** meaningful gaps, not yet a crisis **7-12** governance problem **0-6** get in touch

# Tested against the bad days?

*Most of this pack has been about steady-state governance. This page is about the bad day. Every scenario below has already happened to somebody in the last eighteen months. If it happened next Tuesday - who would do what, and how quickly?*

## COST RUNAWAY

*Your AI bill has tripled this month. Nobody is yet sure why.*

1. Who notices, and when?
2. Who has authority to switch the offending workflow off?
3. Can you reconstruct what drove the spike line-by-line?

## MODEL MAKER DISAPPEARS

*Your primary provider files for administration. Or export controls come in. Production access ends in 14 days.*

1. Which processes depend on that model, listed and ranked?
2. Have you tested another provider's model on the same workflow?
3. What's the contractual position on the data they hold?

## OUTAGE

*Provider down at 09:00 Tuesday. ETA unclear. By Friday you've lost three days.*

1. Which workflows degrade gracefully? Which fail hard?
2. Do your staff still know how to do the job manually?
3. What is the customer-communication plan?

## REGULATOR AT THE DOOR

*An inspector arrives wanting your AI risk register, data-flow mapping, model-governance log, and oversight evidence. By end of play.*

1. Which of these four is closest to what you have already? What's the gap to a regulator-ready version?
2. Are they current? Within the last quarter? Or older?
3. Who owns each, and are they in the building today?

## THE PUBLIC FAILURE

*Your AI tool has produced a high-profile bad answer. Trending on social media. Three journalists calling.*

1. Who is briefed and ready to speak by lunchtime?
2. Is the workflow paused, kept running, or rolled back?
3. What does the Board statement say & who pre-approved it?

***If any one of those leave you uncomfortable, your Exec hasn't done the homework***

# Now do something about it.

*If this briefing landed, take one of three next steps.*

PRIMARY

## COMMISSION OUR AI HEALTH CHECK FOR YOUR BOARD

A two-week diagnostic. Structured interviews with your Board and Executive. Document review. Written, board-ready report. 90-minute live debrief. 90-day prioritised action plan.

This is what ENOS does.

SECONDARY

## COMMISSION ONE FOR YOUR LEADERSHIP TEAM

Same diagnostic, framed for executive teams who want an independent read - and a Critical Friend in the room - before they take the next plan to their Board.

EXPLORATORY

## TALK TO US FIRST

A 30-minute exploratory call with Geoff.

No commitment. No slide deck. No follow-up sales sequence.

**Typical engagement £8k - £25k, depending on organisation size and depth required**

Existing ENOS 360° Health Check clients: from £6,500

*We don't sell videos. We sell X-ray diagnostics. That's the next conversation.*

**Geoff Porritt** - *Director and Principal Consultant, ENOS Limited*

geoff@enos.ltd.uk - +44 (0) 7946 314 316 - www.enos.ltd.uk