

Thought leadership paper

6 Questions Directors Should Be Asking About AI

Author: Simon Poole, AI for Boards expert
Prepared for distribution • 14 February 2026

In some organisations AI is now embedded in everyday operations—from customer service and marketing to pricing, hiring, credit decisions, cybersecurity and software development. Other organisations are just getting started.

The Board's job isn't to pick models or review prompts. It's to set intent, approve guardrails, demand visibility, and ensure accountability.

Use these questions to guide effective, non-technical Board oversight.

Where exactly are we using AI today—and what type is it?

Why this matters:

“AI” can mean very different things (traditional machine learning, generative AI, or more autonomous “agentic” workflows). Each carries different risks, controls, and regulatory expectations. If the Board can’t see where AI is used and what kind it is, it can’t oversee it.

What business outcomes are we targeting, and how does AI support our strategy?

Why this matters:

Strong organisations don’t pursue an “AI strategy” in isolation—they use AI to deliver the business strategy more effectively (growth, service, speed, cost, risk reduction). Clarity here prevents expensive experimentation that doesn’t move the needle and helps management prioritise high-value use cases. There may be a case for AI experimentation when you are just getting started and building organisational knowledge and capacity, but the end goal remains to use AI to achieve your business strategy.

What is our AI governance structure—and is it working in practice?

Why this matters:

Good governance typically includes: a Board-approved AI usage policy, an executive steering committee with clear decision rights, and a living register of material AI use cases with controls and status reporting. Without this, AI adoption becomes fragmented and proceeds without proper control and Board oversight. As Directors you will be ‘flying blind’. (See my thought leadership paper on AI governance for more on this)

What would a “bad AI incident” look like for our organisation?

Why this matters:

Boards manage downside risk best when it’s made concrete. Incidents might include harmful customer outcomes, discrimination, misleading advice, data leakage, fraud enablement, regulatory breach, cyber compromise via AI tools, or uncontrolled autonomous actions in internal systems. Defining plausible scenarios drives the right controls, testing, monitoring, and crisis playbooks.

What are the performance measures and success criteria for our AI use cases?

Why this matters:

AI should be measured like any other material capability: effectiveness (quality/accuracy), efficiency (time/cost saved), risk (error rates, bias indicators), customer impact (complaints, churn, NPS), and operational resilience. Naturally, you'll want to know your AI program's return on investment as with any other investment. The Board needs to define a clear reporting structure and rhythm. What do they expect to hear and when from the executive management team?

How are we managing data confidentiality when staff use generative AI tools?

Why this matters:

Generative AI changes the data risk profile: employees can paste sensitive data into tools, and outputs can inadvertently include confidential information. The Board should ensure there are practical rules (enterprise-grade tooling, access controls, training, and monitoring) rather than ineffective blanket bans that push usage underground.

About the author



For thirteen years, until 2024, Simon Poole was an Operating Partner with the PE fund Helios Investment Partners where he worked with portfolio companies to develop and deliver value creation plans. In this role he served on the Boards of many portfolio companies including three that then listed on the London Stock Exchange. Prior to Helios he was a CFO, mainly working in telecoms. He qualified as a Chartered Accountant with PW. Since retiring from Helios, Simon has completed a Certificate in Machine Learning and AI with Imperial College, London. He now leads Board workshops for Boards that are implementing AI. The workshops are pitched so that they provide a balance of some technical clarity about AI but are accessible to non-technical experts.

Simon is also a current Board Director, serving on three Boards and an active angel investor.

If you would like him to work with your Board on developing your approach to AI please reach out.

simon@groundedthinking.ai