What Boards Are Being Asked About AI — and the Evidence Management Should Be Ready to Provide

An independent review of AI governance guidance from standards bodies, regulators, institutional investors, governance organizations, and board advisors.

About the Analysis

After completing my review of how public companies describe AI strategy, governance, and risk in investor materials, SEC filings, and governance disclosures, I was left with a practical question:

If Boards are increasingly expected to oversee AI risk, what exactly are they being told to ask management, and what evidence should management be prepared to provide?

To answer that question, I reviewed guidance from standards bodies (NIST, ISO, COSO), regulatory developments, board-governance organizations, institutional investors, legal advisors, and practitioner communities. While the sources approached AI from different perspectives, several consistent themes emerged.

Research Question

What should Boards be asking management about AI, and how should management demonstrate that governance is working?

The Observation That Surprised Me

The gap between AI adoption and AI oversight appears to be widening.

According to the Conference Board's 2026 review of S&P 500 governance disclosures, the percentage of companies identifying AI as a material risk increased dramatically over the last several years. Yet AI-specific expertise at the Board level remains relatively uncommon, and fewer than one-quarter of governance leaders surveyed described their Boards as highly fluent in AI.

The implication is not that Boards are failing. The implication is that the expectations placed on Boards are growing faster than the mechanisms used to support them.

Three Observations

One — Effective AI oversight is inherently cross-functional.

No credible source suggested that AI governance belongs exclusively within a single Board committee. Audit, Risk, Technology, Governance, Quality, Legal, Privacy, Security, Data, and Business leaders all contribute different perspectives.

The question is not who owns AI.

The question is whether the governance structure enables the integration AI requires.

Two — AI literacy is becoming an ongoing governance requirement.

Across multiple sources, AI literacy was treated less like a training event and more like a continuing governance responsibility.

Unlike many traditional risks, AI capabilities, regulations, business models, and market expectations continue to evolve rapidly. Boards should expect ongoing education, periodic reassessment, and regular discussion.

Three — Boards increasingly need evidence, not assurances.

The strongest sources consistently moved beyond governance structures and principles toward a more practical question:

How does management demonstrate that governance is operating effectively?

This includes visibility into AI inventories, data fitness, model oversight, third-party dependencies, incident management, pilot governance, regulatory obligations, and organizational readiness.

What surprised me was not how different the sources were. It was how often they converged on the same themes: accountability, AI literacy, data fitness, third-party risk, governance effectiveness, and management's ability to demonstrate that controls are operating as intended.

Questions for Management and Boards

The following questions appeared repeatedly across the sources reviewed and illustrate the types of discussions Boards should expect to have with management.

• What are the highest-risk AI systems operating within the organization, and who is accountable for them?

• How does management demonstrate that the data supporting AI initiatives is reliable, fit for purpose, and governed appropriately?

• How are AI literacy, change management, and responsible use being reinforced across the workforce?

• How do we know our AI initiatives are creating business value without creating unmanaged risk?

• How dependent is the organization on third-party AI providers, and what risks accompany that dependence?

A Final Observation

Most organizations do not struggle to define AI governance principles. They struggle to operationalize them.

The questions in this analysis are ultimately less about AI and more about whether the organization can consistently produce the evidence needed to support effective oversight, decision-making, and accountability.

That evidence does not come from a single function. It is produced across data governance, privacy, security, legal, compliance, technology, and the business itself.

As AI adoption accelerates, the challenge for many organizations will not be identifying the right questions. It will be building the capabilities needed to answer them consistently and with confidence.