15-MINUTE BRIEFING · NO. 1

AI: What Every Director Needs to Know

AI is on the agenda. What has changed, where it creates exposure, what good oversight looks like, and what to ask at the next board meeting.

What has changed

AI can introduce risks that familiar technology controls do not fully address. Systems that generate content, learn from data, or take actions need oversight suited to their use. Six changes matter to a board.

Fluent answers can be incorrect. Generative AI can produce plausible text without establishing that it is true. Repeating a request can produce a different answer. Neither fluency nor consistency alone establishes accuracy.

They draw on data, and inherit its faults. A model trained on past decisions can reproduce patterns in them, including unfair or inappropriate ones. Its behavior also depends on the information supplied when it is used. Change either, and the output can change even when the surrounding software does not.

Some act as well as advise. Systems called agents can take steps on their own: sending email, updating records, or making payments. Each step may be reasonable. The sequence may exceed what anyone intended or approved.

Providers can change what you rely on. Organizations can use a provider's model directly or through an existing product. Changes to the model, its terms, or its availability can affect a business process even when the organization has not changed its own software.

They can arrive outside procurement. AI can appear in email, document, customer service, and recruitment software through an update or an activated feature. Staff can also use public tools without making a purchase. An inventory based only on purchases may miss those uses.

They operate at scale. A person usually applies a bad judgment one case at a time. A system can apply it thousands of times before anyone notices, and every instance looks like normal output.

These changes make a familiar oversight question more urgent: how would management know that a system is no longer working as intended, and what would happen next?

15-MINUTE BRIEFING · 02

Where AI creates exposure

The harms are mostly familiar. What's new is how quietly they arrive.

Bad decisions, made at scale. A model screening loan applications, job candidates, or insurance claims can be wrong in a consistent direction. By the time the pattern is visible, it has touched thousands of people, and each of them may have a claim.

Confidential data leaving. Staff may paste customer records, contracts, and board papers into public AI tools to save time. Depending on the provider's terms, that material may be stored, read, or used to train a model.

Security compromise. AI systems can be manipulated through the content they read. A carefully worded email or web page can instruct an assistant to reveal data or take an action its owner never intended. Security controls need to address these interactions as well as familiar threats.

Agents doing more than intended. A system permitted to act can act wrongly, quickly and repeatedly. After an incident, the board needs to understand both what happened and who approved the system's authority to act.

Output that sounds right and is wrong. Generative AI can produce invented citations, incorrect figures, or misleading interpretations. Without effective review, people may rely on that output in customer communications, legal work, or business decisions.

Legal and regulatory exposure. Existing obligations concerning discrimination, consumer protection, privacy, and product liability can apply to AI use. Some AI-specific rules also reach across borders. Using a vendor does not remove the need to assess the organization's own obligations.

Dependence on a vendor. When a critical process runs on a provider's model, that provider's prices, outages, and policy changes become your operational risk. Some organizations find they cannot switch without rebuilding the process.

Reputational harm. An AI failure can quickly become a question about the board's oversight. Be ready to explain what management knew, what controls were in place, and how the organization responded.

15-MINUTE BRIEFING · 03

What good oversight looks like

Overseeing AI calls for what a board expects of any material risk, applied to systems whose behavior can change without an obvious outage. Management should be able to show eight things.

A complete list. How many AI systems are in use, including those built into other products, and how that count was checked.

A named owner for each. One person, senior enough to act, answers for what each system does. A committee is a forum. An owner is a person.

A judgment of what could go wrong. Each system is rated by the harm it could cause, and the scrutiny it receives follows that rating. A tool that drafts internal notes and one that declines customers deserve different treatment.

Clear human authority. It's written down what each system may decide on its own, what needs a person's approval, and who that person is.

Independent testing. Someone other than the people who built or bought the system has checked that it works as claimed, including for the people it could treat unfairly, and reported what they found.

Security designed for AI. Each system can reach only what it needs, and has been tested against attempts to manipulate it.

Ongoing monitoring. Performance is measured in operation as well as at launch, because AI systems drift as the world and their data change.

A plan for when it fails. Someone has the authority to switch the system off, at any hour, and the business knows how to carry on without it.

Descriptions establish intent. Ask for evidence that these controls work, and follow up where the evidence is missing or incomplete.

EIGHT QUESTIONS, IN ORDER

What to ask at your next board meeting

Eight questions, in order. Each answer depends on the one before it: there is little point asking who owns a system until you know it exists. Treat a confident answer as the start of the conversation, and ask how it is known.

01 VISIBILITY
How many AI systems are in production in this organization, who owns each one, and when was that list last checked against something other than itself?

02 ACCOUNTABILITY
For each material AI system, who is the single person answerable for its outcomes, and what has this board delegated as against retained?

03 VALUE
What did we say this would deliver, and what has it actually delivered since?

04 EXPOSURE
For this system, what is the worst realistic outcome, by what sequence would it occur, and what stands in the way?

05 FOUNDATIONS
Where did the data come from, were we entitled to use it for this, and whose model is underneath?

06 OBLIGATIONS
Which obligations attach to our use of AI, who confirmed that list, and when was it last reviewed?

07 ASSURANCE
Who, independent of the people who built it, has tested that this works as claimed, and what did they find?

08 RESILIENCE
If this system began producing wrong answers tomorrow, how would we find out, how long would that take, and who can switch it off?

Explore the free Board Questions preview, including the complete Visibility example. The Board Edition includes the follow-up questions and answers to examine for all eight domains. The framework sets out the requirements behind those questions.

These questions follow the eight domains of the framework. Read more at theassuranceline.com/framework.