CLEMENZA LAW GROUP

NEW YORK, NY · (212) 365-4875 · CLEMENZALAW.COM

THE PRIVATE BRIEF · AI & IP

Board Governance in the Age of AI: Five Questions Every Director Should Ask

ANTHONY CLEMENZA · MANAGING PARTNER

Updated September 2026

IN BRIEF

- A board's legal duty is oversight: making a reasonable effort to know about the significant risks the company faces and to see that someone is managing them. AI is now one of those risks for almost every company, whether or not it sells anything technical.
- Directors do not have to understand the technology. They have to ask where the company uses it, what information goes into it, what has been promised to customers about it, who is accountable, and what the plan is if it goes wrong.
- Asking the questions and recording the answers in the minutes is most of the job. A board that can show it asked is protected in a way a board that assumed management had it covered is not.

Most directors we talk to have the same private worry about artificial intelligence: they do not understand it well enough to oversee it. That worry is understandable, and it is also mistaken about what oversight requires. A board is not asked to understand how a technology works. It is asked to know where the company is exposed, to satisfy itself that someone competent is managing that exposure, and to write down that it did so.

That is the standard for every other risk a board oversees, from cybersecurity to product safety, and it is the standard for this one. This piece gives you the five questions that meet it. They can be asked in a single meeting. The answers, and the fact that you asked, go in the minutes.

Why this is a board matter at all

Directors of a New York corporation, and of a New York nonprofit, owe a duty of care: to act with the diligence an ordinarily prudent person in their position would use. Delaware's courts, in a line of cases that began with a 1996 decision called Caremark, have held that directors must also make a good faith effort to oversee the company's operations, and a careful New York board holds itself to that standard too. A board cannot sit back and assume management is handling the significant risks. It has to make a reasonable effort to be informed and to see that reporting systems exist.

AI has become a significant risk for nearly every organization, not because every organization builds AI, but because nearly every organization now uses it, often without a decision having been made anywhere. Employees use it to write, analyze, summarize, and decide. Vendors have built it into products the company already licenses. Customers ask about it. Regulators have started to enforce against companies that overstated what their AI could do. New York City, for one, requires employers using automated tools in hiring to have them independently audited and to tell candidates. A board that has never discussed any of this is not, in the legal sense, overseeing it.

A board is not asked to understand how a technology works. It is asked to know where the company is exposed, to satisfy itself that someone competent is managing that exposure, and to write down that it did so.

Question one: where are we using it?

Start with an inventory. Ask management to list every place AI is currently in use: the tools employees use, the features vendors have turned on inside existing software, and anything the company has built or bought that makes or shapes decisions (pricing, hiring, credit, customer service, content). Most boards are surprised by the length of the list, and by how much of it was never approved by anyone.

You are not looking for a technical document. A one-page table with three columns (what, where, who owns it) is what a board needs.

Question two: what goes in, and what comes out?

Every AI tool takes information in and produces something out, and each direction carries risk. On the way in: confidential company information, customer data, personal information protected by privacy law, material the company licensed but does not own. On the way out: text or images that may not be protected by copyright, decisions that may be biased, and answers that may simply be wrong.

Ask: what kinds of information are our people allowed to put into these tools, and is that written down anywhere? Do our vendor agreements say our data will not be used to train their models? Where AI output is used in something that matters, does a person review it first? If the answers are “we are not sure,” that is a finding, and it is a fixable one.

Question three: what have we told people?

Companies talk about AI in three directions: to customers (in contracts and marketing), to investors (in decks and disclosures), and to employees and job candidates. Each is a place where a statement can be untrue.

Regulators have begun bringing cases for what they call AI washing: claiming a product uses AI when it does not, or claiming capabilities it does not have. Customer contracts may contain warranties, written years ago, that no longer describe how the work is produced. Job candidates in some places are entitled to notice when an automated tool screens them.

Ask management to show the board what the company currently says about AI, in its marketing, its standard contracts, and its investor materials, and to confirm that each statement is accurate today.

Question four: who is accountable?

Risk that belongs to everyone belongs to no one. The board should be able to name the executive who owns AI risk across the company, and that person should be able to describe the policy, the approved tools, the training people receive, and how new uses get evaluated before they are adopted. It does not need to be a chief AI officer. It needs to be a named human being with authority.

Ask, too, how this reaches the board. Once a year is usually the right cadence for a report on the inventory, the incidents, and the changes. More often if the company builds AI into its products.

Question five: what happens when it goes wrong?

At some point, something will. A tool will leak confidential information, an automated decision will be challenged, an employee will publish something the tool invented, a vendor will change its terms. The question for the board is whether the company will find out quickly and know what to do.

Ask whether the incident-response plan covers AI failures. Ask whether the insurance does; cyber and errors-and-omissions policies vary on this, so ask to see what the policy language actually says. Ask who would make the call to pause a tool, and whether they know they have that authority.

Risk that belongs to everyone belongs to no one.

What to do with the answers

Record them. The minutes should reflect that the board discussed AI use and risk, what management reported, what the board asked, and what follow-up was agreed. Then set a date to return to it. This is not about generating paper for its own sake. When something does go wrong, the difference between a board that can show a record of informed attention and a board that cannot is the difference between a difficult week and a lawsuit against the directors personally.

Where to start

Put the five questions on the agenda for the next meeting and send them to management in advance, with the request for a one-page inventory. That single step moves the company from unmanaged to managed, and it takes an hour of the board's time.

This is general information about board oversight, not legal advice about your company or your duties as a director.

WRITTEN BY

Anthony Clemenza

MANAGING PARTNER

Anthony is the Managing Partner of Clemenza Law Group and serves as general counsel to a wide range of companies and philanthropic organizations. Nearly every essay here began as a real question from a real client: the contract that landed at 4 p.m., the board that meets on Thursday. He writes them the way he would answer you across the table, plainly and with care.

CLEMENZA LAW GROUP

NEW YORK, NY · (212) 365-4875 · CLEMENZALAW.COM

The current version of this essay is kept at clemenzalaw.com/insights/board-governance-ai-questions