KENNEDY ADVISORY

BOARD CHECKLIST

10 Questions Your Board Should Ask About AI

A board-ready checklist: why each question matters, what a good answer sounds like, and the red flags that warrant a deeper look.

Michael Kennedy · Former Group CIO, Robert Walters · Singapore / Global · July 2026

WHY THIS CHECKLIST

Most boards are asking about AI. Far fewer are asking the right questions — and almost none know what a good answer sounds like.

AI oversight fails in a predictable way: management presents an enthusiastic update, the board nods, and nobody in the room can tell the difference between genuine control and confident narrative. These ten questions are designed to make that difference visible. They need no technical background — only the willingness to ask for evidence rather than assurance.

Each question comes with the reason it matters, what a substantive answer sounds like, and the red-flag answers that should prompt a deeper look. They are drawn from taking AI into production in a live global business — not from theory.

> **How to use it:** don’t table all ten at once. Pick two or three per board meeting, ask for evidence, minute the answers — and revisit them in six months. The comparison over time tells you more than any single answer.

01 Where is AI already in use in our organisation — including where we haven't approved it?

Shadow AI is the default state, not the exception. In most organisations, staff adopted public AI tools long before anyone wrote a policy, and boards typically discover this only after sensitive data has already left the building. You cannot govern what you cannot see.

A good answer sounds like:

“We maintain a living AI inventory — approved tools, AI features embedded in vendor products, and monitored shadow use — and it's reviewed quarterly.”

A red flag sounds like:

“We have a policy against unapproved tools.” A policy is not visibility. Ask when the inventory was last updated; if there isn't one, that is the answer.

02 Who is accountable when an AI-assisted decision goes wrong?

AI diffuses accountability. When a model contributes to a bad outcome, everyone points at the vendor, the data, or the algorithm. The board should expect a single named executive owner for each material AI use case — not a working group.

A good answer sounds like:

“Each material use case has a named executive owner, and the escalation path has been tested, not just documented.”

A red flag sounds like:

“The AI steering committee owns it.” Committees don't take the phone call when a client has been harmed. Names do.

03 Which decisions have we agreed AI must never make alone?

The human-in-the-loop line should be drawn deliberately, in advance, by the business — not discovered after an incident. Decisions affecting individuals' rights, money, health or employment deserve explicit treatment.

A good answer sounds like:

“We hold a documented list of decisions requiring human sign-off — hiring, credit, pricing, anything affecting individual customers — and controls that enforce it.”

A red flag sounds like:

“Our people always review the output.” At volume, review becomes theatre. Ask how often reviewers actually override the machine; if the answer is ‘almost never’, nobody is reviewing anything.

04 What data is leaving the organisation through AI tools — ours and our vendors'?

Contracts, client records, source code and unannounced results flow into prompts every day. The exposure isn't just your staff — it's every vendor that quietly routes your data through a model.

A good answer sounds like:

"We have technical controls on AI endpoints, enterprise agreements with no-training clauses, and AI data-handling terms in vendor contracts."

A red flag sounds like:

"We've blocked ChatGPT." Blocking one tool pushes usage onto personal devices. The goal is governed access, not prohibition.

05 How would we know if one of our AI systems started failing quietly?

Traditional systems fail loudly — they crash. AI fails silently: accuracy drifts, bias creeps in, and outputs degrade while everything looks operational. Silent failure is the defining risk of production AI.

A good answer sounds like:

"Each production use case has defined performance metrics, live monitoring, and thresholds that trigger human review."

A red flag sounds like:

"The vendor monitors it." The vendor monitors their model, not your outcomes. Those are different things.

06 What are our vendors and suppliers doing with AI on our behalf?

Your risk surface now includes every supplier that added AI features to their product without asking you. Legal, HR, finance and CRM platforms have all done exactly that in the last two years.

A good answer sounds like:

"AI questions are embedded in vendor due diligence and contract renewals, and we keep a register of vendor AI that touches our data."

A red flag sounds like:

"That's covered by our existing vendor management." Ask when the vendor questionnaire was last updated. If it predates 2023, it doesn't mention AI.

07 Does what we do with AI match what we tell customers and regulators?

The gap between public statements and internal practice is where regulatory and reputational risk lives. Privacy notices, client contracts and marketing claims were mostly written before AI arrived.

A good answer sounds like:

"We've reviewed our privacy notices, client contracts and public claims against actual AI use — and closed the gaps we found."

A red flag sounds like:

Nobody in the room knows what the privacy notice says about automated processing. That means nobody has checked.

08 What would an AI incident actually cost us — and do we have a response plan?

AI incidents don't fit the standard cyber playbook: hallucinated advice given to a client, a biased decision at scale, confidential data surfaced through a prompt. Different failure modes need different responses.

A good answer sounds like:

"AI scenarios are in our incident response and crisis communication plans, and we've exercised at least one of them."

A red flag sounds like:

"Our cyber incident plan covers it." It probably doesn't. Ask who has the authority to switch a production model off — if the answer takes more than five seconds, there is no plan.

09 Are we investing in AI to a strategy, or to a fear of missing out?

FOMO spend produces pilots, not P&L.; AI investment deserves the same discipline as any other capital allocation: a named business outcome, a baseline, and the courage to stop what isn't working.

A good answer sounds like:

"Every initiative has a named business outcome and baseline metrics, and pilots have kill criteria — several have already been stopped."

A red flag sounds like:

A count of pilots presented as the success metric. Pilots are cost. Production is value. A hundred pilots and nothing in production is a red flag, not a programme.

10 What is AI doing to our people and skills — and what is the plan?

Capability either walks out the door or quietly atrophies. If AI is drafting the work your juniors used to do, where do your seniors come from in five years? Augmentation needs deliberate design.

A good answer sounds like:

"We have a role-level view of exposure, a reskilling investment, and we've been straight with our people about intent."

A red flag sounds like:

Silence — or treating AI purely as a headcount opportunity. Announce that, and your best people (the ones with options) leave first.

THE QUICK VERSION

Scorecard for the boardroom

Tick each question your board has asked in the last twelve months — and received an evidenced answer to, not just assurance.

- ☐ **01** Where is AI already in use in our organisation — including where we haven't approved it?
- ☐ **02** Who is accountable when an AI-assisted decision goes wrong?
- ☐ **03** Which decisions have we agreed AI must never make alone?
- ☐ **04** What data is leaving the organisation through AI tools — ours and our vendors'?
- ☐ **05** How would we know if one of our AI systems started failing quietly?
- ☐ **06** What are our vendors and suppliers doing with AI on our behalf?

- ☐ **07** Does what we do with AI match what we tell customers and regulators?
- ☐ **08** What would an AI incident actually cost us — and do we have a response plan?
- ☐ **09** Are we investing in AI to a strategy, or to a fear of missing out?
- ☐ **10** What is AI doing to our people and skills — and what is the plan?

> **Scoring:** 8–10 evidenced answers — your oversight is ahead of most boards. 4–7 — typical, with meaningful blind spots. 0–3 — your AI risk is being governed by luck. That is worth a conversation.

ABOUT THE AUTHOR

Michael Kennedy is the founder of Kennedy Advisory and former Group CIO of Robert Walters, where he led a 150+ person technology organisation across 14 countries, delivered £3M+ in annualised savings, achieved ISO/IEC 27001 across all jurisdictions and took AI safely into production. He is CISSP certified and based in Singapore with a global remit.

Working through this for real?

Independent judgement on technology, cyber risk and AI — as an interim or fractional CIO, or at the board table. Engagements start with a conversation, not a proposal.

michael@kennedyadvisory.co · kennedyadvisory.co · Singapore / Global