Five Questions Boards Should Ask On AI

By Victor P. Fetter III

As fast as boards have learned about artificial intelligence, the AI tools and capabilities your company is putting to work are evolving even faster. Boards (as well as top executives) are increasingly surprised by both the depth and dangers of AI adoption company wide. Here are the questions your board needs to ask to get ahead of this stampede.

When a tech executive who leads AI deployment also serves on a public company board and audit committee, two fundamentally different conversations about artificial intelligence collide. On the operating side, the questions are granular. Which processes are changing, how fast, and what breaks when the model is wrong?

In the boardroom, the conversation carries more gravity, but often less precision. The board asks whether management has a strategy for AI. It rarely asks whether anyone has cataloged the decisions that AI is *already* making.

This is a collision between building the technology and governing it. Having led teams that deploy production AI models, and having served on an audit committee reviewing their outputs, one lesson stands out. Operational visibility reveals what opacity hides. In one instance, a governance review uncovered an AI model producing confident outputs while optimizing for the wrong objective entirely.

That insight only surfaced because the governance architecture required the review. Without that architecture, the same model produces outputs that look authoritative, but erode trust in ways that only become visible after the damage is done.

Board oversight of AI tripled from 16 percent to 48 percent of Fortune 100 companies between 2024 and 2025. Latham and Watkins reported in February 2026 that nearly half of Fortune 100 boards have formally delegated AI oversight to a specific board committee, and that nominating and governance committees are integrating AI fluency into director succession criteria. The SEC's 2026 examination priorities name AI compliance as a focus area, with examiners evaluating whether firms have implemented adequate policies to monitor and supervise its use.

Yet a National Association of Corporate Directors' survey found that, while 62 percent of boards set aside agenda time for AI, only 23 percent have assessed how AI disruption might affect their strategy. Only 27 percent have formally incorporated AI governance into their committee charters. Setting aside time is not the same as building governance architecture. The attention is there, and the architecture is catching up. Yet cataloging AI tools is not the same as understanding how AI-driven decisions affect business strategy and operational controls. That distinction will define the next phase of board-level AI governance.

"AI will transform our industry" is language that makes executives feel savvy and boards feel reassured without specifying what has actually changed.

Management presents AI initiatives in one of two languages. The first is deeply technical. Model architectures, training data volumes and tuning methodologies that few outside the engineering team can evaluate. The second is aspirational. "AI will transform our industry" is language that makes executives feel savvy and boards feel reassured without specifying what has actually changed.

The operational conversation sits between these two languages, and it is the one that each board committee should be conducting. What processes changed? Who approved the change? What happens when the

Victor P. Fetter III *is a technology executive and NACD-certified public company director. [linkedin.com/in/fetter]*

AI output is wrong? Who bears accountability when the AI makes a decision that exposes the company to financial, regulatory, or reputational risk?

Pace compounds the problem. AI systems move from pilot to production faster than oversight structures adapt. Boards that were deliberate and methodical about prior digital transformation initiatives must now govern capabilities that were deployed before the governance conversation began.

PwC's 2024 survey found that 57 percent of directors assign AI oversight to the full board. That instinct is correct. AI governance requires full board engagement at the strategic level. Only 17 percent point to the audit committee as playing a defined role.

That gap is an opportunity. Full board oversight sets direction; committees then execute the operational detail within their mandates. Audit committees carry fiduciary responsibility for internal controls over financial reporting, enterprise risk architecture, and asset governance. When AI systems influence revenue forecasting, generate investor materials, or drive material decisions, they fall within that mandate by definition.

The NIST AI Risk Management Framework provides the operational structure. Over the next twelve months, five questions will define whether a board's AI governance has moved from awareness to architecture. These are not questions designed to slow deployment. They are the questions that make deployment possible at scale.

> ***An AI model enters production with a defined scope. Over six months, it expands, without revisiting the original mandate. Nobody authorized the current scope.***

☐ Where is AI making decisions?

Most organizations adopt AI one tool at a time without recognizing the aggregate governance implications. A department finds a workflow accelerated by machine learning. A team discovers that AI drafts contract language faster than paralegals. A supply chain group implements an optimization algorithm. Each adoption makes sense locally. Aggregated across the enterprise, however, they represent something that was never explicitly authorized. Institutional judgment has been incrementally delegated to machines.

Every other category of enterprise technology has an asset management discipline. Servers have utilization metrics. Databases have performance baselines. Software licenses have compliance audits. Yet AI models have no equivalent discipline in most companies.

The delegation problem compounds quietly. An AI model enters production with a defined scope. Over six months, operational teams expand its inputs, extend its reach into adjacent workflows, and increase its decision authority without revisiting the original mandate. Nobody authorized the system's current scope. It accumulated through dozens of incremental choices, each individually defensible, but collectively an unauthorized transfer of institutional judgment.

Boards will increasingly expect a *decision* inventory rather than a *tool* inventory. That inventory does not need to be comprehensive on day one. The practical starting point is decisions that touch financial reporting, regulatory compliance, or client-facing output. Undetected errors here compound fastest.

For each significant decision made or influenced by AI in those domains, the board needs to know what decision is being made, who approved the delegation, and who is responsible for detecting and correcting errors before they compound. A useful diagnostic is to ask management which decisions made by AI last quarter would have required a human signature twelve months ago.

☐ How are we measuring data quality?

When a model trained on incomplete or biased data makes thousands of operational decisions per day, data quality problems compound into systemic risk faster than any human process would allow.

This is where the audit committee's mandate becomes urgent. What data assets does the company possess, how are these curated, and is data quality measured with the same rigor applied to other operational metrics? Are there baseline accuracy metrics for training data? Is there a process for identifying data drift (the gradual degradation that occurs as the operating environment changes but the training data does not)?

Data provenance adds a second layer. Every model consumes data with a history—where it originated, who curated it, when it was last validated, and whether it contains proprietary information that could flow into third-party platforms. Companies that cannot trace the lineage of their training data cannot assess the risk those models carry. The audit committee should expect the same chain-of-custody discipline for AI training data that it already expects for data underlying financial statements.

Consider this test. Ask management to identify a single data quality failure that would erode model accuracy by 10 percent without triggering a correction. If management cannot name it, the detection infrastructure is not yet mature.

Boards learned a decade ago with cybersecurity that a breach does not come through your front door. It comes through your vendor's vendor. The same principle now applies to AI.

☐ How are we managing third-party AI risk?

This may be the most consequential of the five questions, and it is the one arriving fastest.

AI systems increasingly operate autonomously, connecting to external data sources, vendor platforms, and third-party tools. An AI agent does not simply process data inside company walls. It executes transactions with third-party platforms and integrates outputs from models your firm did not build and does not control.

Boards learned a decade ago through cybersecurity oversight that a breach does not come through your front door. It comes through your vendor's vendor. The same principle now applies to AI. The risk lives in the third-party data the model consumes, the external systems the AI agent accesses, and decisions that depend on outputs you cannot independently verify.

The question for the board is whether management can map the AI supply chain. If management cannot produce that map, the third-party AI risk is unquantified. The dependency runs deeper than data. When a vendor updates a model you rely on, performance can degrade without notice. Switching to an alternative may require rebuilding the integration architecture from scratch.

☐ What happens when AI fails?

Third-party risk compounds when failures are invisible. The defining question for AI performance is not whether the system will fail. It is whether staff would know. How would it detect if its AI systems were systematically producing wrong or degraded outputs?

Context determines the threshold. A hallucination rate that is tolerable in a marketing content tool is unacceptable in a system generating financial projections or medical recommendations. The board does not need to set the thresholds. Rather, it needs to confirm that thresholds exist, that someone is monitoring them, that there is a defined response when they are breached. Are the thresholds recalibrated as operational experience reveals what the initial assumptions missed?

The financial reporting dimension deserves particular attention. When AI models influence revenue forecasting, generate investor relations materials, or shape management guidance, they become part of the financial reporting ecosystem. Under Sarbanes-Oxley (SOX) and the internal control framework, the integrity of those outputs is not merely a technology question. It is a controls question.

If an AI system shapes a number that appears in an earnings release or SEC filing, the audit committee needs to understand whether that system falls within the scope of internal controls over financial reporting, and whether the external auditor has assessed the control environment. Most audit firms are building AI competency, but few have developed standardized protocols for testing AI-driven financial data as part of the control environment.

The audit committee that raises this question before the auditor does will shape the terms of the conversation. The one that waits will discover the gap during a restatement or a regulatory inquiry.

No Suprises

AI Questions Your Board Should Ask

As AI systems move deeper into enterprise decision-making, it will become increasingly important for boards to ask the following questions with specificity and operational rigor. These are the questions that will shape how boards govern AI at scale.

☐ ***1. Defining and mapping AI.***
Has the full board assessed which strategic decisions are now influenced by AI, and does the governance architecture match the scale of that delegation? Can management produce a decision inventory rather than a tool inventory?

☐ ***2. Data quality.***
Has the audit committee confirmed that data feeding AI models are curated with the same rigor applied to other operational metrics? Are there baseline accuracy standards and a process for detecting data drift before it compounds into systemic risk?

☐ ***3. Third-party AI risk.***
Can management map the AI supply chain to its first material third-party dependency, specifying what data flows outward and what protections govern that exposure? The audit committee and the nominating and governance committee share oversight here, as vendor risk spans both operational controls and board-level competency.

☐ ***4. Performance and failure.***
Has the audit committee confirmed that AI systems that affect financial reporting fall within the scope of internal controls? Are failure thresholds and response protocols defined before the first error compounds? The compensation committee should evaluate whether incentive structures reward responsible governance rather than speed of adoption alone.

☐ ***5. Talent and accountability.***
Can management identify a named individual with the authority to halt an AI deployment on governance grounds? Can any committee point to a specific instance where that authority was exercised? The nominating and governance committee should incorporate AI fluency into succession criteria for both the board and senior leadership.

Too often, deployment architectures are presented with impressive sophistication, while the failure architecture amounts to "we will handle it." Boards increasingly expect both, side by side. The principle extends beyond AI. Any domain where the company faces a crisis faster than humans can deliberate requires pre-built architecture, not improvised response.

☐ Who has the authority to challenge AI?

Latham and Watkins report that nominating and governance committees are integrating AI fluency into director succession criteria for the board itself. The same question applies throughout the enterprise, and with greater operational urgency.

Is there a named person with operational accountability for AI risk, the way someone is explicitly responsible for information security? More importantly, has this staffer demonstrated the ability to challenge a technical recommendation on business grounds?

The distinction matters. Many companies have appointed AI leads who can explain what the technology does. Fewer have empowered someone who can tell the engineering team "this model is not ready for production," and have the authority to make that judgment stick.

> ***Ask management to identify the person who has explicit authority to delay an AI deployment. Then ask when that authority was last exercised.***

Russell Reynolds found that only 25 percent of boards possess complementary technology depth —someone who pairs technical knowledge with operational P&L experience. Private company boards face the same challenge. The Private Directors Association has found that many lack the technical fluency to evaluate AI recommendations independently.

This gap often exists across the C-suite. AI oversight may rest with a chief data officer who lacks business context, or a business leader who lacks technical depth. When governance is distributed across multiple functions without that combined capability, do you have true coverage, or merely the appearance of it?

Ask management to identify the person whose explicit mandate includes the authority to delay an AI deployment on governance grounds, independent of the technical team's recommendation. Then ask when that authority was last exercised. If no one holds the mandate, or if it has never been tested, the challenge function does not really exist in practice.

The incentive structure reinforces this point. If compensation rewards speed of AI adoption without weighting governance quality, the company is paying people to bypass the very controls the board is trying to build.

If your board found gaps in these five questions, the response is not to slow down. Rather, it is to build the governance architecture that makes speed sustainable.

Most companies are still early in answering these questions with the specificity that governance demands. Tools have been deployed faster than decision skills have been built. Production AI systems have outpaced failure protocols. Models have been embedded without systematic tracking of whether they still perform as they did at deployment. AI governance has been assigned to boards still developing operational mandates and technical fluency. Finally, none of these governance structures has yet been validated by actually catching something before it caused harm.

What will separate boards that are ready for this next phase from those still finding their footing? Operational discipline. Not strategy decks. Not workforce planning presentations. Not director education sessions. The specific, uncomfortable, operational questions that distinguish governance with substance from governance that exists only in the proxy statement.

Beneath all five questions sits a more fundamental obligation. The board must understand what technologies are reshaping the operating environment, which third-party solutions the company depends on, and how AI is embedded in its products and services.

That obligation extends beyond the boardroom. The final discipline that validates architecture is external—what the board discloses to shareholders about AI oversight. AI shareholder proposals more than doubled between 2023 and 2024. Glass Lewis now expects companies using or developing AI to provide clear disclosure of the board's oversight role. ISS reported that S&P 500 disclosure of board AI oversight reached 31.6 percent in 2024, an 84 percent increase from the prior year.

The proxy statement is where governance architecture becomes visible to the market. Boards that can clearly articulate which committee holds AI oversight, what decision inventory exists, and how failure thresholds are monitored will find that the discipline of disclosure strengthens governance itself. Where that clarity is lacking, expectations from investors and other stakeholders will likely bring increased scrutiny over time.

The boards that build this governance muscle now will find something counterintuitive. Well-governed AI moves faster, not slower. When data provenance is already verified, deployment teams stop relitigating inputs. When failure thresholds are predefined, production decisions that once required weeks of internal debate resolve in days.

I have seen this pattern repeat across organizations. In one case, management pre-validated its data pipeline and defined model failure thresholds before deployment. Production approval that had historically taken six weeks of committee review cleared in nine days because every stakeholder question had already been answered by the governance architecture. The governance structure becomes the accelerant because the company spends its energy scaling models that work rather than remediating ones that never should have reached production.

Boards that build this architecture now will set the terms for how AI governance matures across their industries. The difference will not be whether the board paid attention to AI. It will be whether the board asked the right questions while there was still time to shape the answers. ■

THE CORPORATE BOARD
4701 SW Admiral Way #4, Seattle, WA 98116
(800) 884-0338 ■ *www.CorporateBoard.com*