A GOVERNANCE PAPER FOR DIRECTORS

AI Oversight as a Fiduciary Obligation

Boards are being asked to oversee a category of risk that arrived inside their organizations without a decision, without a policy, and in most cases without a standing report. This paper sets out why that is a governance exposure rather than a technology problem, what the current evidence shows about where boards actually stand, and what a board should do at its next meeting.

Prepared for boards of private companies, portfolio companies, and nonprofit organizations.

Andrew Brümmer
The Ardunan Village
abrummer@ardunan.com · +1 (404) 915-2027 · ardunan.com
August 2026

Purpose and recommendation

This paper is written for directors of private companies, private equity portfolio companies, and nonprofit organizations. It addresses a question most boards have not formally answered: what is this board's oversight obligation with respect to artificial intelligence, and is the board currently discharging it?

The short answer is that AI oversight is best treated as a mission-critical risk under the same fiduciary framework that governs financial reporting, safety, and cybersecurity. That framework does not require the board to understand model architecture. It requires the board to ensure a reporting system exists, to receive information through it, and to act on what it shows. The evidence set out below suggests most boards have not yet established that system, and a meaningful proportion have approved substantial AI investment without doing so.

> **RECOMMENDATION**
>
> That the board assign AI oversight to a named committee, direct management to produce a decision authority map and accountability framework within 90 days, and establish AI as a standing agenda item with a defined reporting cadence.
>
> Detailed recommended actions, drafted so that each can be minuted, appear at section 5.

1. Why this sits with the board

Under Delaware law, the duty of oversight originates in In re Caremark International Inc. Derivative Litigation, 698 A.2d 959 (Del. Ch. 1996), which held that directors have an obligation to make a good faith effort to ensure a corporate information and reporting system exists, and that a sustained or systematic failure to do so may give rise to liability.

Two later decisions matter more than Caremark itself for present purposes.

STONE V. RITTER, 911 A.2D 362 (DEL. 2006)

The Delaware Supreme Court confirmed that the oversight duty is grounded in the duty of loyalty rather than the duty of care. The distinction is practical rather than academic. Charter provisions authorized by DGCL section 102(b)(7), which most corporations adopt, exculpate directors from personal liability for duty of care violations. They do not exculpate loyalty claims. A director who consciously disregards the need for an oversight system cannot rely on an exculpatory charter provision.

MARCHAND V. BARNHILL, 212 A.3D 805 (DEL. 2019)

The Delaware Supreme Court held that where a risk is mission-critical to the business or to its compliance obligations, the board must establish board-level oversight of that risk and cannot treat management's discretionary reporting as a substitute. Marchand concerned food safety at an ice cream manufacturer following a listeria outbreak. Delaware courts have since entertained oversight claims concerning cybersecurity on the same reasoning.

A third development is frequently cited and should be read together with its outcome. In the McDonald's litigation, the Delaware Court of Chancery recognized in early 2023 that officers, as

agents who report to the board, owe an oversight obligation to identify red flags, report upward, and address those falling within their own responsibility. The reasoning is real, but its precedential weight is limited. It is a trial-level Chancery opinion, it was not tested on appeal, the claims against the director defendants failed, and the claim against the officer was subsequently dismissed on demand futility grounds. It is best treated as an indication of direction rather than settled authority.

Two qualifications matter here, and a board's counsel will raise both. First, oversight claims are among the hardest in corporate law to plead successfully. The standard requires bad faith rather than negligence, and the overwhelming majority of such claims fail. Marchand drew attention precisely because it was exceptional. Second, in a closely held or sponsor-backed company, the party who would ordinarily bring a derivative claim frequently sits on or controls the board, which reduces the practical litigation risk further still.

The argument of this paper therefore does not rest on litigation exposure, and a board should be sceptical of any adviser who tells it that it should. It rests on the observation that the same evidentiary question, namely whether the organization can show that a system existed, that information reached the board, and that the board acted, is asked in transaction diligence, in insurance underwriting, by lenders, and by counterparties. In those settings the question arrives without a pleading standard to protect you, and the answer has immediate commercial consequence.

Whether AI qualifies as mission-critical is company-specific and has not been settled by a court. The relevant question for a board is not whether a case exists on point, but whether the board could demonstrate, on the documentary record, that it established a system, received information through it, and responded. Recent academic commentary makes the same observation: AI does not destabilize the oversight doctrine so much as activate it, and good faith is increasingly demonstrated through governance design, documentation, and escalation structures rather than through the absence of failure.

> **A NOTE FOR PRIVATE AND NONPROFIT BOARDS**
>
> For private stock corporations the duty attaches to the corporation rather than to whether its shares are traded, so private company directors owe substantively the same oversight obligation as their public company counterparts. In practice the exposure arrives through different channels: directors and officers insurance underwriting, lender covenants, customer and partner contract representations, and transaction diligence.
>
> The position for nonprofits is different and weaker. Nonprofit corporations are governed by state nonprofit corporation statutes rather than by Delaware corporate law, enforcement generally rests with the state attorney general rather than with derivative plaintiffs, and a number of states apply a more forgiving standard. Nonprofit directors should read what follows as prudent governance practice and as a response to funder, insurer, and partner expectations, rather than as a statement of legal duty.
>
> This paper is a governance discussion document and is not legal advice.

2. Where boards actually stand

The gap between adoption and oversight is now well documented, and the figures are consistent across independent sources.

ADOPTION IS CLOSE TO UNIVERSAL

Stanford HAI's 2026 AI Index reports that 88 percent of organizations used AI in 2025, with generative AI live in at least one business function at 70 percent.

BOARD-LEVEL OVERSIGHT IS NOT

Institutional Shareholder Services reviewed 3,048 U.S. companies across the Russell 3000 and S&P 500 in January 2026 and found that roughly 8 percent disclosed any board-level AI oversight and 9 percent reported established AI policies. Sixteen percent disclosed a director with specialized AI skills, heavily concentrated in five sectors. The aggregate conceals a wide spread: ISS separately reported around 22 percent of S&P 500 companies disclosing board oversight of AI against roughly 6 percent across the Russell 3000, so the low headline figure is driven by smaller companies.

Two cautions apply to that figure, and this board should apply both. It measures disclosure in proxy statements rather than practice, and the absence of disclosure is not evidence of the absence of oversight. It also says nothing about private companies or nonprofits, which do not file proxies at all. It is offered here as an indication of direction among companies obliged to report, and not as a benchmark against which this board should measure itself.

BOARDS ARE APPROVING SPEND WITHOUT GOVERNANCE

Grant Thornton's 2026 AI Impact Survey found that three in four boards have approved major AI investments, while 48 percent have not set AI governance expectations and 46 percent have not integrated AI risk into ongoing oversight. Seventy-eight percent of business executives surveyed lacked strong confidence that their organization could pass an independent AI governance audit within 90 days.

CADENCE CORRELATES WITH RETURN

Protiviti's 2026 board governance research found that approximately 26 percent of boards discuss AI at every meeting. Among organizations reporting strong returns from AI, roughly 63 percent place AI on every board agenda; among those reporting weak returns, roughly 13 percent do. The direction of causation is not established, but the correlation is strong enough to warrant the board's attention.

DIRECTORS THEMSELVES REPORT THE GAP

Gartner's Board of Directors Survey, fielded in mid-2024 among 328 non-executive directors, found that 80 percent considered board practices and structures inadequate to oversee AI. That figure is the oldest cited in this paper and should be read as directional rather than current. More recent work is consistent with it: a majority of directors report limited or no AI knowledge, and only about half of boards have held or been offered AI-focused director education.

Two conclusions follow. First, a board that has not formalized AI oversight is presently in the majority, which means the absence of oversight is not yet a competitive disadvantage but is an

increasingly visible one. Second, the disclosure and education data suggest that where oversight does exist, it frequently rests on a single director, which is a concentration risk in itself.

3. The exposure is already concrete

Three channels have moved from anticipated to active.

Securities enforcement and disclosure accuracy

The Securities and Exchange Commission began enforcement on overstated AI claims in March 2024 with settled actions against two registered investment advisers, Delphia (USA) Inc. and Global Predictions Inc. In January 2025 it charged Presto Automation Inc., the first such action against a public company, in relation to representations about its voice AI technology, including the extent of third-party provenance and the level of human intervention actually required. The Commission has established a Cyber and Emerging Technologies Unit within its Division of Enforcement, and its fiscal 2026 examination priorities direct examiners to assess whether AI representations are accurate and whether operations align with disclosures.

There is no AI-specific SEC disclosure rule as of mid-2026. Enforcement has proceeded under existing anti-fraud provisions, which is the point of consequence for private companies: the same substantiation logic is applied by acquirers in diligence, by insurers in underwriting, and by counterparties enforcing contractual representations. A private company that cannot evidence its AI claims faces the commercial version of the same problem.

United States state law

For a U.S. organization, state statutes are more likely to bite than the EU AI Act, and the position changed materially over the past eighteen months.

- The Texas Responsible Artificial Intelligence Governance Act (HB 149), signed 22 June 2025, took effect 1 January 2026. It applies an intent-based liability model, is enforced exclusively by the Attorney General with no private right of action, and provides a 60-day period to cure.
- Illinois HB 3773, addressing AI in employment decisions, took effect 1 January 2026. California and Utah obligations are also in force, with California transparency requirements becoming operative in August 2026.
- Colorado's SB 26-189, signed 14 May 2026, repealed and reenacted the state's AI Act framework and moved core automated decision-making duties to 1 January 2027.

Reach generally attaches through the residents affected rather than through the organization's address, so a company headquartered elsewhere can still be within scope.

One provision bears directly on the recommendations in this paper. TRAIGA establishes an affirmative defense for organizations that substantially comply with the most recent NIST AI Risk Management Framework, including its generative AI profile, and that identify issues through internal testing, red-teaming, or documented internal review. In Texas, documented governance is not merely prudent. It is a statutory defense, and the documentation is the thing that has to exist in advance.

A federal preemption effort is live but has not succeeded. Executive Order 14365 of 11 December 2025 directed the Department of Justice to establish an AI Litigation Task Force to challenge state AI laws, and the White House published a National Policy Framework on 20 March 2026 recommending that Congress legislate preemption. As at August 2026 no federal preemption of

state AI law has been enacted, and the obligations above remain in force. A board should not treat the preemption debate as a reason to defer.

Regulatory obligations with extraterritorial reach

The EU AI Act entered into force on 1 August 2024 on a phased schedule. Its position in mid-2026 requires care, because the headline coverage has been misleading.

- Following the Digital Omnibus package, endorsed by the European Parliament on 16 June 2026 and approved by the Council on 29 June 2026, obligations for high-risk systems under Annex III are deferred to 2 December 2027, and for AI embedded in regulated products under Annex I to 2 August 2028.
- The Article 50 transparency obligations and the Commission's general-purpose AI enforcement powers were not deferred and applied from 2 August 2026.
- The Act has extraterritorial reach. A U.S. organization can fall within scope where its AI system's output touches the EU through sales, access, or downstream integration.

The practical implication for a board is that the deferral applies to one tier only. An organization concluding from press coverage that it has until December 2027 to consider the Act may already be out of step with obligations that took effect this month.

Commercial and transactional exposure

The most likely route by which a private company board first encounters this is not a regulator. It is a customer contract containing AI representations, a diligence questionnaire in a sale or financing process, an insurer's renewal questions, or a lender's covenant. In each case the organization is asked to evidence something it may never have documented. Answering badly once is expensive, and the first occasion should not be during a transaction.

4. What adequate oversight looks like

Adequacy here is a governance standard, not a technical one. The NIST AI Risk Management Framework places Govern as its foundational function, and its content maps closely to what the oversight duty requires: defining who has authority to approve deployments, who owns each system once deployed, how AI risk is reported to management and the board, what the organization's risk tolerance is by category of use, and what the incident response requires when a system fails or produces harmful output.

For most private company and nonprofit boards, that resolves into three artifacts and one habit. None of them requires a platform purchase.

ARTIFACT ONE · A DECISION AUTHORITY MAP

A written record of which decisions AI informs, which it effectively makes, and which must remain explicitly and documentably human. This is the foundational control, because every other measure is theoretical without it. It is usually the first occasion on which a leadership team sees the full picture at once, and it commonly reveals that the boundary was never set, only accumulated.

ARTIFACT TWO · AN ACCOUNTABILITY FRAMEWORK

Who owns the outcome when AI is in the loop, what happens when it is wrong, and how accountability flows through an AI-assisted decision chain. The framework should identify where the chain currently breaks. Accountability assigned to a team rather than a person is the most common failure and is functionally equivalent to no assignment.

ARTIFACT THREE · A RECONSTRUCTION STANDARD

The organization's ability to explain, six or twelve months later, how a specific material decision was reached. This is the board's position in litigation, in regulatory review, and in diligence, and it is also the mechanism by which the organization learns from failure. The output of a system is not the reasoning behind a decision.

THE HABIT · STANDING CADENCE AND NAMED OWNERSHIP

A defined reporting cadence to a named committee with at least two directors able to interrogate what they are shown. Ownership without capability is worse than no ownership, because it closes the question. Where fluency is absent, the board's options are director education, an independent director recruited for this capability, or an external advisor retained to the committee.

A note on proportionality and failure modes

Governance can be built badly, and over-built governance fails in a characteristic way. The Grant Thornton research cited above for the oversight gap also records that centralized review bodies become overwhelmed as use cases multiply, creating bottlenecks that slow the business without reducing risk. The same research found that organizations which built governance first adopted AI faster rather than slower. The distinction that matters is between governance which establishes who decides and what gets documented, and governance which requires a committee to approve each individual use. The first scales. The second collapses and is then abandoned, which leaves the organization worse placed than before, because it now has a failed process on the record.

Proportionality applies to everything in this section. For a forty-person nonprofit or an early-stage company, the decision authority map may run to two pages and the accountability framework may be a single named owner with an escalation path to the chief executive. The three artifacts scale down without losing their function. What does not scale down is the requirement that a named person owns it and that the board sees it on a defined cadence.

5. Recommended actions

The following are drafted so that each can be minuted as a resolution or a direction to management.

1. **Assign ownership.** Name the committee responsible for AI oversight and record the assignment in that committee's charter. Confirm that at least two directors are able to interrogate management on the subject, and if not, adopt a plan to close that gap through education or recruitment.
2. **Direct a decision authority map.** Instruct management to produce, within 90 days, a written map of which decisions AI currently informs, which it effectively makes, and which are designated as requiring human judgment. Ask specifically for override rates where AI output is characterized as advisory.
3. **Direct an accountability framework.** Instruct management to identify, for each material AI-assisted decision category, the individual accountable for the outcome, the escalation path when the system errs, and the incident response. Require that accountability be assigned to named individuals rather than to functions.
4. **Establish reporting cadence.** Adopt AI as a standing item on the responsible committee's agenda with a defined frequency, and specify the content of the report rather than leaving its scope to management's discretion. Marchand is explicit that discretionary management reporting does not discharge the board's obligation for mission-critical risk.
5. **Review external representations.** Direct a review of what the organization currently claims about its AI capabilities in marketing materials, customer contracts, financing documents, insurance applications, and grant submissions, and confirm each claim can be substantiated.
6. **Establish a documentary record.** Ensure that the board's consideration of this paper, the assignment of ownership, and the directions above are minuted. Where good faith is demonstrated through governance design rather than the absence of failure, the minute book is the evidence.

A board that completes these six actions will have converted an undocumented exposure into a governed one, and will hold a documentary record if it is ever asked to demonstrate as much. None of it requires the board to become technical.

In closing

AI governance is not principally a question about technology. It is a question about whether the board can demonstrate that it established an oversight system proportionate to the organization's use of AI, that information reached it, and that it acted on what it saw. Directors do not need to become experts in the underlying systems. They do need to ensure that authority, accountability,

evidence, and reporting are settled deliberately, before the absence of those structures becomes visible during a transaction, an insurance renewal, a regulatory enquiry, or a failure.

Appendix: fifteen questions for management

The following is a discussion tool for the committee holding AI oversight. It is not a compliance checklist and is not intended to be worked through in a single sitting. Each question is followed by a note on what an unsatisfactory answer tends to sound like.

DECISION AUTHORITY

1. **Which decisions does AI currently inform, and which does it effectively make?**
   "It is only advisory" invites the follow-up: how often is the recommendation overridden? An override rate near zero indicates the decision has moved.

2. **Which decisions must remain human, and is that recorded in writing?**
   An unwritten boundary is a shared assumption, and shared assumptions do not survive turnover.

3. **Who approved the current arrangement, and when?**
   A reference to vendor selection is not responsive. Procurement is not governance, and the absence of an approval decision is itself the finding.

4. **Where in the business would an AI error be least likely to be caught?**
   "We review everything" is rarely true at scale. The risk concentrates where output is plausible, unverified, and consumed downstream as fact.

5. **What has AI changed about the speed of decisions, and was the removed deliberation serving a control function?**
   Cycle-time improvement reported without reference to what the time was doing.

ACCOUNTABILITY AND EVIDENCE

6. **When an AI-assisted decision proves wrong, who owns it?**
   "The team." Accountability assigned to a group is not assigned.

7. **Can we reconstruct how a specific material decision was reached six months after the fact?**
   "We would still have the output." The output is not the reasoning.

8. **What is our disclosure position with customers, partners, insurers, and lenders?**
   "It has not come up yet." It will, and the first occasion should not be during a transaction.

9. **Can every public and contractual claim we make about our AI capability be substantiated?**
   Confidence without documentation. This is the substance of the SEC's enforcement theory and of most diligence questionnaires.

10. **Has anything been reversed, escalated, or disciplined as a result of an AI-related error?**
    A wholly clean record in a heavily adopting organization usually indicates that reporting is not working, rather than that performance is flawless.

DEPENDENCY, PEOPLE, AND THE BOARD

11. **Which capabilities now exist only because AI performs them, and what happens if access is lost?**
    "We would revert to the manual process." Ask who still knows it.

12. **Is the judgment of our managers improving or atrophying?**

Any answer framed solely in productivity terms. The long-term risk is a cohort of managers who never developed the judgment to recognize a bad output.

13. **What are our people doing with AI that policy does not yet cover?**
    "We prohibit it." A prohibition without a sanctioned alternative produces concealment rather than compliance.

14. **What should management report to this board about AI, how often, and in what form?**
    Ad hoc updates when something notable occurs. That is management's discretion, not the board's oversight.

15. **Does the committee that owns this have the fluency to interrogate what it is shown?**
    "The full board handles it," which in practice means no one prepares.

Sources

Legal: In re Caremark Int'l Inc. Derivative Litig., 698 A.2d 959 (Del. Ch. 1996); Stone v. Ritter, 911 A.2d 362 (Del. 2006); Marchand v. Barnhill, 212 A.3d 805 (Del. 2019); In re McDonald's Corp. Stockholder Derivative Litig. (Del. Ch. 2023); DGCL section 102(b)(7).

Regulatory, federal and international: SEC enforcement actions against Delphia (USA) Inc. and Global Predictions Inc. (March 2024) and Presto Automation Inc. (January 2025); SEC Division of Examinations fiscal 2026 priorities; Executive Order 14365 of 11 December 2025; White House National Policy Framework for Artificial Intelligence, 20 March 2026; Regulation (EU) 2024/1689 (the AI Act) and the Digital Omnibus amendments endorsed by the European Parliament on 16 June 2026 and approved by the Council on 29 June 2026; NIST AI Risk Management Framework (AI RMF 1.0).

Regulatory, U.S. state: Texas Responsible Artificial Intelligence Governance Act (HB 149), effective 1 January 2026; Illinois HB 3773, effective 1 January 2026; Colorado SB 26-189, signed 14 May 2026; California and Utah AI provisions in force during 2026.

Data: Stanford HAI, 2026 AI Index; ISS-Corporate review of 3,048 U.S. companies (January 2026); Grant Thornton 2026 AI Impact Survey; Protiviti 2026 board governance research; Gartner survey of non-executive directors; Private Directors Association and Arizona State University New Governance Lab, Private Company Governance 2026 Report.

Figures are as reported by the cited sources and reflect the position as at August 2026. Regulatory positions continue to move, particularly in the European Union and at U.S. state level, and should be confirmed with counsel before being relied upon.

About the author

Andrew Brümmer is an operating executive with more than twenty years of leadership across technology, healthtech, and professional services in the United States, Europe, the United Kingdom, and Africa. He served as Board Observer and Executive Liaison to the board of LifeQ, a venture-backed healthtech company, from 2020 to 2025, and as its Chief Operating Officer and Chief of Staff from 2015 to 2025, where he led the introduction of AI-enabled decision support across an organization operating on three continents and scaled the company from 42 to 114 people in eighteen months.

He is the founder of The Ardunan Village, advising founders, chief executives, and boards on AI adoption and governance, and he built and delivers the AI Operator Method, a cohort programme for senior operators. He is the author of Leading Magnanimously and You Decide, and is a member of the Private Directors Association.

He delivers this paper to boards and leadership teams as a structured session covering AI governance risk, oversight design, and the questions to be asking now. He is not an AI vendor and does not sell a platform, implement tools, or manage vendors.

Readers should note that the recommendations set out here align with work the author is engaged to perform. The argument is offered on its evidence and its citations, and should be weighed accordingly.

abrummer@ardunan.com · +1 (404) 915-2027 · ardunan.com

© 2026 The Ardunan Village. This paper may be shared freely with directors and leadership teams in its complete form. It is a governance discussion document and does not constitute legal advice.