Governance Alerts (Short-Form Analysis)

Alert 24-02

New Fiduciary Implications for AI-Driven Decision Models

Category: Corporate Governance | Audience: Boards, Audit Committees, Executive Management

Purpose: To identify why AI-enabled decision models now require board-visible governance, controls, and evidence.

Board-Level Alert

AI-driven decision models are no longer experimental tools sitting inside information technology. They are becoming operating infrastructure for credit, pricing, hiring, compliance, financial forecasting, customer segmentation, claims handling, fraud monitoring, legal review, and executive decision support. That change creates a fiduciary issue: directors and senior managers may now need evidence that consequential AI systems are inventoried, risk-ranked, tested, monitored, and escalated when model behavior creates legal, financial, operational, or reputational exposure.[1-5]

Why This Alert Matters Now

For boards, the central question is not whether artificial intelligence is valuable. It is whether the company is using AI in ways that create consequential decisions without an equally mature governance architecture. Decision models can influence who receives employment opportunities, credit, insurance, pricing, service priority, compliance clearance, fraud suspicion, medical or benefit triage, capital allocation, or legal-risk recommendations. When those models are wrong, biased, insecure, undocumented, or overstated in investor-facing or customer-facing communications, the resulting exposure is not merely technical. It can become a disclosure, discrimination, consumer-protection, internal-control, privacy, cybersecurity, contract, employment, or fiduciary oversight problem.[3-10]

The regulatory direction is clear even though U.S. federal AI legislation remains fragmented. The SEC has already brought AI-washing cases for false and misleading AI claims.[6] The FTC has pursued enforcement against deceptive AI claims and AI-enabled unfair or deceptive practices.[7] The EU AI Act has entered into force and applies in phases, including prohibited practices, AI literacy, general-purpose AI obligations, and broader high-risk system requirements.[8] Colorado has adopted comprehensive high-risk AI requirements, now extended to June 30, 2026, and

California employment regulations make clear that automated-decision systems may violate state civil-rights law when they produce discriminatory employment effects.[9,10]

The board implication is straightforward: an AI system that materially affects customers, employees, financial reporting, compliance judgments, safety, cybersecurity, or public statements can no longer be treated as a back-office tool. It should be treated as a governed decision system.

The Fiduciary Shift: From Technology Adoption to Oversight Evidence

Traditional technology oversight often asks whether a system was purchased, implemented, secured, and budgeted. AI oversight requires a different evidentiary posture because model outputs may be probabilistic, dynamic, vendor-dependent, data-sensitive, and difficult to explain. A board does not need to validate every algorithm. It does need to know that management has implemented a reliable system for identifying AI use, assigning ownership, testing performance, documenting limitations, controlling vendor risk, monitoring drift, and escalating material failures.

Delaware oversight doctrine is relevant because AI can become a mission-critical compliance and reporting risk. Stone v Ritter confirmed that oversight liability is tied to a failure to make a good-faith effort to implement and monitor reporting systems.[1] Marchand v Barnhill reinforced that boards must pay particular attention to mission-critical risks and board-level reporting systems.[2] In the AI context, the fiduciary question is not whether a model made a bad recommendation. It is whether the board and management had a good-faith oversight system capable of surfacing model risks before they became institutional failures.

The practical fiduciary standard is not perfection. It is disciplined visibility. Directors should be able to ask management for an AI inventory, a risk classification, a model owner, a vendor owner, a validation record, a human-review protocol, an incident/escalation pathway, and a board-level dashboard for material AI use cases. If those materials do not exist, the company may be making consequential AI decisions without governance evidence.

What Makes AI-Driven Decisions Different

AI-driven models differ from ordinary software in several ways that matter to boards. First, they can produce outputs that vary with prompts, context, data drift, model updates, or vendor changes. Second, they may rely on training data or proxy variables that embed bias or create unexpected disparate impacts. Third, management may not fully understand the model because it was acquired from a vendor, embedded in a cloud platform, or adopted informally by employees. Fourth, the company may make public claims about AI capabilities that exceed what the system actually does. Fifth, AI systems can generate plausible but false information, leak confidential data, or be manipulated through adversarial inputs. NIST identifies AI risk management as a continuous lifecycle activity organized around governance, mapping, measurement, and management; its generative-AI profile adds risks such as confabulation, data privacy,

information integrity, information security, intellectual property, and value-chain/component integration.[3,4]

These characteristics produce a governance problem: AI risk is distributed across the enterprise, while accountability often remains undefined. Human resources may deploy hiring analytics. Finance may use forecasting models. Legal may use AI-assisted contract or discovery tools. Marketing may use generative AI for claims and content. Operations may use AI for supply chain or quality triage. Cybersecurity may rely on AI detection tools. Each function may believe the model is someone else's responsibility. That is the oversight gap.

Current Regulatory Signals Boards Should Not Ignore

The SEC's AI-washing actions are especially important for public companies and private companies preparing for capital formation, acquisition, or debt financing. In March 2024, the SEC settled charges against two investment advisers for allegedly making false and misleading statements about their purported use of AI and machine learning; the firms agreed to pay a total of $400,000 in civil penalties.[6] The governance lesson is broader than investment-adviser regulation: AI claims must be substantiated. A company should not say that AI improves accuracy, reduces risk, personalizes service, predicts outcomes, or automates expert judgment unless those claims are supported by documented evidence.

The FTC's Operation AI Comply reinforces a similar point for consumer-facing businesses: AI does not excuse deception, unfair practices, or unsubstantiated performance claims.[7] If AI is used to sell, screen, advise, rank, recommend, influence, or deny access, management should be prepared to explain what the model does, what it does not do, how it was tested, and how consumers or employees can obtain appropriate human review where required or prudent.

State and international regimes are moving toward documentation, risk management, notices, impact assessments, and human accountability. Colorado's high-risk AI law requires reasonable care to protect consumers from known or reasonably foreseeable risks of algorithmic discrimination and requires deployers of high-risk AI systems to implement risk management policies and programs, conduct impact assessments, provide disclosures in certain consequential-decision contexts, and maintain evidence.[9] California's employment regulations emphasize that automated-decision systems can violate employment discrimination law when they adversely affect applicants or employees based on protected characteristics, and they require retention of employment records, including automated-decision data, for at least four years.[10]

Board-Level Red Flags

Red Flag	Why It Matters	Board Response
No enterprise AI inventory	Management cannot govern systems it has not identified.	Require a current inventory of AI systems, owners, vendors, data sources, and business purposes.
No risk tiering	Low-risk productivity tools and consequential decision systems require different controls.	Classify models by legal, financial, employment, consumer, safety, cybersecurity, and disclosure impact.
Vendor black box	A vendor's model can create the company's liability without giving the company adequate evidence.	Require vendor documentation, audit rights, data-use terms, performance metrics, security controls, and change notices.
Unsupported AI claims	Marketing, investor, lender, or customer claims can become misleading disclosures.	Require substantiation files for material AI claims and legal review before publication.
No human-review path	Automated decisions can create fairness, due-process, and customer-trust failures.	Require human review for consequential adverse decisions and escalation for contested outcomes.
No model monitoring	Models can drift after launch as data, users, products, or vendor systems change.	Require monitoring for performance, bias, security events, data leakage, and unexpected outputs.

Minimum Governance Architecture

A practical board approach is to require management to build an AI control architecture around five elements.

Inventory: Identify AI systems in use, including vendor tools, embedded platform features, internally developed models, generative AI tools, and informal departmental deployments.

Risk classification: Distinguish ordinary productivity tools from systems that affect consequential decisions, financial reporting, compliance judgments, employment, consumer access, safety, privacy, or public statements.

Control ownership: Assign an executive owner, a business owner, a technical owner, and a legal/compliance owner for each material AI system.

Validation and monitoring: Require documented pre-use testing, performance thresholds, bias/fairness review where relevant, cybersecurity assessment, data-protection review, model-change controls, drift monitoring, and post-implementation review.

Escalation and reporting: Define what must be reported to management, the audit/risk committee, or the full board, including material model failures, adverse regulatory correspondence, AI-related complaints, privacy incidents, unsupported claims, and high-risk vendor changes.

COSO's 2026 guidance on effective internal control over generative AI is important because it frames GenAI governance as an internal-control issue, not merely a technology-policy issue.[5] That framing should resonate with boards: if AI affects operations, reporting, or compliance, it should be auditable. Evidence should exist before the company relies on the model, not after a regulator, plaintiff, lender, auditor, or board committee asks for it.

Questions Directors Should Ask at the Next Meeting

Where are AI-driven decision models currently being used, including vendor tools and embedded software features?

Which AI systems influence consequential decisions involving employment, customers, credit, pricing, benefits, compliance, fraud, legal review, financial forecasting, or safety?

Do we have documentation showing the purpose, data sources, limitations, validation results, monitoring controls, and responsible owners for each material system?

What public, investor, lender, customer, or employee claims have we made about AI capabilities, and can we substantiate them?

What AI-related incidents, complaints, overrides, false positives, bias concerns, data leaks, or vendor changes have occurred in the last two quarters?

Which board committee owns AI oversight, and what metrics will appear on the recurring dashboard?

Recommended 90-Day Response

In the next 30 days, management should identify all AI systems in material use and freeze any new consequential AI deployment until ownership and risk classification are documented. In days 31 through 60, management should classify systems by risk, obtain vendor documentation,

review data inputs, review public AI claims, and identify systems requiring human review, impact assessment, or heightened monitoring. In days 61 through 90, management should deliver a board-level AI governance dashboard covering inventory, risk tiering, owners, testing status, unresolved gaps, incidents, regulatory exposure, vendor risk, and next-quarter remediation.

This response does not slow innovation. It protects it. The companies most likely to scale AI successfully will be those that can prove what their systems do, where they are used, who is accountable, what evidence supports reliance, and when leadership is notified that the model has exceeded its governance boundary.

Bottom Line

AI-driven decision models are becoming part of the institutional nervous system of the modern enterprise. Boards should not manage the models, but they must require architecture: inventory, ownership, validation, monitoring, disclosure discipline, vendor controls, and escalation. The fiduciary risk is not that AI will occasionally be wrong. The fiduciary risk is that the company will rely on AI without a board-visible system for knowing when, where, how, and why it is making decisions that matter.

For boards and business managers, the immediate message is simple: AI governance is now corporate governance. If a model can affect legal rights, financial outcomes, regulatory compliance, public statements, employee opportunities, customer access, or enterprise risk, it belongs within the company's fiduciary oversight architecture.

Sources

1. Stone v Ritter, 911 A2d 362 (Del 2006). Supreme Court of Delaware. Accessed April 28, 2026. https://courts.delaware.gov/opinions/download.aspx?id=84060

2. Marchand v Barnhill, 212 A3d 805 (Del 2019). Supreme Court of Delaware. Accessed April 28, 2026. https://courts.delaware.gov/Opinions/Download.aspx?id=291200

3. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. Published January 2023. Accessed April 28, 2026. https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf

4. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1. Published July 2024. Accessed April 28, 2026. https://doi.org/10.6028/NIST.AI.600-1

5. Committee of Sponsoring Organizations of the Treadway Commission. Achieving Effective Internal Control Over Generative AI. Published 2026. Accessed April 28, 2026. https://www.coso.org/generative-ai

6. US Securities and Exchange Commission. SEC Charges Two Investment Advisers with Making False and Misleading AI Claims. Press Release 2024-36. Published March 18, 2024. Accessed April 28, 2026. https://www.sec.gov/newsroom/press-releases/2024-36

7. Federal Trade Commission. FTC Announces Crackdown on Deceptive AI Claims and Schemes. Published September 25, 2024. Accessed April 28, 2026. https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-announces-crackdown-deceptive-ai-claims-schemes

8. European Commission. AI Act. Shaping Europe's Digital Future. Accessed April 28, 2026. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

9. Colorado General Assembly. SB25B-004 Increase Transparency for Algorithmic Systems. Approved August 28, 2025; effective November 25, 2025. Accessed April 28, 2026. https://leg.colorado.gov/bills/sb25b-004

10. California Civil Rights Department. Civil Rights Council Secures Approval for Regulations to Protect Against Employment Discrimination Related to Artificial Intelligence. Published June 30, 2025. Accessed April 28, 2026. https://calcivilrights.ca.gov/2025/06/30/civil-rights-council-secures-approval-for-regulations-to-protect-against-employment-discrimination-related-to-artificial-intelligence/

Note: This alert is an informational governance briefing and is not legal advice. Boards should consult counsel for jurisdiction-specific advice.

About the Author:

George (Keoki) Wallace, PhD, LLM, JD, is a multi-disciplinary executive with an extensive background spanning finance, law, and corporate leadership. As a licensed attorney, CPA candidate, CEO, CFO and General Counsel, he specializes in bridging the gap between complex financial data and legal compliance.