FAIR LABS FRONTIER SERIES

VOL 10 · 2026

The Director's Dilemma

What boards owe oversight now that every company is at least dependent on advanced AI

PREPARED FOR BOARDS OF DIRECTORS · GENERAL COUNSEL · C-SUITE EXECUTIVES · INSTITUTIONAL INVESTORS

FAIR Labs
FAIR ARTIFICIAL INTELLIGENCE RESEARCH LABS

May 2026
FL-2026-10

FAIR LABS FRONTIER SERIES · VOLUME 10

The Director's Dilemma: What boards owe oversight now that every company is at least dependent on advanced AI

The tenth and closing volume of the Frontier Series is written for the people who will be asked, after the fact, what the board knew and when. It argues that AI oversight is now squarely a board duty for every company — not just the ones building frontier models, since even the most AI-distant firm is at minimum dependent on vendors and infrastructure this series treats as a monoculture risk. The report names three postures a company can hold, builds the due-diligence playbook and committee architecture each requires, and closes the series with the self-assessment a board can run before its next AI-related decision.

REPORT CODE	FL-2026-10
PUBLICATION DATE	May 2026
PRIMARY AUDIENCES	Boards of Directors · General Counsel · C-Suite Executives · Institutional Investors
SERIES	FAIR Labs Frontier Series, Fair Artificial Intelligence Research Labs
SUGGESTED CITATION	*The Director's Dilemma: What boards owe oversight now that every company is at least dependent on advanced AI.* FAIR Labs Frontier Series, Vol. 10. Fair Artificial Intelligence Research Labs, 2026.
KEYWORDS	corporate governance, board oversight, fiduciary duty, due diligence, AI risk disclosure

IN BRIEF · EXECUTIVE SUMMARY

Every company is now at least dependent on advanced AI. That makes oversight a board duty, not a management memo.

You do not need to build frontier models to carry frontier-scale risk. A company that merely runs on AI-dependent vendors and infrastructure inherits the concentration and continuity exposure this series spends nine volumes mapping. This closing volume asks what a board specifically owes oversight, once "we don't really do AI" stops being a defense.

01 There is no fourth posture called "uninvolved."

Every company is a developer, a deployer, or at minimum a dependent — reliant on vendors and infrastructure whose failure modes are now the company's own.

02 The board usually knows less than it thinks.

Confident management demos are not evidence. The information gap between what leadership knows and what the board is told is often widest exactly where the risk is highest.

03 "We have an AI committee" is not an answer by itself.

Committee architecture matters less than whether it has real cadence, real expertise, and a real reporting line — decorative committees provide no assurance at all.

04 Due diligence has a specific, checkable shape.

Assurance case reviewed, concentration mapped, incident plan tabletopped — not a vague comfort briefing from the team proposing the deployment.

05 Disclosure is a market-pricing problem, not a PR problem.

Investors and insurers cannot price a risk the company will not name. Silence is not neutral; it is a bet that nothing goes wrong before someone asks.

06 Posture shifts — the assessment has to repeat.

A company that was purely "dependent" last year may be a deployer this year. The due-diligence exercise is annual, not a one-time board resolution.

CONTENTS

Inside this report

01

The new fiduciary terrain

Board oversight of major risk categories is not a new duty. What is new is that AI now qualifies as one, for nearly every company, whether or not it appears in the firm's own description of what it does.

Directors have always been expected to oversee the risks material to their company — that is what a risk committee, an audit committee, and the board's general oversight duty already exist to do. This volume does not propose a new species of legal obligation; it argues that advanced AI has become material fast enough that most boards' existing oversight machinery has not yet caught up to it. The gap is not one of law; it is one of practice, expertise, and information flow — and it is closing far more slowly than the exposure it is meant to cover is growing.

What makes this volume different from a generic "technology risk" briefing is the same principle running through the whole series: **every company now holds one of three postures** with respect to advanced AI, and each posture carries duties distinct enough that a single generic "AI policy" satisfies none of them well.

> **THE THREE POSTURES, IN BRIEF**
>
> **Developer** boards oversee companies building frontier or near-frontier models — the safety-scale duties Volume 01 addresses. **Deployer** boards oversee companies putting AI to work in their own products or operations — Volume 06's deployment-risk territory. **Dependent** boards oversee companies that do neither, yet still run on AI-dependent vendors and infrastructure — inheriting Volume 02 and Volume 03's concentration and continuity exposure whether they sought it or not.

Almost no board is purely one posture; most are a dependent board with pockets of deployment. The chapters ahead take this seriously rather than collapsing it into one-size-fits-all guidance. Chapter 2 maps duties to posture. Chapter 3 names the information-asymmetry problem underneath all three. Chapter 4 builds committee architecture and a due-diligence playbook. Chapter 5 turns to disclosure. Chapter 6 is an eight-point agenda; Chapter 7 closes the series with a self-assessment for the next board meeting where this comes up.

> *"We don't really do AI" stopped being a governance answer the day the company's critical vendors started answering it for them.*
>
> THE PREMISE OF THIS VOLUME

02

Three postures, three duties

A developer board, a deployer board, and a dependent board are answering different questions, even when they use the same word — "AI risk" — to describe them.

2.1 Developer boards: safety at the scale of the company's own models

A board overseeing a frontier or near-frontier developer carries the assurance duties Volume 01 addresses directly: is the company's own capability-assurance gap closing or widening, is evaluation and interpretability work funded at a level proportionate to what the company is building, and does the board receive that information from a source independent of the team racing to ship the next release.

2.2 Deployer boards: the assurance case as the actual approval gate

A board overseeing a company deploying AI into its own products, operations, or customer-facing decisions carries Volume 06's deployment-risk duties: has a signed assurance case been produced before go-live, does the deployment have a graceful-degradation plan rather than a binary on/off switch, and is fallback capacity funded as standing insurance rather than assumed away.

2.3 Dependent boards: continuity and concentration, inherited not chosen

A board whose company does neither still depends on vendors, cloud providers, and infrastructure running on the same handful of base models as everyone else — Volume 03's monoculture risk and Volume 07's concentration problem, inherited rather than chosen. The duty here is unglamorous and easy to skip: mapping vendor concentration the way a competent board already maps single-supplier risk in any other critical input, and asking what happens if two or three key vendors have a correlated bad week at the same time.

EXHIBIT 10.1 · FAIR LABS ASSESSMENT (ILLUSTRATIVE INDEX)

Which duties bind which posture

	Capability & safety oversight	Deployment risk	Concentration & continuity risk	Disclosure duty
Developer	primary	material	light	primary
Deployer	light	primary	material	material
Dependent	—	light	primary	material

FAIR Labs analytical construct for orientation; most companies sit across more than one posture and should read more than one row.

THE POSTURE QUESTION IS NOT RHETORICAL

Ask it plainly at the next board meeting: are we a developer, a deployer, a dependent, or — most likely — some mix, and has that mix changed since we last asked? Companies drift from "dependent" toward "deployer" quietly, one internal tool at a time, without a single board resolution marking the transition.

03

What the board doesn't know it doesn't know

The confident demo is the most dangerous artifact in AI governance, because it is designed, deliberately or not, to answer a different question than the one the board needs answered.

3.1 The demo problem

A live demonstration shows a system succeeding on the cases chosen to show it succeeding. It answers "can this work" — a question no one seriously disputed — rather than "how does this fail, how often, and how would we know." A board that leaves an AI briefing reassured by a smooth demo has usually been shown evidence calibrated to produce exactly that reassurance, not evidence calibrated to surface the failure modes oversight exists to catch.

3.2 Where the gap is widest

Awareness gaps are not uniform across topics. Directors are told about budget — AI and compute spend shows up in every board deck, because management wants approval for it. Directors hear far less about near-misses, vendor concentration, and workforce impact planning, precisely the categories least flattering to report and least likely to surface unprompted.

EXHIBIT 10.2 · FAIR LABS ASSESSMENT (ILLUSTRATIVE INDEX)

The awareness gap is not evenly spread across topics

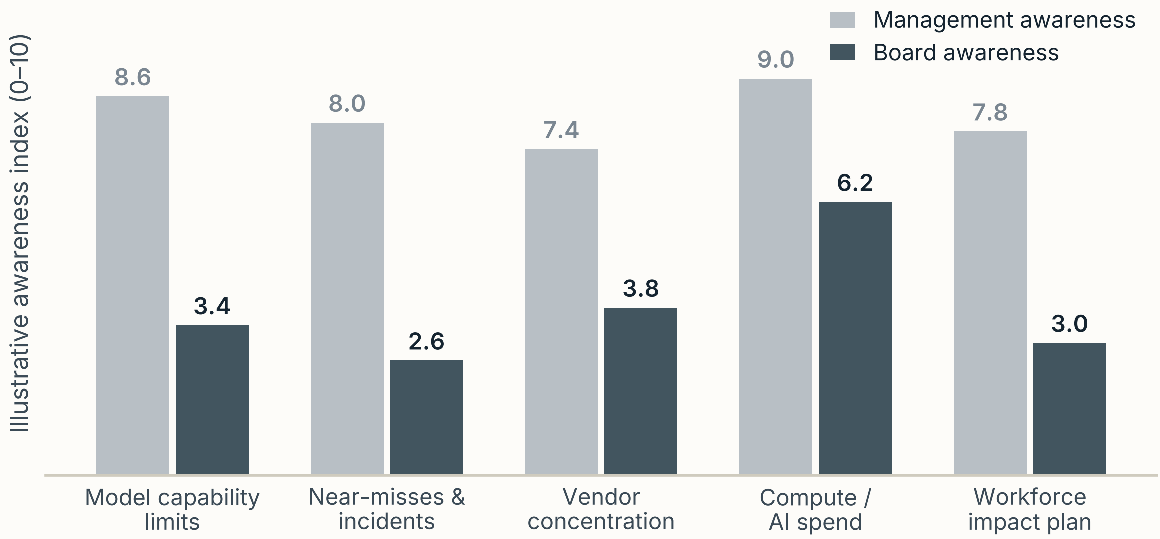

Illustrative FAIR Labs index for orientation; actual awareness varies enormously by company and is not a measured survey finding.

3.3 Questions that unmask the gap

Three questions reliably separate genuine assurance from a well-rehearsed briefing: what is the worst documented failure of this system to date, and who outside the team that built it confirmed that account? What would the board be told if the answer were embarrassing rather than reassuring? And who, specifically, is accountable if this goes wrong in a way nobody in the room anticipated? A management team that answers fluently has probably done the work. One that reaches for the demo again has not.

04

Committee architecture and due diligence

Where oversight sits matters less than whether it has cadence, expertise, and a reporting line independent of the team being overseen.

4.1 Where AI oversight should sit

Three architectures are defensible, and the right choice depends on the company's posture mix more than any universal rule. An existing risk or audit committee can absorb AI oversight if it has the bandwidth and is willing to develop real fluency rather than treating it as one line item among many. A dedicated technology or AI committee signals seriousness and can develop deeper expertise, at the risk of siloing AI risk away from the board's broader risk picture. Full-board ownership works only for companies small enough that every director can plausibly stay current. What fails in every configuration is a committee that meets annually, hears a single management presentation, and produces no independent question all year.

4.2 The due-diligence playbook

Concrete board-level due diligence, whatever the committee architecture, has a checkable shape: has a signed assurance case (Volume 06) been reviewed before this specific deployment, not just described in the abstract; has vendor and model concentration been mapped the way any other single-supplier risk would be; has the incident response plan been tabletopped at board level, not just filed; and — for workforce-affecting adoption decisions — has the augmentation-versus-replacement planning Volume 08 describes actually been done before the decision, not retrofitted after complaints arrive.

EXHIBIT 10.3 · ILLUSTRATIVE

What due diligence looks like in practice today — most of it doesn't happen yet

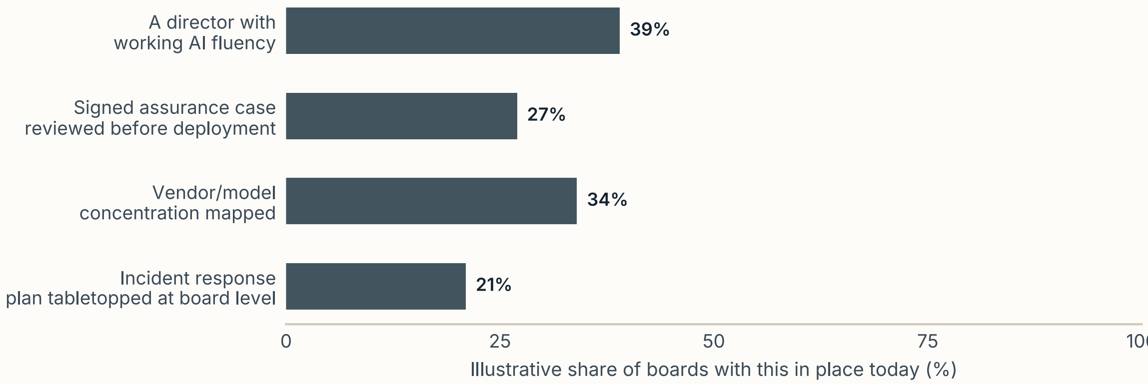

Illustrative index for orientation, not a survey finding about any named company or population of boards.

4.3 M&A and vendor diligence

The same playbook extends to acquisitions and major vendor relationships: an AI-exposed acquisition target should be asked for its own assurance case and concentration map before close, not after — the liabilities in Volume 03's atlas transfer with the company, whether or not the diligence process asked about them.

05

Disclosure and the market for trust

A market cannot price a risk the company will not name. Silence is not a neutral choice; it is a bet that the question never gets asked before something goes wrong.

5.1 Four levels of disclosure maturity

Disclosure practice today ranges across a rough ladder. Some companies remain silent — no mention of AI dependency or exposure anywhere material. Many have graduated to boilerplate: generic risk-factor language naming "artificial intelligence" without firm-specific detail, satisfying a checklist without informing anyone. A smaller set produce quantified disclosure — firm-specific concentration, incident history, and spend an investor could actually use. Fewer still reach assured disclosure, where those figures are reviewed by an independent party the way financial statements are, per Volume 04's verification-layer agenda.

EXHIBIT 10.4 · CONCEPTUAL FRAMEWORK

Disclosure maturity: from silence to assured, audited figures

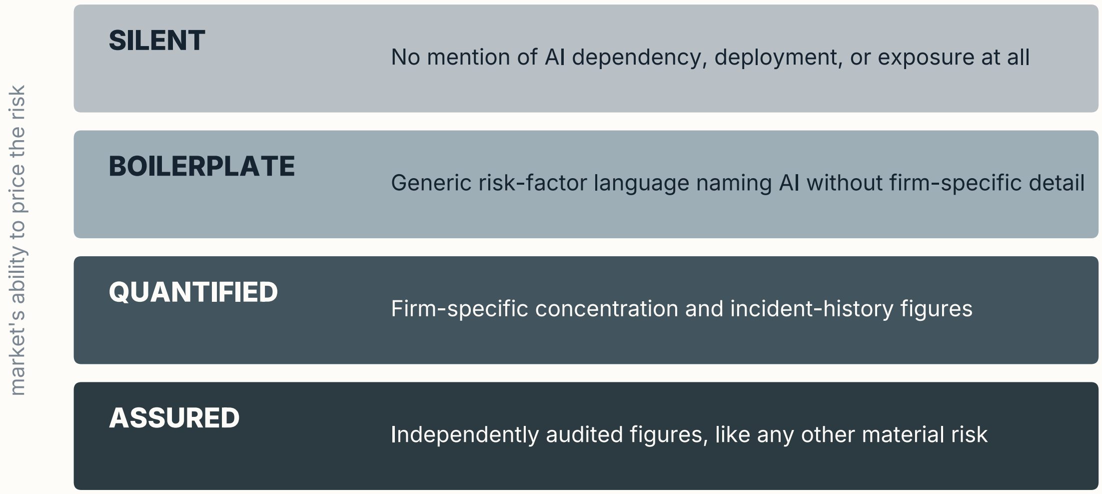

Illustrative construct describing a maturity progression, not a claim about the current practice of any named company or market.

5.2 Insurance follows disclosure

Insurers can only price what companies disclose, and Volume 03's watchtower board already lists insurers and reinsurers among the owners who most need this information. A company at the "silent" or "boilerplate" rung is not avoiding scrutiny; it is simply pricing its own risk badly, by default, in a market where competitors further up the ladder will look — and eventually be priced — as the more credible bet.

5.3 What good disclosure says, specifically

Useful disclosure names the company's posture (developer, deployer, dependent, or a stated mix); quantifies vendor and model concentration in its critical functions; states whether a governance committee with real cadence exists; and reports material incidents on a timeline consistent with Volume 04's reporting bar, rather than only when a disclosure becomes unavoidable.

06

Recommendations

Eight moves that convert this volume's argument into standing board practice, each with an owner.

BD-01

Name a committee that owns AI oversight, and give it real cadence

Existing risk committee, a dedicated technology committee, or the full board — any of the three works if it meets, questions, and reports independently.

OWNERS: **BOARDS · GENERAL COUNSEL**

BD-02

Demand the signed assurance case before approving deployment

Not a demo, not a management assurance — the structured document Volume 06 specifies, with a named signatory.

OWNERS: **BOARDS · RISK OFFICERS** · DEEP DIVE: VOL 06

BD-03

Map vendor and model concentration like any other supply-chain risk

Treat a shared base model across critical vendors the way procurement already treats a single-source part.

OWNERS: **BOARDS · PROCUREMENT** · SEE ALSO: VOLS 03, 07

BD-04

Fund one director's real fluency, not everyone's polite familiarity

At least one board member with working technical literacy, resourced to develop and maintain it, not a single onboarding session.

OWNERS: **BOARDS · NOMINATING COMMITTEES**

BD-05

Require workforce due diligence before adoption decisions, not after

Augmentation-versus-replacement planning done in advance, with a named owner accountable for the transition.

OWNERS: **BOARDS · PEOPLE LEADERS** · DEEP DIVE: VOL 08

BD-06

Disclose AI exposure like any other material risk

Firm-specific concentration, incident history, and governance structure — not boilerplate naming the technology without naming the exposure.

OWNERS: **BOARDS · INVESTOR RELATIONS**

BD-07

Tabletop the incident before it happens

A board-level exercise, on a fixed cadence, rehearsing the response to a plausible AI-related failure — not a policy document nobody has opened.

OWNERS: **BOARDS · GENERAL COUNSEL** · SEE ALSO: VOL 06

BD-08

Re-assess the company's posture annually

Developer, deployer, dependent, or a shifting mix — re-ask the question every year, since drift between postures rarely announces itself.

OWNERS: **BOARDS · RISK COMMITTEES**

07

A board self-assessment

Eight questions for the next board meeting where AI comes up — worth answering honestly, and worth revisiting every year after.

QUESTION	GREEN ANSWER	RED FLAG
What is our posture — developer, deployer, dependent, or a mix?	Named explicitly, reassessed on a schedule	Never asked as an explicit question
Who owns AI oversight, and how often do they meet?	A named committee, meeting on a real cadence, with independent questions	An annual slide, presented once, unchallenged
Have we seen a signed assurance case for our deployments?	Yes, reviewed before go-live, with a named signatory	A demo stood in for a safety case
Do we know our vendor and model concentration?	Mapped like any other supply-chain risk	"We use several vendors" with no further detail
Does a director have real technical fluency?	Resourced, ongoing, more than one onboarding session	Polite familiarity substituting for literacy
Has the incident response been tabletopped?	Rehearsed at board level, on record, on a cadence	A policy document nobody has opened
Do we disclose exposure, or just mention the technology?	Firm-specific figures an investor could use	Boilerplate risk-factor language only
Was workforce impact planned before the adoption decision?	Done in advance, with a named owner	Retrofitted after complaints arrive

FAIR Labs diagnostic, intended for any board to run against its own practice; not an assessment of a named company.

THE SERIES, IN ONE PAGE

This closes the Frontier Series where Volume 01 opened it: a widening gap between what advanced AI can do and our collective capacity to assure, govern, deploy, and oversee it well. Ten volumes, one shared evidence base, ten vantage points — regulators to Volume 04, deployers to Volume 06, compute policymakers to Volume 07, and now boards, whose due diligence closes the loop on every recommendation the other nine volumes make. The gap narrows only if every reader with a hand on the wheel — including the one holding this volume — takes their assigned page seriously.

Fair Artificial Intelligence Research Labs (FAIR Labs) is an independent research institute working to ensure that advanced artificial intelligence is developed safely, governed wisely, and shared fairly. Through rigorous technical research, policy analysis, and public engagement, FAIR Labs works with scientists, governments, companies, and communities to widen the circle of people who benefit from — and have a say in — the most consequential technology of our time.

The **Frontier Series** is FAIR Labs' flagship collection of stakeholder briefings on advanced AI — one shared evidence base, ten vantage points. Each volume stands alone; together they form a common map for scientists, executives, policymakers, workers, and citizens navigating the same transition.

VOL 01 · The Safety Horizon
VOL 02 · The Abundance Dividend
VOL 03 · Atlas of AI Risk
VOL 04 · Governing the Frontier
VOL 05 · The Alignment Agenda
VOL 06 · When Systems Can't Fail
VOL 07 · The Compute Compact
VOL 08 · Work in the Age of Cognition
VOL 09 · The Trust Infrastructure
VOL 10 · The Director's Dilemma