



The Center for
AI Oversight

CENTER PAPER · SEPTEMBER 2026

Is AI Mission Critical?

The Caremark question for directors and officers: which uses of AI create a duty of oversight, and what the record must show

Delaware courts let oversight claims proceed against directors who fail to oversee a company's "mission critical" risks (the duty is Caremark's; the phrase is Marchand's), but have never said what makes a risk mission critical. Seven years of cases show what they look at. This paper sets out the three determinations a board records for each use of AI, the criteria the courts have used for each, the obligation that follows, and the program and record that make a board's position defensible.

By **Brian J. Allen**, Executive Director, The Center for AI Oversight

Written for directors, audit and risk committee members, general counsel, and the risk, compliance and audit executives who report to them. The argument runs through Part V. The terms, the worksheet, the model resolution and the case table are for reference.

EXECUTIVE SUMMARY

Delaware requires a board to make a good-faith effort to have a board-level reporting system for the compliance risks at the center of its business, and to act on what it reports. The courts have never said which risks those are. A chartered program that makes three determinations for each AI use, and keeps the record, is the evidence a court will ask for.

THE RIGHT LAST QUESTION

"Is AI mission critical?" is the question directors ask. It is the right question to end with and the wrong one to start with. The duty to have a board-level reporting system attaches to any compliance risk central to the business, whether or not a court would call it mission critical; a plaintiff need not plead mission criticality to state a claim. A board that asks only the mission-critical question and answers "no" will hear "no duty," and for most companies most uses of AI fall exactly there: in the middle of the three circles this paper draws (business risk in the outer ring, central compliance risk in the middle, mission critical at the center), where a system is required. The analysis a board records therefore has three determinations, made in order, of which mission critical is the third.

First, is this use a compliance risk or a business risk? If a business risk, no Delaware court will impose oversight liability for failing to monitor it; how the board is informed about it is its business judgment, and the record is the decision it made. A business risk the company could not survive still goes on the inventory, and Part II says why. Second, is the compliance risk central to the business or peripheral? If central, the board must have a reporting system for it; this is the determination that creates the obligation. Third, is it mission critical: is the regulated activity the company's core product, service or purpose, and would a failure harm the people the company serves? If so, the system must exist in fact, an unaddressed warning will be read as bad faith more readily, and the record must show the system ran.

The five questions in Part II are the evidence for these determinations. The three circles below are the picture of them; Part I gives their sources.

KEY TAKEAWAYS

- 1 Since 2019, Delaware courts have decided more than twenty board-oversight cases without defining "mission critical." Five questions drawn from those cases supply the evidence; three determinations, made

in order, supply the answer. Mission critical is the third. The second, whether the risk is a compliance risk central to the business, is the one that creates the duty.

- 2 The operative legal category is the "central compliance risk": a risk of failing to meet a legal or regulatory requirement, sitting at the center of the business. "Mission critical" is the courts' word for the subset where the regulated activity is the company's core product, service or purpose and a failure would harm the people the company serves: Blue Bell's ice cream, Boeing's airplanes, a bank's lending, a health plan's coverage decisions.
- 3 The company need not have violated anything. Blue Bell's nominal compliance with FDA rules did not save its board. What the courts hold a board to is a system for the compliance risk, and a response to what it reports. Business risks, however large, are outside the liability. They are not outside the board's job, and one the company could not survive can become a compliance risk by three routes (Part II).
- 4 AI is not one risk. AI that decides who gets care, credit or a job under a regulator's rules is a compliance risk. AI capital spending is a business risk; only its disclosure is governed by law.
- 5 A mission-critical use needs a board-level system that exists in fact, not on paper: a committee with responsibility, a schedule, reporting management cannot skip, an escalation process for bad news, and a response when it arrives. It does not need omniscience.
- 6 Oversight liability requires bad faith, not negligence. A chartered program that records the three determinations for each use, and its oversight of the result, is the board's strongest evidence on either side of the line. Never having asked the question is the opposite.

THE FIVE QUESTIONS (PART II)

- 1 Is this a compliance risk, or a business risk?
- 2 How central is it: is the regulated activity the company's core product, service or purpose?
- 3 Which rule governs it, and who enforces it?
- 4 What harm follows a failure, to whom, and at what scale?
- 5 What has the company said about it, including about the board's oversight of it?

THE THREE DETERMINATIONS (PART II)

DETERMINATION	ANSWERED BY	CONSEQUENCE
1. Compliance risk, or business risk?	Question 1	No oversight liability for a business risk; the board decides how it will be informed, and records the decision.
2. Central to the business, or peripheral?	Question 2, with 3 and 5 as evidence	A central compliance risk requires a board-level reporting system. This creates the duty.
3. Mission critical?	Questions 2 and 4 together	The same system, read more strictly: it must exist in fact, and an unaddressed red flag is read as bad faith more readily.



The three circles. Determination 1 or 2 answered no puts a use in the outer ring; Determination 2 puts it in the middle; Determination 3 puts it at the center.

RECOMMENDATION

Charter the program now, use by use, rather than wait for a court to describe what was missing.

- 1 **Charter it.** Assign AI oversight to a named committee, amend its charter to say so, set the escalation path for bad news, and align public statements about AI oversight with what the program will actually do.
- 2 **Inventory.** Management lists every AI use that makes or informs a decision about a person under a rule, controls a safety function, performs a compliance function, handles protected data, depends on a third-party model or vendor-supplied system, supports a public statement, or is one the company's revenue or operations cannot run without.
- 3 **Answer the five questions** for each use, in writing. Counsel does this; Appendix A is the worksheet.
- 4 **Make the three determinations** for each use, in order: compliance risk or business risk; central or peripheral; mission critical or not. Each has its own consequence, and the second creates the duty. The placement in a circle follows from the three.
- 5 **Oversee.** Each mission-critical use goes on the committee's schedule with a protocol that fixes what management must report, and tolerances the board has ratified. Internal audit confirms the reporting runs.
- 6 **Record and revisit.** Record the determinations in a board resolution (a model is in Appendix B). Revisit annually, or on any new use, new rule, incident or material public statement.

The legal floor is lower than this program (Part IV). The program is what makes the good-faith effort provable. We are aware of no court decision on a board-oversight claim about AI. By the Center's survey, about ten derivative complaints had been filed over a company's use of AI as of September 2026. Four plead a theory about the use itself, and each has a compliance hook: Medicare coverage rules at UnitedHealth; copyright at Adobe, Microsoft and Nvidia. The rest are follow-on suits to securities class actions pleading misstatements and internal-control failures, and the one built on an operational failure alone, the July 2024 CrowdStrike outage, was voluntarily dropped in April 2026. None has been tested on a motion to dismiss. Each of the four will turn on exactly this record.

TERMS, AS THIS PAPER USES THEM

Each term carries the case or instrument it comes from. Full citations are in Appendix C and the sources.

THE ANALYSIS

The three determinations. The conclusions a board records for each AI use, in order: whether the use presents a compliance risk or a business risk; whether a compliance risk is central to the business or peripheral; and whether a central compliance risk is mission critical. The five questions in Part II are the evidence for them. The second determination creates the duty to have a reporting system; the third changes how strictly a court reads the record. (Center framework, drawn from Citigroup (2009), McDonald's (Del. Ch. Mar. 1, 2023) and Marchand (2019).)

THE DUTY

Fiduciary duties: care and loyalty. Delaware directors and officers owe two fiduciary duties. The duty of care requires an informed decision; a breach is gross negligence, and a charter may exculpate directors (and, for direct claims since 2022, officers) from damages for it (8 Del. C. § 102(b)(7)). The duty of loyalty requires acting in the corporation's interest and in good faith, and cannot be exculpated. Good faith is not a third duty: "the obligation to act in good faith does not establish an independent fiduciary duty that stands on the same footing as the duties of care and loyalty." (Stone v. Ritter (2006).)

Oversight duty. A real duty, but not a third one. It is the part of the duty of loyalty that requires directors to make a good-faith effort to see that the corporation has a reasonable system for bringing information about legal compliance to the board, and to respond to what it reports. Stone v. Ritter placed it inside loyalty: the duty of loyalty "is not limited to cases involving a financial or other cognizable fiduciary conflict of interest. It also encompasses cases where the fiduciary fails to act in good faith." That placement is why exculpation does not reach it and why the standard is bad faith. Also called the Caremark duty; a suit alleging its breach is a Caremark claim. (Caremark (1996); Stone v. Ritter (2006); Marchand (2019).)

Information-systems claim (the first prong). A claim that the directors "utterly failed to implement any reporting or information system or controls" for a central compliance risk. Marchand, Boeing (2021), Abbott and Wells Fargo were sustained on this theory. (Stone v. Ritter (2006).)

Red-flags claim (the second prong). A claim that the directors, "having implemented such a system or controls, consciously failed to monitor or oversee its operations," so that warnings reaching the board went unaddressed. Clovis, Chou and Regions were sustained on this theory. (Stone v. Ritter (2006).)

Good-faith effort. What the first prong requires. The board "must make a good faith effort—i.e., try—to put in place a reasonable board-level system of monitoring and reporting." It does not require that the system catch every failure: "Caremark does not demand omniscience." (Marchand (2019); Boeing (Del. Ch. Aug. 13, 2026).)

Bad faith. The mental state an oversight claim must plead: a conscious or knowing failure to act, not carelessness. Because a Delaware charter may exculpate directors for breaches of the duty of care, only a claim of bad faith survives. "Bad oversight is not bad-faith oversight." (8 Del. C. § 102(b)(7); Stone v. Ritter (2006); Richardson v. Clark (MoneyGram) (2020).)

Officer oversight duty. Officers owe the same duty within their area of responsibility. The Chief Legal Officer "is responsible for legal oversight and for making a good faith effort to establish reasonable information systems to cover that area." An officer must "make a good faith effort to monitor central compliance risks within her remit that pose potential harm to the company or others." (McDonald's (Del. Ch. Jan. 26, 2023); Segway v. Cai (2023).)

THE KINDS OF RISK

Compliance risk. The risk that the company fails to meet a legal or regulatory requirement: a statute, a regulation with the force of law, a consent order, or court-made law such as negligence or defamation ("the common law is just as much 'the law' as statutory law"). Agency guidance and industry frameworks are context, not requirements (SEC guidance "does not establish positive law"). (Fox (2024); SolarWinds (2022).)

Business risk. The risk that a lawful business decision turns out badly. No court has imposed oversight liability for failing to monitor it, but it is not outside the board's job: Caremark's reporting system was one that lets the board reach informed judgments about "business performance" as well as compliance. "Absent statutory or regulatory obligations," how much to spend against a threat "requires an evaluation of business risk, the quintessential board function." "Business risks are shades of gray; legal compliance risks are black and white." (Citigroup (2009); SolarWinds (2022); ProAssurance (2023); B. Riley (2026).)

Business judgment rule. The presumption that directors making a business decision acted on an informed basis, in good faith, and in the honest belief that the decision was in the company's best interests (Aronson v. Lewis (1984)). A court that applies it will not second-guess the decision. A plaintiff rebuts it by pleading a breach of care (an uninformed decision, judged by gross negligence and usually exculpated) or of loyalty (a conflict, or bad faith). The rule protects a decision, including a decision about how much to monitor a business risk; it does not protect the absence of one. It is a standard of review, not a kind of risk: this paper's outer circle is "business risk," and the rule is why it is safe. (Aronson v. Lewis, 473 A.2d 805 (Del. 1984); Citigroup (2009); McDonald's (Del. Ch. Mar. 1, 2023).)

Central compliance risk. A compliance risk that is central to the company's business. This is the category the oversight duty attaches to: the board must have a reporting system for it. "Essential and mission critical risks necessarily qualify as central compliance risks," but not every central compliance risk is mission critical. (McDonald's (Del. Ch. Mar. 1, 2023); Walton (Del. Ch. Apr. 12, 2023).)

Mission-critical risk. A compliance risk in an activity that is the company's core product, service or purpose, whose failure would harm the people the company exists to serve. A subset of central compliance risks. The label changes three things: oversight must be "more rigorously exercised"; a court looks past the existence of a system to "some indicia of effectiveness"; and an inference of bad faith from an unaddressed red flag is "easier to draw." (Marchand (2019); Clovis (2019); Chou (2020); Wells Fargo (N.D. Cal. 2024); Teligent (2025); McDonald's (Del. Ch. Mar. 1, 2023); Segway (2023).)

Corporate trauma. A loss to the company caused by the oversight failure: fines, settlements, lost revenue, or physical or financial harm to the people it serves that the company must answer for. Without one there is no claim; it "is a predicate to a Caremark claim." (Plug Power (Del. Ch. May 2, 2025); Fox (2024).)

Red flag. Information reaching the board, or that a reasonable system would have brought to it, indicating a violation of law or a serious risk of one in a central compliance area. "If a red flag concerns a central compliance risk, then it is easier to draw an inference that a failure to respond meaningfully resulted from bad faith." (McDonald's (Del. Ch. Mar. 1, 2023); Clovis (2019); Chou (2020).)

Yellow flag. The Boeing (2026) court's term for reports of general operational risk that reached the board with management's response attached. They show a system working, not a board ignoring it, and do not support an inference of bad faith. (Boeing (Del. Ch. Aug. 13, 2026).)

THE SYSTEM

Reporting system. Marchand's "reasonable board-level system of monitoring and reporting." In the cases it has four parts: a committee, or the full board, with direct responsibility for the risk; a regular schedule; a protocol that requires management to report, so that reporting is not at management's option; and an escalation path for bad

news, with a response when it arrives. (Marchand (2019); Boeing (2021); Wells Fargo (2024); Teligent (2025). Part IV gives the authority for each.)

Charter; chartered. A committee's written statement of its responsibilities, adopted by the board. A risk is chartered when a committee's charter names it. Courts read the charters first: "every committee charter was silent as to airplane safety"; "the Board had no committee charged with direct responsibility to monitor fair lending compliance." (Boeing (2021); Wells Fargo (2024).)

Indicia of effectiveness. Evidence that a system operated in fact rather than on paper. A court "must look beyond the mere existence of a system to some indicia of effectiveness when determining whether a board made the required good faith effort." "Nominal" compliance with a regulator's rules is not a board-level system. (Teligent (2025); Marchand (2019).)

Reliance. Directors may rely in good faith on reports from officers, employees, committees and experts. Reliance is not a substitute for a system: Caremark "envision[s] some degree of board-level monitoring system, not blind deference to and complete dependence on management." (8 Del. C. § 141(e); Hughes v. Hu (Kandi) (2020).)

Inventory. Management's list of every AI use in the categories the board has named, with an owner for each. A use not on the list is a use the board did not consider. (Center term.)

Tolerance; escalation threshold. A tolerance is a board-ratified limit on a measurable output of a mission-critical use, reported against on schedule; for a coverage-determination model, for example, the rate at which adverse determinations are reversed on appeal. An escalation threshold is the event or measurement that requires management to bring a matter to the committee chair between scheduled reports. (Center terms.)

THE LITIGATION

Derivative suit. A suit by a stockholder on the corporation's behalf against its directors or officers. Any recovery goes to the company. Every oversight decision in Appendix C but one (Teligent, brought directly by a bankruptcy plan administrator) arose in a derivative suit.

Demand; demand futility. Before suing derivatively, a stockholder must ask the board to sue or plead that asking would be futile because a majority of the directors face a substantial likelihood of liability or are not independent. Most oversight cases are decided at this step, on a motion to dismiss. (Court of Chancery Rule 23.1; UFCW v. Zuckerberg (2021).)

Pleading stage; sustained; dismissed. Oversight cases are decided almost entirely on motions to dismiss, where the court takes the complaint's well-pleaded facts as true. A claim that survives is sustained and proceeds to discovery, and in practice to settlement; one that does not is dismissed. No decision in Appendix C imposed liability after a trial.

Section 220 demand. A stockholder's statutory right to inspect the corporation's books and records, including board and committee minutes, before suing. The minutes a court reads in an oversight case usually arrive this way, which is why the record the program keeps is the board's evidence. (8 Del. C. § 220; Marchand (2019).)

THE TECHNOLOGY

AI use. For this paper, a deployment of a system whose outputs (predictions, scores, classifications, recommendations, decisions or generated content) make or materially inform a decision, action or statement of the company. The classification does not depend on whether the system is "artificial intelligence" in a technical sense; a proprietary predictive algorithm and a generative model are analyzed the same way. Where a statute defines the term, the statute's definition governs the compliance analysis under Question 3. (Center term. Statutory definitions: Colorado SB 26-189 (automated decision-making technology); Regulation (EU) 2024/1689, art. 3(1); Texas HB 149.)

Model; third-party model. A model is a quantitative method that turns input data into estimates; a third-party model is one the company licenses or receives from a vendor or affiliate rather than builds. The federal banking agencies' revised model-risk guidance keeps vendor models inside model risk management even where a bank cannot fully validate them ("the principles of model risk management remain applicable"). A vendor's assurances about a model are a control to be tested, not a report to be relied on. (Revised Guidance on Model Risk Management (Fed SR 26-2; OCC Bulletin 2026-13) (Apr. 17, 2026); Hughes v. Hu (Kandi) (2020).)

Generative AI; agentic AI. Generative systems produce text, images, code or other content; agentic systems plan and take actions across tools with limited human intervention. The 2026 interagency model-risk guidance excludes both: "Generative AI and agentic AI models are novel and rapidly evolving. As such, they are not within the scope of this guidance." A board's own system therefore has to reach them. (Revised Guidance on Model Risk Management (Apr. 17, 2026).)

Model drift; secondary red flags. Model drift is the decline in a model's performance as the data it sees departs from the data it was built on. It sends no letter. Commentators call the signals that stand in for one "secondary red flags": model drift, performance degradation, bias amplification, declining alert rates, and divergences between a model's outputs and external indicators such as regulatory inquiries, enforcement actions and whistleblower complaints. The Center adds a rising reversal rate on appeal and divergence from external benchmarks. The reporting protocol has to surface them. (Matera, CLS Blue Sky Blog (Mar. 10, 2026); Center addition as marked.)

PART I

What the law requires, and why the board answers first

A court decides whether a risk was central to the company only after something has gone wrong, from the record the board's system kept. A board with a chartered process that has already asked the question sets the terms of that review.

THE DUTY

Why it matters. The duty has three authors. *Caremark* (1996) held that directors must make a good-faith effort to assure that a reporting system exists, one that lets the board reach "informed judgments concerning both the corporation's compliance with law and its business performance." *Stone v. Ritter* (2006) made it a two-part test, and the two parts name the two kinds of claim: liability follows where directors "utterly failed to implement any reporting or information system or controls" (the information-systems claim), or, having one, "consciously failed to monitor or oversee its operations" (the red-flags claim). Delaware directors owe two fiduciary duties, care and loyalty, and Stone placed the oversight duty inside loyalty, as a failure to act in good faith. That placement does the work: corporate charters exculpate directors for breaches of care, so an oversight claim survives only as a claim of bad faith, which is why the standard is so demanding and why the record of a good-faith effort matters so much. *Marchand* (2019) gave the duty teeth. The Delaware Supreme Court held that a complaint adequately alleged that the board of an ice-cream maker with one line of products had no board-level system at all for food safety after a listeria outbreak killed three people: no committee, no protocol, no schedule. The court called food safety "essential and mission critical" and did not define the phrase. Most of the sustained oversight cases since have been arguments about whether some other risk belongs in the same category.

Two things that are easy to get wrong. First, the duty does not depend on a proven violation. *Marchand* said so: "the fact that Blue Bell nominally complied with FDA regulations does not imply that the board implemented a system to monitor food safety at the board level." What the court asked was whether food safety was a compliance risk central to the business, whether the board had any system for it, and whether harm followed. It called food safety "the obviously most central consumer safety and legal compliance issue facing the company." Second, the category is compliance, not importance. No Delaware court has sustained an oversight claim for a risk that is only a business risk, however central: not for Citigroup's subprime exposure (2009), and not for SolarWinds' cybersecurity (2022), which the court called "mission critical" and still treated as "an evaluation of business risk, the quintessential board function" in the absence of a governing rule.

The two terms that do the most work. A *compliance risk* is the risk that the company fails to meet a legal or regulatory requirement: a statute, a regulation with the force of law, a consent order, or court-made law such as negligence or defamation (*Fox*: "the common law is just as much 'the law' as statutory law"). Agency guidance and industry frameworks are context, not requirements (*SolarWinds*: SEC guidance "does not establish positive law"). A *reporting system* is *Marchand*'s "reasonable board-level system of monitoring and reporting."

In the cases it has four parts: a committee, or the full board, with direct responsibility for the risk; a regular schedule; a protocol that requires management to report, so that reporting is not at management's option; and an escalation path for bad news, with a response when it arrives. Part IV gives the authority for each.

THE MAP: THREE CIRCLES

The map. In *McDonald's* (March 1, 2023), the Court of Chancery distinguished "central compliance risks" from the "essential and mission critical" subset, and placed decisions about monitoring anything else under the business judgment rule; *Walton* (2023) and *Plug Power* (2025) restated the distinction. This paper draws those distinctions, together with the business-risk line from *Citigroup* and *SolarWinds*, as three concentric circles. The operative category is the middle one. In the court's words, "the phrase 'essential and mission critical' deploys more intense terms to explain why those risks qualified as 'central compliance risks,'" and a plaintiff need not plead mission criticality to state a claim. What the label changes is the inference: "if a red flag concerns a central compliance risk, then it is easier to draw an inference that a failure to respond meaningfully resulted from bad faith," and where regulators govern a mission-critical operation the board's oversight "must be more rigorously exercised" (*Clovis*).

	WHAT SITS CIRCLE THERE	WHAT THE BOARD OWES	DETERMINATION
Outer	Business risk, and compliance risk that is peripheral to the business	Oversight by business judgment: the board decides how it will be informed, and records the decision. Courts will not second-guess it: FedEx's contraband-cigarette shipments were a compliance matter, but "an infinitesimal fraction of the overall business" (2021).	Determination 1 or 2 answered no
Middle	Every compliance risk that is central to the business: the "central compliance risks"	A reporting system. This is where the duty attaches.	Determination 2
Center	The subset where the regulated activity is the company's core product, service or purpose and a failure would harm the people it serves: "mission critical"	The same system, which must exist in fact and not on paper; oversight "more rigorously exercised"; an easier inference of bad faith from an unaddressed red flag	Determination 3

The circles are drawn in the executive summary.

WHERE THE QUESTION IS DECIDED

A court sorts a risk into a circle at the pleading stage, usually on a motion to dismiss for failure to make a demand on the board, where the question is whether a majority of the directors face a substantial likelihood of liability. It answers from facts the company itself created: its business model, its regulators, its filings and charters, and what its board's system received and did. The board is the first mover whether it knows it or not. And for AI the gap between what companies say and what boards do is already on the public record. Eighty-three percent of S&P 500 companies disclosed AI as a risk in their filings in 2025 (The Conference Board); 22 percent disclosed board oversight of AI (ISS-Corporate). The two figures come from different studies and measure different things, and that is the point: a plaintiff pleads the first and asks what the minutes show about the second. The company's own filings are Question 5.

READERS OUTSIDE THE PATTERN: OTHER STATES, AND OFFICERS

Two notes for readers outside the pattern. Most of the paper's authority is Delaware's, and most states follow it on oversight: the *Abbott* court applied Illinois law, which "follows Delaware law in establishing demand futility requirements," and then applied *Marchand* and *Boeing* to the substance; a California court adopted the *Caremark* framework in *Kanter v. Reed* (2023). A board incorporated in Texas or Nevada, which have diverged from Delaware since 2024, should have counsel confirm the standard. Second, the officers who will read this paper have their own exposure. *McDonald's* (January 26, 2023) held that officers owe the oversight duty within their area of responsibility, and used the chief legal officer as its example; *Segway* (2023) confirmed the duty reaches central compliance risks "within her remit that pose potential harm to the company or others." The record this paper describes protects the general counsel and the chief risk officer personally as well as the board.

AUTHORITY

Marchand v. Barnhill, 212 A.3d 805 (Del. 2019): the board must "make a good faith effort—i.e., try—to put in place a reasonable board-level system of monitoring and reporting"; "food safety was essential and mission critical." *In re McDonald's Corp. Stockholder Derivative Litig.*, 291 A.3d 652 (Del. Ch. Mar. 1, 2023): "Although it is fair to infer that all 'essential and mission critical risks' qualify as 'central compliance risks,' it is also possible that some 'central compliance risks' may not reach the level of 'essential and mission critical'"; "Outside of central compliance risks, including essential or mission critical risks, a plaintiff will have difficulty rebutting the business judgment rule where officers or directors have made a good faith decision regarding the level of monitoring resources, if any, to assign to a risk" (quoted in *In re Plug Power Inc. Stockholder Derivative Litig.* (Del. Ch. May 2, 2025)). *In re Clovis Oncology* (Del. Ch. Oct. 1, 2019): "when a company operates in an environment where externally imposed regulations govern its 'mission critical' operations, the board's oversight function must be more rigorously exercised." *SolarWinds* (Del. Ch. Sept. 6, 2022): "[n]o case in this jurisdiction has imposed oversight liability based solely on failure to monitor business risk," though an "extreme hypothetical" might one day do so.

PART II

The five questions and the three determinations

No court has announced a test. These five questions are the Center's distillation of what the courts have relied on, in the order the courts reach them. They are evidence. The three determinations at the end of this Part are the conclusions a board draws from them: Questions 1, 2 and 4 decide the determinations; Questions 3 and 5 decide how hard a court will look, and what it will quote back.

THE FIVE QUESTIONS

1. Is this a compliance risk, or a business risk? The threshold. A compliance risk is the risk of failing to meet a legal or regulatory requirement, and it is the only kind of risk the oversight duty has reached. A soured investment or a missed forecast is a business risk, and the business judgment rule protects the board's decision about it; the closing section of this Part says what that means for a business risk the company could not survive. The company need not have violated anything; the question is what kind of failure this would be. One conversion to watch: a business bet stays a business risk, but the disclosure of it is governed by securities law. That is where AI-washing lives.

What counts as a legal requirement. It is broader than a statute and narrower than everything that matters. The sources, with the cases that applied each:

SOURCE OF THE REQUIREMENT	FOR AN AI USE, WHAT IT LOOKS LIKE	WHERE THE COURTS HAVE APPLIED IT
A statute or regulation governing the activity	Coverage rules for a plan's determinations; fair-lending law for a credit model; product-safety regulation for a device; employment law for a screening tool	<i>Marchand</i> (FDA); <i>Abbott</i> (FDA); <i>Wells Fargo</i> (fair lending); <i>Chou</i> (FDA)
A consent order, deferred prosecution agreement or settlement the company operates under	The terms are the company's own rule, and the regulator is already watching	<i>Walton</i> (DEA settlement); <i>Regions</i> (CFPB consent order)
Court-made law: negligence, defamation, product liability	Liability for what a model says or does to a third party	<i>Fox</i> (defamation)
Disclosure law: Regulation S-K Items 105 and 407(h), Rule 10b-5, Delaware's duty of candor	A business bet becomes a compliance risk the moment it is described to investors; so does a statement that the board oversees AI	<i>CVS</i> (2026); the SEC's AI-washing orders (2024); <i>Malone v. Brincat</i>

SOURCE OF THE REQUIREMENT	FOR AN AI USE, WHAT IT LOOKS LIKE	WHERE THE COURTS HAVE APPLIED IT
Law that attaches to inputs: copyright, biometric and consumer-privacy statutes	Training data, fine-tuning data, and the personal data a model consumes	Adobe, Microsoft and Nvidia complaints (alleged; untested)
A rule requiring the board itself to oversee risk (some regulated sectors)	The obligation to have a system is the requirement; a business risk the board never saw is a compliance failure at one remove	Counsel confirms whether one applies
Not a legal requirement: strategy, capital allocation, pricing, competitive position, vendor concentration, and ethical or reputational harm with no rule behind it	Business risk, the outer circle, unless one of the routes above reaches it	<i>Citigroup</i> ; <i>SolarWinds</i> ; <i>ProAssurance</i> ; <i>B. Riley</i>

The company's own policies are not law, but a court notes whether one was broken when it weighs bad faith. In *Skechers*: "The Company was not violating an internal policy or any regulations, which can support an inference of bad faith."

2. How central is it: is the regulated activity the company's core product, service or purpose? This is the first half of what "essential and mission critical" means; the second half is Question 4. *Marchand* asked whether the compliance issue was "intrinsically critical to the company's business operation," and pointed to "the only product it makes." The test is not confined to one-product companies: *Chou* applied it to a diversified drug distributor because drug-safety law went to "the central purpose of ABC's business"; *Wells Fargo* applied it to fair lending at a bank with many lines of business; *Teligent* to a manufacturer whose FDA approval meant "if no products, then no revenue." The indicators are the product or service the company is known for, a decision that *is* the business, the share of revenue or of decisions in the business line that depend on the activity, and whether it could be stopped without stopping the business. A license or approval the business depends on is one such indicator, not the test: a company can operate lawfully while the activity at its center fails the people it serves. A significant compliance risk that is not central is the middle circle. A risk "contained" to a discrete group (*Skechers*) or "an infinitesimal fraction of the overall business" (*FedEx*) is the outer.

3. Which rule governs it, and who enforces it? The statute, regulation, consent order or common-law duty, and the regulator. Where "externally imposed regulations govern its 'mission critical' operations, the board's oversight function must be more rigorously exercised" (*Clavis*). A named regulator is also who sends the notices a board's system must catch: at Blue Bell, "official notices of food safety deficiencies" went unseen at the board for years.

4. What harm follows a failure, to whom, and at what scale? Every oversight claim requires harm to the company (*Plug Power*: "corporate trauma" is "a predicate to a Caremark claim"). Harm to people outside the company, to "the company or others" (*Segway*), has marked every mission-critical case and none of the declined ones: three deaths at Blue Bell, 346 at Boeing, infants at Abbott, borrowers at Wells Fargo, patients at AmerisourceBergen; against a tax gross-up at Skechers, an investment at B. Riley, a restatement at Kandi. With Question 2, it classifies the risk. It also drives what the courts describe in "more intense terms" (*McDonald's*): oversight that "must be more rigorously exercised" (*Clovis*) and an inference of bad faith that is "easier to draw." For an AI use, record scale and reversibility as well: how many decisions the system makes or informs in a period, how many people they reach, and whether an erroneous output can be undone. A model's errors are systematic, and a systematic error is what a plaintiff pleads.

5. What has the company said about it, including about the board's oversight of it? Risk factors, proxy statements, marketing, charters, minutes. A company's own words can put a risk in the middle circle at the pleading stage. Summarizing the *McDonald's* record, the *Walton* court wrote: "When a company makes statements like the examples above, it becomes reasonably conceivable that protecting employees is a central compliance risk." A statement that the board oversees AI is a representation the minutes must support. If they do, it is the board's best evidence. If they do not, it is the plaintiff's best exhibit, and possibly a securities claim of its own.

WHAT THE COURTS HAVE SAID, QUESTION BY QUESTION

Q A COURT THAT SAID YES

A COURT THAT SAID NO

1 <i>Fox</i> (2024): defamation liability; "the common law is just as much 'the law' as statutory law." <i>B. Riley</i> (2026): the doctrine "addresses failures to monitor internal corporate compliance with positive law."	<i>ProAssurance</i> (2023): "Business risks are shades of gray; legal compliance risks are black and white." <i>B. Riley</i> : credit and valuation warnings "are quintessential business risks." <i>Citigroup</i> (2009): no oversight liability for failing to monitor business risk.
2 <i>Marchand</i> : a "monoline company" making one product. <i>Chou</i> (2020): drug-safety law goes to "the central purpose of ABC's business." <i>Teligent</i> (2025): "if no products, then no revenue."	<i>Skechers</i> (2024): the risk "was contained; it was limited to the use of two corporate assets by a discrete group of individuals." <i>FedEx</i> (2021): "an infinitesimal fraction of the overall business."
3 <i>Wells Fargo</i> (N.D. Cal. 2024): "fair lending compliance is a mission critical risk to [a] bank like food safety compliance is to an ice cream company." <i>Abbott</i> (N.D. Ill. 2024): product safety is "externally regulated and 'essential and mission critical.'"	<i>SolarWinds</i> (2022): "Absent statutory or regulatory obligations, how much effort to expend to prevent criminal activities by third parties against the corporate interest requires an evaluation of business risk, the quintessential board function."
4 <i>Marchand</i> : three deaths. <i>Boeing</i> (2021): 346 deaths. <i>Abbott</i> : infants. <i>Wells Fargo</i> : borrowers denied on race.	<i>Skechers</i> : a tax gross-up on aircraft use; no one outside the company harmed. <i>Plug Power</i> (2025): plaintiffs "fail to allege any type of corporate trauma."

Q A COURT THAT SAID YES**A COURT THAT SAID NO**

5 *Boeing*: "every committee charter was silent as to airplane safety." *Wells Fargo*: "Defendants do not dispute" that fair lending was mission critical. *Plains* (2020): the CEO's testimony that pipeline integrity was "not discussed at the board level."

Marriott (2021): the board "consistently ranked cybersecurity as a primary risk," which showed attention, not neglect. *Boeing* (2026): the plaintiffs' theory risked "recasting the volume and depth of Boeing's reporting from a best practice into evidence of disloyalty."

THE THREE DETERMINATIONS

The five questions are evidence. The conclusions a board draws from them are three, made in order, and each has its own consequence.

First: compliance risk, or business risk? (Question 1.) A business risk goes to the outer circle. No Delaware court has imposed oversight liability for failing to monitor one. How much to monitor it is a business decision protected by the business judgment rule, and the record is that decision: what the board decided it would be told, how often, and against what limits. The section below on the business risk that could sink the company says why that record is not a formality. A compliance risk continues.

Second: central to the business, or peripheral? (Question 2, with Questions 3 and 5 as evidence of centrality.) A peripheral compliance risk, one "contained" to a discrete group (*Skechers*) or "an infinitesimal fraction of the overall business" (*FedEx*), joins the outer circle. A central one is the middle circle: the board must have a reporting system for it. This is the determination that creates the obligation, and it does not depend on the third. A plaintiff need not plead mission criticality to state a claim (*McDonald's*). "Central" means a regulated area the company necessarily operates in, at scale, across the business: harassment law for an employer of hundreds of thousands (*McDonald's*), cyber for a hotel chain holding guest data (*Marriott*), overdraft rules for a retail bank (*Regions*).

Third: mission critical? (Questions 2 and 4 together.) A central compliance risk in an activity that is the company's core product, service or purpose, whose failure would harm the people the company serves, is the center. "Core" means the regulated activity is the product or service itself: ice cream, airplanes, lending, coverage decisions. The obligation is the same system. What changes is how a court reads the record: oversight "more rigorously exercised" (*Clovis*), "some indicia of effectiveness" rather than a system on paper (*Teligent*), and an inference of bad faith from an unaddressed red flag that is "easier to draw" (*McDonald's*).

Questions 3 and 5 do not move a use between circles. Question 3 fixes which rule and regulator the system is built around and how rigorously it will be judged; Question 5 fixes what a court will quote back. A use that fails the first determination needs no further analysis. A use that passes the second has already created the duty, whatever the answer to the third.

THE BUSINESS RISK THAT COULD SINK THE COMPANY

The line above leaves a question every board will ask: what about an AI use that is a business risk, not a compliance risk, and that the company could not survive losing? The answer has three parts.

First, as the law stands, no Delaware court will impose oversight liability for failing to monitor it.

Citigroup said so about subprime exposure that nearly destroyed the company: "Oversight duties under Delaware law are not designed to subject directors, even expert directors, to personal liability for failure to predict the future and to properly evaluate business risk." The rule has held through *SolarWinds*, *ProAssurance* and *B. Riley*. Two qualifications keep it from being a license. *Caremark* itself described a reporting system that lets the board reach informed judgments about "the corporation's compliance with law and its business performance"; the courts have confined the liability to compliance, not the system. And *SolarWinds* left the door ajar: "While no case in this jurisdiction has imposed oversight liability based solely on failure to monitor business risk, it is possible, I think, to envision an extreme hypothetical involving liability for bad faith actions of directors leading to such liability." The path through that door is *Disney's* second definition of bad faith, "intentional dereliction of duty, a conscious disregard for one's responsibilities." No court has taken it. A board that knowingly leaves an existential dependence unexamined is the board that would test it.

Second, the business judgment rule protects a decision. It presumes the board acted on an informed basis. A board that never considered the risk has not made a decision to protect. Exculpation means its directors will not pay damages for the carelessness, but this paper is not written to tell boards what they can survive. For a business risk the company could not survive, the outer circle's "documented decision" is a decision about what the board will be told, how often, against what limits, and by whom, and the reporting that follows is reporting on performance, resilience and alternatives, not on compliance. The record is short. Its absence is what a plaintiff will plead as conscious disregard.

Third, an existential business risk converts into a compliance risk by three routes, and the conversion is what a plaintiff pleads.

- *Disclosure.* For a public company a known material risk belongs in the risk factors (Regulation S-K Item 105), the board's role in overseeing risk is a required disclosure (Item 407(h)), and a misstatement is Rule 10b-5. Neither route is confined to public companies: Rule 10b-5 reaches any purchase or sale of a security, including a private financing, and Delaware's own duty of candor reaches directors who "knowingly disseminate false information that results in corporate injury or damage to an individual stockholder" (*Malone v. Brincat*, 1998), "with or without a request for shareholder action."
- *A rule that requires the board to oversee risk.* In some regulated sectors the obligation to maintain a board-level risk-governance framework is itself a legal requirement. Where it is, a business risk the board never saw is a compliance failure at one remove: the rule required a system that captured it. Counsel should ask whether such a rule applies before Determination 1 is recorded.
- *Insolvency.* If the failure is fatal, the plaintiff arrives with it. Creditors of an insolvent Delaware corporation may bring derivative claims on its behalf (*Gheewalla*, 2007), and a bankruptcy trustee or plan administrator sues directly, which is how *Teligent* reached the Court of Chancery. Exculpation still covers care claims, so

they will plead bad faith and look for a compliance hook. What the board knew and decided is what they will find.

"Material" is a disclosure word; the oversight cases say "central" and "mission critical." A disclosed material risk that the board's minutes never mention is a plaintiff's exhibit on both the securities claim and the bad-faith inference. The paper keeps the two words apart.

What the program does with this is simple. Any AI use the company's revenue or operations cannot run without goes on the inventory, whatever Determination 1 says. The board records what it has decided it will be told about it. If a rule reaches the use, it moves to the middle circle. The duty itself applies to every Delaware corporation, public or private; what changes with public status is only who the plaintiff is likely to be: stockholders and regulators first, creditors if the failure is fatal.

WHERE THE LINE RUNS, AND WHY SETTLEMENTS DO NOT MOVE IT

Why the line runs where it does. A court can measure compliance against an objective standard, the requirement itself. It will not substitute its judgment for the board's on a business bet, and the business judgment rule exists to keep it from trying. That is why every oversight claim sustained at the pleading stage since 2019 rests on a legal requirement, and why no Delaware court has extended the duty to a business risk, however large. Cybersecurity shows the line at work. *SolarWinds* called it "mission critical" for an online service provider and still dismissed: the board had a system, no red flag reached it, and without a governing rule how much to spend against criminals was a business judgment. *Marriott* called it "a central compliance risk deserving of board level monitoring at companies across sectors" and dismissed for the same reasons. A pipeline outage with no rule behind it is business risk; after the TSA's security directives of May and July 2021 it is compliance. The same is now happening to AI. Ethical or reputational harm with no governing rule sits in the outer circle until a rule reaches it, and rules are reaching it: Colorado's re-enacted act on automated decision-making technology (SB 26-189, effective January 1, 2027), Illinois's employment-AI law and Texas's Responsible Artificial Intelligence Governance Act (both effective January 1, 2026), New York City's hiring-tool law, the NAIC's insurer bulletin in 25 states and the District of Columbia, and the EU AI Act, whose high-risk obligations the July 2026 Digital Omnibus deferred to December 2027 and August 2028. A December 2025 executive order set out to challenge such state laws; in the xAI challenge to Colorado's first act, in which the Department of Justice intervened, a federal court entered a stipulated order staying enforcement, and the legislature replaced the act. The map will keep moving. The board's system has to keep up with it.

Settlements confirm the line rather than narrow it. Because most oversight cases settle, the decided cases could be a biased sample. The objection is fair, and the settlements answer it. Every eight-figure oversight settlement in Appendix C followed a claim a court had sustained on a compliance theory: Blue Bell (\$60 million), Boeing (\$237.5 million), Walmart (\$123 million), AmerisourceBergen (\$111.25 million), Wells Fargo (\$110 million), Abbott (\$40 million). The ruling created the value. Claims built on an operational or business failure with no governing rule went the other way: the Evolv derivative suit settled for governance reforms and a fee paid by insurers, and the derivative suit over the July 2024 CrowdStrike outage was voluntarily dropped in April 2026 after the companion securities suit was dismissed. Plaintiffs' firms price a theory by whether it

survives a motion to dismiss. An empirical study of every Delaware oversight decision through 2024 found that stand-alone red-flag claims failed in all 28 cases in which they were brought (Shiman, 2026). The compliance line is not an artifact of which cases reach decision. It is the reason the others do not.

PART III

Applying the questions to AI, use by use

AI is a technology, not a risk. It becomes a risk of a particular kind depending on what it is used for. The answer is a list the program maintains, and it is almost never company-wide.

THE USES, SORTED

Why it matters. The *Wells Fargo* court sorted risks within one company, and its reasoning is the template. Fair-lending compliance answered every question it asked: federal law governs it, borrowers are harmed by it, the bank cannot lend without complying with it, and the defendants did not dispute its centrality. Diversity hiring, though "undisputably crucial," was "not precedentially mission critical to its operations like fair lending compliance." One caution: *Wells Fargo* applied Delaware law and *Abbott* applied Illinois law, which "follows Delaware law in establishing demand futility requirements"; neither binds a Delaware court. The table applies the five questions to categories of use; a given system can sit in more than one row. The last column mixes decided cases, pending complaints and settlements, and says which is which.

AI USE	QUESTIONS ANSWERED	DETERMINATION AND CIRCLE	LITIGATION TO DATE (STATUS)
Decisions about people under a regulator's rules where the decision is the product: coverage and prior authorization, credit, underwriting and pricing	1 (compliance: CMS, ECOA, state insurance law), 2 (the decision is the business), 3 (the regulator is named), 4 (patients, borrowers, applicants; at scale), often 5 (the 10-K describes the platform)	Mission critical for insurers, health plans, lenders and any business where the decision is the product	UnitedHealth derivative suit (D. Minn., amended Aug. 7, 2026; pending); <i>Wells Fargo</i> on fair lending (sustained 2024; settled 2026); CVS securities ruling (S.D.N.Y. Aug. 27, 2026; claims over statements about AI tools in post-acute prior authorization proceed; reconsideration motion pending)
Decisions about people under a regulator's rules where the decision is not the product: hiring, promotion, tenant screening	1 (Title VII, Illinois HB 3773, NYC Local Law 144, Colorado SB 26-189), 3, 4 (applicants), 5	Central compliance risk: a reporting system, under the committee with the compliance mandate. Mission critical only where screening or staffing is the business (<i>Wells Fargo</i> : hiring "not precedentially mission critical")	<i>Wells Fargo</i> (hiring claims dismissed 2024); no AI-specific oversight decision

AI USE	QUESTIONS ANSWERED	DETERMINATION AND CIRCLE	LITIGATION TO DATE (STATUS)
AI that controls a physical system or a safety function: driver assistance, medical devices, weapons screening, industrial control	1, 2, 3 (NHTSA, FDA, FTC), 4 (life and injury), 5	Mission critical; the direct analogue of <i>Marchand</i> , <i>Boeing</i> and <i>Abbott</i>	No oversight decision. The Tesla robotaxi suit (W.D. Tex., filed Aug. 2025; pending) and the Evolv derivative settlement (D. Mass.; preliminarily approved Aug. 6, 2026; final hearing Oct. 21, 2026) are disclosure cases about statements made, not oversight of a safety system
AI embedded in compliance and safety monitoring: anti-money-laundering, fraud detection, content moderation, complaint triage	1, 3, 4; question 2 where the underlying compliance is mission critical	As critical as the compliance function it performs	Analogues only: <i>Block</i> (Cash App compliance failures; motion to dismiss denied N.D. Cal. Jan. 6, 2026; pending); <i>Chou</i> ; <i>Richardson v. Clark</i> (MoneyGram)
AI supplied by a vendor or affiliate and used in any of the rows above	The questions of the underlying use; plus reliance	Takes the determination of the use it performs. The board's system has to reach it: contractual rights to performance data, validation and audit, tolerances the vendor reports against, a named internal owner	The UnitedHealth complaint concerns a tool supplied by an affiliate (alleged); the 2026 interagency model-risk guidance keeps vendor models inside model risk management
Training data and model provenance: copyright, biometric and privacy law	1 and 3 (Copyright Act, BIPA, state privacy law; court-made liability counts), 5 (public statements about "licensed" data)	Central compliance risk for any company that trains or fine-tunes; mission critical for model developers	Adobe, Microsoft and Nvidia derivative complaints (Apr. to July 2026; all pending; none pleads <i>Caremark</i> by name; the Microsoft suits consolidated)

AI USE	QUESTIONS ANSWERED	DETERMINATION AND CIRCLE	LITIGATION TO DATE (STATUS)
Statements about AI to investors and customers, including statements about the board's oversight of it	1 and 3 (securities and consumer-protection law), 5	Central compliance risk; owned by disclosure controls and by the officers who make the statements	46 AI-related securities class actions from 2024 through mid-2026 by the Center's count from Cornerstone Research's reports; in the first half of 2026 AI-related filings were 13 percent of filings and 73 percent of the Disclosure Dollar Loss index (Cornerstone); CVS (S.D.N.Y. Aug. 27, 2026) sustained claims that the company concealed its use of AI to deny post-acute care at scale; SEC AI-washing orders (Mar. 18, 2024)
AI capital allocation and strategy: data centers, compute commitments, vendor dependence	None of 1 to 4; only 5	Business risk. No oversight liability; the board decides how it will be informed and records it. A dependence the business cannot run without goes on the inventory (Part II)	<i>Citigroup, ProAssurance, B. Riley</i> ; the Oracle, Microsoft and CoreWeave securities suits are disclosure cases
Enterprise productivity AI: copilots, drafting, internal automation	Usually none, unless the use drifts into a regulated decision or a safety function	Business risk with a compliance edge; a <i>SolarWinds</i> -level system suffices	No case yet; one firm flags data leaks from agentic systems as a possible basis for a future derivative complaint

CASE IN POINT

The amended derivative complaint against UnitedHealth's directors (D. Minn., Aug. 7, 2026) alleges that the company cut off post-acute care for Medicare Advantage members "in lockstep with a proprietary algorithm's predictions," that members who appealed "prevailed 99.7% of the time," and that leadership bet most would never appeal. The consumer class action over the same tool alleges a roughly 90 percent reversal rate and a 0.2 percent appeal rate. Run the three determinations on those allegations: coverage decisions are governed by Medicare rules, so this is a compliance risk; deciding coverage is what a health plan does, so it is central; the people harmed are the members the plan exists to serve, and the tool decided at scale, so it is mission critical. If the allegations are proved, this would be the first AI use a court applying Delaware law calls mission critical, and the question will be what the board's minutes show. The complaint says "algorithm," not "AI"; the analysis does not change.

FOUR WORKED EXAMPLES

A regional bank's AI underwriting model for consumer loans. Determination 1: compliance risk (ECOA and Regulation B, the Fair Housing Act, state fair-lending law). Determination 2: central (lending is what a bank does; the model touches every application in the line). Determination 3: mission critical (lending is the core product; borrowers are the people harmed; *Wells Fargo* is on point). The use goes on the risk committee's schedule with tolerances on approval-rate disparities and adverse-action accuracy, reported by system and by protected class.

An industrial manufacturer's résumé-screening tool. Determination 1: compliance risk (Title VII; Illinois HB 3773; New York City Local Law 144 where it hires there; Colorado SB 26-189 from 2027). Determination 2: central (employment law reaches the whole enterprise, at scale) so a reporting system is required, under the committee with the compliance mandate. Determination 3: not mission critical (hiring is not the product; *Wells Fargo*). The record is a chartered reporting line and an annual bias audit where the law requires one.

An engineering firm's document copilot. Determination 1: business risk. No rule governs drafting itself; the compliance edge (confidentiality, privilege, export-controlled data) is a policy question for management. The record is a documented, informed decision about the controls, and the use is revisited if it drifts into a regulated deliverable.

A specialty retailer with no mission-critical use. Demand forecasting: business risk. Marketing personalization: business risk, unless it uses protected data, in which case privacy law makes it a central compliance risk under the audit committee's mandate. A customer-service chatbot: business risk with a consumer-protection edge (FTC Act section 5 and state deceptive-practices law), which becomes a central compliance risk if the chatbot makes commitments the company must honor. The resolution records that no use is mission critical, that two are central compliance risks with a chartered reporting line, and that the rest are business decisions. That record is worth exactly as much as the record of a "yes."

THE RULES QUESTION 3 ASKS FOR, BY USE

Question 3 is answerable only if counsel can name the instrument. The table lists the AI-specific instruments in force or scheduled as of September 2026 for the uses most often at the center; it is a starting point, not a survey, and each entry should be pulled in its current form before it is cited to a board.

USE	GOVERNING INSTRUMENTS (AS OF SEPT. 2026)
Health-plan coverage and prior authorization	42 C.F.R. § 422.101(b)–(c) (Medicare Advantage coverage criteria; CMS-4201-F, April 2023); CMS FAQs on coverage criteria and utilization management (Feb. 6, 2024), stating that an algorithm predicting the length of a post-acute stay cannot be the sole basis for ending coverage and that determinations must rest on the individual patient's circumstances; California SB 1120 (2024), requiring physician review of utilization decisions that use AI
Credit, underwriting and pricing (lenders)	ECOA, 15 U.S.C. § 1691(d), and Regulation B, 12 C.F.R. § 1002.9 (adverse-action notices with specific reasons); the Fair Housing Act; the Fair Credit Reporting Act. The CFPB withdrew its 2022 and 2023 circulars on adverse-action notices for algorithmic decisions on May 12, 2025; the statute and regulation, not the circulars, govern
Insurance underwriting and pricing	NAIC Model Bulletin on the Use of Artificial Intelligence Systems by Insurers (Dec. 4, 2023), adopted in 25 states and D.C. as of Aug. 31, 2026; New York Insurance Circular Letter No. 7 (2024) (July 11, 2024), which places responsibility for the outcomes of AI use with the board and senior management; Colorado SB 21-169 and its implementing regulations
Hiring and promotion	Title VII; Illinois Human Rights Act as amended by HB 3773 (P.A. 103-0804; eff. Jan. 1, 2026); New York City Local Law 144 of 2021 (bias audits and notice; enforced from July 5, 2023); Colorado SB 26-189 (eff. Jan. 1, 2027); Regulation (EU) 2024/1689, Annex III (employment as high-risk; obligations deferred to Dec. 2, 2027)
Medical devices and safety functions	Federal Food, Drug, and Cosmetic Act and 21 C.F.R. Part 820; FDA final guidance on predetermined change control plans for AI-enabled device software functions (Dec. 2024) and draft guidance on lifecycle management and marketing submissions for AI-enabled device software functions (Jan. 2025); NHTSA Standing General Order 2021-01 on crash reporting for automated driving and Level 2 systems (June 2021, as amended); Texas HB 149 (eff. Jan. 1, 2026)
Compliance functions (AML, fraud)	Bank Secrecy Act, 31 U.S.C. § 5318(h), and the AML program rules of the functional regulator; for banks, the Revised Guidance on Model Risk Management (Fed SR 26-2; OCC Bulletin 2026-13; Apr. 17, 2026), which excludes generative and agentic AI from its scope
Training data and provenance	Copyright Act, 17 U.S.C.; Illinois Biometric Information Privacy Act, 740 ILCS 14; state consumer-privacy statutes; court-made liability where a suit is pending
Statements about AI	Regulation S-K Item 105 and Exchange Act Rule 10b-5; FTC Act section 5, 15 U.S.C. § 45; state deceptive-practices law

PART IV

What a mission-critical use requires: the legal floor, and the program above it

The law requires a good-faith effort at a reasonable board-level system, and a response to what it reports. The program the Center recommends goes further, because it makes the good-faith effort provable. Neither requires the board to catch every failure.

WHAT A BOARD CAN SHOW

Why it matters. *Marchand* listed what Blue Bell's board lacked, and the *Wells Fargo* and *Abbott* courts applied that list to a bank and a formula maker. It reduces to four things a board can show. The program in Part V adds a fifth.

- 1 **A committee, or the full board, with direct responsibility for the risk**, named in a charter. In *Boeing* (2021), "every committee charter was silent as to airplane safety"; in *Wells Fargo*, "the Board had no committee charged with direct responsibility to monitor fair lending compliance." One committee can own the whole AI program; what matters is that each mission-critical use is on its charter and its schedule.
- 2 **A regular schedule.** *Marchand*: "no schedule for the board to consider on a regular basis, such as quarterly or biannually." NiSource's safety committee held five formal meetings a year and the case was dismissed. Boeing's 2026 benchmark: the board "met at least bimonthly," airplane safety "was discussed at every meeting," and the safety committee "met at least twenty-three times between January 2022 and July 2024."
- 3 **A protocol that requires management to report**, so that reporting is not at management's option. *Marchand*: "no regular process or protocols that required management to keep the board apprised." In *Boeing* (2021), management's reports "focused primarily on the business impact" of safety events, and the chief compliance officer's risk update "did not identify product safety as a 'compliance risk.'"
- 4 **An escalation process for bad news, and a response when it arrives.** *Boeing* (2021): "the Board did not have a means of receiving internal complaints about airplane safety." The company had internal channels; none reached the board. *Clovis*: the board received reports and "did nothing." *Regions* (2025): the bank had a system, and the claim survived on warnings the board did not act on.
- 5 **A record that the board decided what belongs on the list.** None of the four cases required it, and the law does not; this is the Center's addition, and the reason this Part distinguishes the floor from the program. The inventory and the three determinations are what put each AI use on a charter and a schedule in the first place, and they are the evidence that the board asked the question for every use, not only for the one that later failed. Without them, a board can show the four things above only for the risks it happened to see.

BETWEEN THE FLOOR AND THE CEILING

The floor and the program. Two recent cases mark the range. The minimum is *Teligent* (2025), a generic-drug maker whose board had no committee for FDA compliance. A committee on paper is not enough: the court will "look beyond the mere existence of a system to some indicia of effectiveness," and a board whose reporting practice "allows management to elect to report (or not)" on a central compliance risk has no system at all. But the floor is low. In *SolarWinds*, a charter amendment naming cybersecurity among the major risks and a management briefing to the committee sufficed to defeat the utter-failure theory at the pleading stage. The maximum is *Boeing* (2026). After the January 2024 door-plug incident, the oversight claim was dismissed because the board had a safety committee of directors with engineering and safety expertise, discussed airplane safety at every meeting, and received a safety report from management each time. Fiduciaries "who 'try' to implement and attend to a 'reasonable board-level system of monitoring and reporting' have met their baseline duty," and "Even in the context of 'mission critical' operations, 'Caremark does not demand omniscience.'" In plain terms: the board does not have to catch every failure, treat every yellow flag as red, or become expert in the technology.

Between the floor and the ceiling sits the program this paper recommends: a chartered committee, an inventory, a written classification, a reporting protocol with tolerances, and internal audit over all of it. The law does not require each of those elements, and the specifics in the model resolution are illustrations, not a standard of care. The reason to build above the floor is evidentiary. A board that can produce the resolution, the charter, the schedule and the reports has made its good-faith effort provable across every AI use it has, and the *Boeing* court's warning against "recasting the volume and depth of Boeing's reporting from a best practice into evidence of disloyalty" means a fuller record is a shield, not a liability.

FOUR THINGS AI CHANGES IN PRACTICE

- **Reports must be framed as risk, not results.** The *Boeing* (2021) court treated reports that "focused primarily on the business impact" as evidence of no oversight. For an insurer that means denial, appeal and reversal rates by system; for a lender, adverse-action and disparate-impact metrics.
- **The red flags are different in kind.** A model does not send a whistleblower letter; it drifts. Commentators call the signals that stand in for one "secondary red flags": model drift, performance degradation, bias amplification, declining alert rates, and divergences between a model's outputs and external indicators such as regulatory inquiries, enforcement actions and whistleblower complaints (Matera, 2026). The Center adds two the cases suggest: a rising reversal rate on appeal, and divergence from an external benchmark. The reporting protocol has to surface them, because nothing else will.
- **Reliance must be tested, and it must reach the vendor.** Delaware law lets a board rely in good faith on officers and experts (8 Del. C. § 141(e)), but *Caremark* "envision[s] some degree of board-level monitoring system, not blind deference to and complete dependence on management" (*Kandi*, 2020). A vendor's assurances about a model are a control to be tested, not a report to be relied on. For a model the company does not own, the system needs contractual rights to performance data, validation and audit, tolerances the vendor reports against, and a named internal owner. The federal banking agencies' 2026 model-risk

guidance makes the same point for banks: even where a bank cannot fully validate a vendor model, "the principles of model risk management remain applicable."

- **Errors are systematic, so reporting is by system and population, not by incident.** A human decision-maker's mistakes are individual; a model's are correlated. The reporting protocol for a mission-critical use reports outcomes by system (determinations made, adverse determinations, appeals, reversals, complaints) and by population affected, against the tolerances the board has ratified. An incident log alone will not show the board a pattern until a plaintiff does.

WHERE THE PROGRAM SITS AMONG THE FRAMEWORKS MANAGEMENT ALREADY RUNS

A chief risk officer already runs a framework: model risk management at a bank, the NIST AI Risk Management Framework (AI RMF 1.0, January 2023) or ISO/IEC 42001:2023 at many companies, COSO-style enterprise risk management at every listed one. Those frameworks are management's system. The program in this paper is the board-level layer above them, and the reporting protocol should draw its metrics from them rather than invent new ones. Two cautions. Frameworks are context, not requirements: a court asked whether the board had a system, not whether management followed a standard (*SolarWinds*). And the frameworks do not always reach the AI a board most needs its system to reach. The Revised Guidance on Model Risk Management that the Federal Reserve, OCC and FDIC issued on April 17, 2026, superseding SR 11-7, applies mainly to banking organizations above \$30 billion in assets, "does not set forth enforceable standards or prescriptive requirements," and states that "Generative AI and agentic AI models are novel and rapidly evolving. As such, they are not within the scope of this guidance." A board that assumes its model-risk function covers those systems should ask.

ONE TOLERANCE, AS AN EXAMPLE

For a coverage-determination model at a health plan, the committee might ratify, for each system: the adverse-determination rate by service line; the appeal rate; the rate at which adverse determinations are reversed on appeal, with an escalation threshold at a stated percentage; the share of determinations reviewed by a clinician before issuance; and the number and disposition of complaints and regulator contacts in the period. Management reports each figure on the schedule, against the prior period and the threshold. A breach goes to the committee chair within the number of days the resolution sets. The numbers are the company's to choose. What the record must show is that the board chose them, received them, and acted when they moved.

PART V

Making and recording the determination

THE SIX STEPS

The determination is one output of a chartered program, minuted, and repeated on triggers. Its record is the board's defense whichever way the answer comes out.

- 1 **Charter the program.** The board assigns AI oversight to a committee, amends its charter to say so, sets the escalation path for bad news, and directs that public statements about AI and its oversight match what the program does.
- 2 **Inventory.** Management lists every AI use in the categories in the Recommendation, including uses supplied by vendors and affiliates and any use the company's revenue or operations cannot run without. A use not on the list is a use the board did not consider.
- 3 **Answer the five questions** for each use, in writing, citing the requirement that governs it and the harm it can cause, at what scale (Appendix A).
- 4 **Make the three determinations** for each use, in order, and record each with its reason. Because oversight liability requires bad faith, a documented, good-faith determination is the board's strongest evidence whichever way it comes out.
- 5 **Oversee.** Each mission-critical use goes on the committee's schedule with a reporting protocol, ratified tolerances and escalation thresholds. Central compliance risks go to the audit or risk committee under its compliance mandate; business risks get a recorded decision.
- 6 **Record and revisit.** The resolution in Appendix B records the determinations. It is revisited on a new AI use, a new rule, an incident, a material public statement about AI, or annually, whichever comes first.

WHO DOES WHAT

The board charters the program, makes the determinations, and owns the system for mission-critical uses. The committee named in the charter runs the schedule and receives the reports. The audit or risk committee owns central compliance risks and receives the assurance. The general counsel applies the questions, identifies the governing requirement for each use, and keeps the record. Officers have their own oversight duty for central compliance risks "within her remit that pose potential harm to the company or others" (*Segway*, 2023), and *McDonald's* gives the chief legal officer as its example of an officer responsible for legal oversight; the record protects the general counsel personally as well as the board. The chief risk officer maintains the inventory and the tolerances, and carries each determination into the enterprise risk taxonomy and the risk appetite statement, so that each mission-critical use has an owner and board-approved limits, and the reporting protocol draws its metrics from the model-risk or AI-risk framework management already runs. Internal audit gives independent assurance that the inventory is complete, the questions were applied, the determinations were made and recorded, and the reporting runs as chartered.

TWO OBJECTIONS, AND WHY THE RECORD DECIDES

What if we classify wrong? A good-faith mistake is not bad faith. But a determination never changes the character of a risk: if a use was in fact a central compliance risk and no system covered it, the board is exposed under the first prong regardless of what it called the use. The rule that follows is simple. When in doubt, put the use in the middle circle. The cost is a reporting line; the determination protects only as evidence of good faith.

Does writing it down hand a plaintiff a concession? In *Wells Fargo* the court noted that the defendants did not dispute that fair lending was mission critical. But the label matters only when the system is missing. Every board that had a functioning system and acted on what it reported prevailed at the pleading stage, however central the risk: Marriott on cybersecurity, NiSource on pipeline safety, Corcept on its primary drug, Boeing in 2026 on airplane safety. The resolution in Appendix B therefore records what the board has elected to oversee and how, not a legal conclusion. A plaintiff can characterize silence. A plaintiff cannot do much with a charter, a schedule and a report.

Why it matters. Every oversight case is decided on the record. *Marchand* and *Boeing* (2021) were sustained on minutes that showed nothing, and then settled for \$60 million and \$237.5 million; *Clovis* and *Chou* on minutes that showed warnings the board did not act on. The boards that will lose the first AI oversight cases at the pleading stage are the ones that never built the system at all, whose filings call AI central and whose minutes never discuss it as a risk. None of this requires predicting what AI will do. It requires deciding what the board must know, and building the system that tells it.

QUESTIONS FOR THE BOARD

- Do we have a chartered system for deciding which AI uses are compliance risks, which are central, and which are mission critical, and for overseeing the ones we decide are?
- Which of our AI uses make decisions about a person under a legal or regulatory requirement, control a safety function, perform a compliance function, or depend on data, models or systems we do not own? Which could the business not run without, and what has the board decided it will be told about them?
- For each: is it a compliance risk, is the regulated activity our core product or service, which rule governs it, who is hurt when it fails and how many, and what have we told investors about it?
- What have we said in our filings about the board's oversight of AI, and would our minutes support it?
- For each mission-critical use: which committee's charter names it, how often is it on the agenda, what must management report, how would a model's drift or a rising reversal rate reach us, and what happens when it does?
- For each vendor-supplied model in a mission-critical use: what performance data are we entitled to, who validates it, and who owns it internally?
- If the general counsel were asked tomorrow to produce the memorandum answering these questions, does one exist? Has internal audit tested that the reporting runs as chartered?

APPENDIX A

Determination worksheet

One row per AI use from the inventory. A compliance risk in an activity that is the company's core product, service or purpose, whose failure would harm the people the company serves, is mission critical; a significant compliance risk that lacks one of those features is a central compliance risk; a business risk, or a peripheral compliance risk, is a business risk decision protected by the business judgment rule. When in doubt, choose the middle. Question 5 is recorded in every case because it will be quoted back.

QUESTION	EVIDENCE TO RECORD	ANSWER AND SOURCE
Program	Committee that owns AI oversight; charter language and date; escalation path; owner of the inventory	
Use under review	System name; function; decisions it makes or informs; population affected; volume per period; owner; whether built, licensed or supplied by an affiliate; whether the business could run without it	
1. Compliance risk or business risk?	Whether a failure would be a failure to meet a legal or regulatory requirement (statute, regulation, consent order, or court-made law such as negligence or defamation) or a commercial outcome; the legal theory a plaintiff would plead; whether the disclosure of the use is itself governed by law	
2. How central?	Whether the activity the use performs or informs is the company's core product, service or purpose: the product or service the company is known for; a decision that is the business; share of revenue or of the decisions in the business line that depend on it; whether the activity could be stopped without stopping the business; whether a license or approval depends on it	
3. Which rule, which regulator?	The statute, regulation, consent order or common-law duty that governs this use (cite it); the regulator; any pending inquiry or notice. Guidance and frameworks noted as context, not requirements	
4. What harm, to whom, at what scale?	Harm to the company from a failure; persons outside the company affected by an erroneous output, and whether they are the people the governing rule exists to protect; nature; scale (decisions per period; population reached); whether an erroneous output can be reversed, and how it would be detected	

QUESTION	EVIDENCE TO RECORD	ANSWER AND SOURCE
5. What have we said?	10-K risk factors, proxy and earnings statements, marketing, committee charters and prior minutes that describe this use, its risk, or the board's oversight of AI	
Determinations	1. Compliance risk or business risk (from Question 1). 2. Central or peripheral (from Question 2, with 3 and 5). 3. Mission critical or not (from Questions 2 and 4). Record each with its reason in one sentence. The placement (outer, middle, center) follows from the three; record it last. For a business risk, record the board's decision about what it will be told, how often, and against what limits	
If mission critical	On which committee's schedule; reporting content and frequency; tolerances ratified; escalation thresholds; accountable executive; for a vendor-supplied model, the contractual rights to performance data, validation and audit, and the internal owner; internal audit review date	
Review triggers	New use; new rule; incident or escalation; material public statement; annual review; date of next review	

APPENDIX B

Model board resolution

The bracketed terms are illustrations. The law requires a good-faith effort at a reasonable system and a response to what it reports; the frequencies, thresholds and reviews below are what makes that effort provable, and a board should set them to fit its own program. The resolution records what the Board has elected to oversee and how. It does not record a legal conclusion, and it does not define "mission critical"; the General Counsel's memorandum (Exhibit B) does.

RESOLUTION OF THE BOARD OF DIRECTORS REGARDING OVERSIGHT OF ARTIFICIAL INTELLIGENCE

WHEREAS management has presented, and the Board has reviewed, an inventory of the Company's material uses of artificial intelligence, attached as Exhibit A, grouped as (a) uses that make or materially inform decisions about individuals subject to regulation; (b) uses that control a physical system or a safety function; (c) uses that perform a compliance or monitoring function; (d) uses that depend on third-party data, content or models, or are supplied by a vendor or affiliate; (e) statements about AI made to investors and customers; (f) uses the Company's revenue or operations cannot run without; and (g) other enterprise uses;

WHEREAS the General Counsel has presented a memorandum, attached as Exhibit B, applying to each such use the questions the Board has adopted for classifying risk: (i) whether a failure of the use would be a failure to meet a legal or regulatory requirement, or a commercial outcome; (ii) whether the activity the use performs or informs is the Company's core product, service or purpose, including whether the activity could be withdrawn without stopping that business; (iii) which statute, regulation, consent order or established common-law duty governs the use, and which regulator or authority enforces it; (iv) what harm would follow a failure, to whom, at what scale, and whether it could be reversed; and (v) what the Company has stated publicly about the use and about the Board's oversight of it; and recording for each use the three determinations the Board has adopted for placing it: whether the use presents a compliance risk or a business risk; whether any compliance risk is central to the Company's business or peripheral; and whether any central compliance risk is mission critical;

NOW, THEREFORE, BE IT RESOLVED that oversight of the Company's use of artificial intelligence is assigned to the [Audit / Risk / Technology] Committee, and that the Committee's charter is amended as set out in Exhibit C to name that responsibility and the uses identified in Exhibit A as mission-critical;

RESOLVED FURTHER that the Board elects to oversee the uses identified in Exhibit A as [list] as mission-critical risks; to oversee the uses identified as [list] as central compliance risks within the Audit Committee's compliance mandate; and to address the uses identified as [list] as business risks through the Board's ordinary decision-making, and, for any such use the Company's revenue or operations cannot run without, directs management to report to the Committee [semi-annually] on its performance, resilience and alternatives and to escalate any failure or material degradation to the Committee chair within [] days;

RESOLVED FURTHER that management shall report to the Committee no less than [quarterly] on each mission-critical use, by system and by population affected, including performance against the tolerances ratified by the Board, [adverse-determination, appeal and reversal rates by system,] incidents, complaints, remediation status and any regulator contact; that [name, title] is the executive accountable for such reporting; that any breach of a ratified tolerance, incident affecting persons outside the Company, whistleblower report or regulator inquiry concerning a mission-critical use shall be escalated to the Committee chair within [] days; and that the Committee shall report to the full Board at each regular meeting;

RESOLVED FURTHER that the Board ratifies the tolerances set out in Exhibit D for each mission-critical use, and directs that any proposal to operate outside them be brought to the Board;

RESOLVED FURTHER that the Chief Risk Officer shall maintain the inventory in Exhibit A, carry each determination into the Company's enterprise risk taxonomy and risk appetite statement, and report to the Committee any new use, material change in an existing use, or change in applicable law that may alter a determination;

RESOLVED FURTHER that, for each mission-critical use supplied by a vendor or affiliate, management shall secure and maintain contractual rights to the performance data, validation and audit access necessary to report against the ratified tolerances, and shall name an internal owner for the use;

RESOLVED FURTHER that Internal Audit shall review [annually], and report to the Audit Committee on, the completeness of Exhibit A, the application of the questions and determinations in Exhibit B, and the operation of the reporting required by this resolution;

RESOLVED FURTHER that the Disclosure Committee and the Audit Committee shall review all statements about AI, and about the Board's oversight of AI, in periodic reports, proxy statements, earnings materials and marketing for consistency with Exhibit A and with this resolution; and

RESOLVED FURTHER that the Board shall revisit these determinations upon any new material AI use, any change in applicable law, any incident or escalation concerning a mission-critical use, any material public statement about AI, and in any event annually, and that Exhibits A through D and the materials supporting these determinations shall be retained with the minutes.

APPENDIX C

The cases behind the five questions

The decisions cited in this paper, June 2019 through September 2026, in date order. Reporter and Westlaw citations are given where confirmed; the rest are cited by docket, court and date and should be pinned to Westlaw or Lexis before use in a filing. **Label** is what the court called the risk. **Q** lists which of the five questions did the work. **Outcome** is the ruling on the oversight claim at the pleading stage, with subsequent history as of September 2026. All were derivative suits except Teligent, brought directly by a bankruptcy plan administrator.

CASE	LABEL	WHAT THE COURT RELIED ON	OUTCOME
<i>Marchand v. Barnhill</i> , 212 A.3d 805 (Del. June 18, 2019)	Mission critical	Monoline ice-cream maker; FDA regulation; three deaths; no committee, protocol, schedule or board-level system (Q1, 2, 3, 4)	Claim sustained; settled (\$60M, 2020)
<i>In re Clovis Oncology, Inc. Derivative Litig.</i> , 2019 WL 4850188 (Del. Ch. Oct. 1, 2019)	Mission critical	Single drug; FDA trial protocol; board received reports and "did nothing" (Q1, 2, 3, 4)	Sustained
<i>Inter-Marketing Grp. USA, Inc. v. Armstrong</i> (Plains All American), 2020 WL 756965 (Del. Ch. Jan. 31, 2020)	Mission critical	Complaint quoting the company's "primary operational emphasis" on pipeline integrity; CEO's criminal-trial testimony that it was "not discussed at the board level" (Q1, 2, 4, 5)	Contract claim against the general partner sustained under a <i>Caremark</i> -style analysis; director claims dismissed
<i>Hughes v. Hu</i> (Kandi), 2020 WL 1987029 (Del. Ch. Apr. 27, 2020)	Not used	Financial reporting; audit committee met "sporadically"; board relied "entirely on management"; the harm alleged was to the company (a restatement)	Sustained
<i>Teamsters Local 443 Health Servs. & Ins. Plan v. Chou</i> (AmerisourceBergen), 2020 WL 5028065 (Del. Ch. Aug. 24, 2020)	Mission critical	Drug health and safety risks called mission critical; drug-safety law goes to "the central purpose of ABC's business"; board never asked for reports on a whistleblower's allegations (Q1, 3, 4)	Sustained

CASE	LABEL	WHAT THE COURT RELIED ON	OUTCOME
<i>Richardson v. Clark</i> (MoneyGram), 2020 WL 7861335 (Del. Ch. Dec. 31, 2020)	Not reached	AML compliance; board and compliance committee met regularly; "Bad oversight is not bad-faith oversight"	Dismissed
<i>Pettry v. Smith</i> (FedEx), 2021 WL 2644475 (Del. Ch. June 28, 2021), aff'd, 273 A.3d 750 (Del. 2022) (TABLE)	Declined	Illegal cigarette shipments and resulting fines "an infinitesimal fraction of the overall business" (n.125); board updated on the litigation on at least eleven occasions (Q2)	Dismissed
<i>In re The Boeing Co. Derivative Litig.</i> , 2021 WL 4059934 (Del. Ch. Sept. 7, 2021)	Mission critical	Airplane safety; FAA; 346 deaths; charters silent; the Board had no means of receiving internal complaints about airplane safety; reports framed as business impact (Q1, 2, 3, 4, 5)	Sustained in part; settled (\$237.5M, approved Feb. 2022)
<i>Firemen's Ret. Sys. of St. Louis v. Sorenson</i> (Marriott), 2021 WL 4593777 (Del. Ch. Oct. 5, 2021)	Not applied; cyber "has increasingly become a central compliance risk"	Board "consistently ranked cybersecurity as a primary risk" and had a system (Q5, for the board)	Dismissed
<i>City of Detroit Police & Fire Ret. Sys. v. Hamrock</i> (NiSource), 2022 WL 2387653 (Del. Ch. June 30, 2022)	Mission critical	Pipeline safety; but a functioning safety committee held five formal meetings a year (Q1, 2, 3, 4)	Dismissed
<i>Constr. Indus. Laborers Pension Fund v. Bingle</i> (SolarWinds), 2022 WL 4102492 (Del. Ch. Sept. 6, 2022), aff'd, 297 A.3d 1083 (Del. 2023) (TABLE)	"Mission critical" in fact; no positive law	SEC guidance "does not establish positive law"; charter reference and a management briefing sufficed against the utter-failure theory (Q1, 3)	Dismissed

CASE	LABEL	WHAT THE COURT RELIED ON	OUTCOME
<i>In re McDonald's Corp. Stockholder Derivative Litig.</i> , 289 A.3d 343 (Del. Ch. Jan. 26, 2023) (officers); 291 A.3d 652 (Del. Ch. Mar. 1, 2023) (directors)	Central compliance; mission-critical subset distinguished	Officers owe oversight within their area of responsibility, the chief legal officer as the example (Jan.); mission-critical risks are a subset of central compliance risks; monitoring outside them is a business judgment; the head of HR's own harassment "the most vibrant of red flags" (Mar.)	Officer claim stated (Jan.); action dismissed for failure to plead demand futility (Mar.)
<i>Ontario Provincial Council of Carpenters' Pension Trust Fund v. Walton</i> (Walmart), 294 A.3d 65 (Del. Ch. Apr. 12, 2023); 2023 WL 3093500 (Del. Ch. Apr. 26, 2023)	Central compliance	Pharmacy controlled-substance compliance; "Essential and mission critical risks necessarily qualify as central compliance risks" (Apr. 12) (Q1, 2, 3, 4)	Sustained in part (Apr. 26); settled (\$123M plus governance terms, approved Dec. 2024)
<i>In re ProAssurance Corp. Stockholder Derivative Litig.</i> , 2023 WL 6426294 (Del. Ch. Oct. 2, 2023)	Declined	Underwriting is business risk: "shades of gray" versus "black and white" (Q1)	Dismissed
<i>Segway Inc. v. Cai</i> , 2023 WL 8643017 (Del. Ch. Dec. 14, 2023)	Officer duty	An officer must monitor central compliance risks "within her remit that pose potential harm to the company or others"; declining sales and receivables are not such a risk	Officer oversight claim dismissed
<i>Lebanon County Emps.' Ret. Fund v. Collis</i> (AmerisourceBergen), 311 A.3d 773 (Del. Dec. 18, 2023)	Central compliance	Opioid distribution under DEA rules. The Court of Chancery had found the directors "were wrapped in" red flags but dismissed on another court's later findings; the Supreme Court held that use of judicial notice improper (Q1, 3, 4)	Dismissal reversed; settled (\$111.25M, 2025)
<i>Conte v. Greenberg</i> (Skechers), 2024 WL 413430 (Del. Ch. Feb. 2, 2024), <i>aff'd</i> , No. 76, 2024 (Del. Jan. 23, 2025) (ORDER)	Declined	Risk "contained" to a discrete group; "relatively minimal magnitude"; no regulation violated (Q2, 3, 4)	Dismissed

CASE	LABEL	WHAT THE COURT RELIED ON	OUTCOME
<i>In re Abbott Labs. Infant Formula S'holder Derivative Litig.</i> , No. 1:22-cv-05513 (N.D. Ill. Aug. 7, 2024)	Mission critical	Product safety "externally regulated and 'essential and mission critical'"; no committee, schedule or protocol; concerns never reached the board (Q1, 2, 3, 4)	Demand excused on the oversight claim; settled (\$40M in plant investment plus reforms, 2026)
<i>In re Wells Fargo & Co. Hiring Practices Derivative Litig.</i> , No. 3:22-cv-05173 (N.D. Cal. Sept. 20, 2024)	Mission critical (fair lending); declined (hiring)	Fair lending "like food safety compliance is to an ice cream company"; hiring "not precedentially mission critical" (Q1, 2, 3, 4, 5)	Sustained in part; settled (\$100M borrower-assistance fund plus \$10M, final approval May 2026)
<i>In re Fox Corp. Derivative Litig.</i> , 2024 WL 5233229 (Del. Ch. Dec. 27, 2024)	Extended to common law	Defamation liability: "the common law is just as much 'the law' as statutory law" (Q1)	Demand excused; claims proceed
<i>In re Plug Power Inc. Stockholder Derivative Litig.</i> , C.A. No. 2022-0569-KSJM (Del. Ch. May 2, 2025)	Restates <i>McDonald's</i>	Insider-trading policy and SEC comment letters; plaintiffs "fail to allege any type of corporate trauma"	Dismissed
<i>Ritchie v. Baker</i> (Corcept), 2025 WL 2048014 (Del. Ch. July 22, 2025)	Not reached	Primary drug; the complaint itself alleged the board was "well-informed" through routine briefings (Q2, 5)	Dismissed
<i>Giuliano v. Grenfell-Gardner</i> (Teligent), C.A. No. 2021-0452-KSJM (Del. Ch. Sept. 2, 2025)	Mission critical	FDA compliance a condition of revenue; no committee; "indicia of effectiveness" required; management cannot elect whether to report (Q1, 2, 3). Direct action by the Chapter 11 plan administrator	Motion to dismiss denied in substantial part
<i>Brewer v. Turner</i> (Regions Financial), C.A. No. 2023-1284-KSJM (Del. Ch. Sept. 29, 2025)	Central compliance	Overdraft practices at a retail bank; the bank had an information system, and the claim survived on warnings the board did not act on (Q1, 3, 4)	Sustained in part (against 9 of 18 directors; officer claims dismissed); interlocutory appeal refused (Del. Dec. 15, 2025)

CASE	LABEL	WHAT THE COURT RELIED ON	OUTCOME
<i>Marchner v. Riley</i> (B. Riley Financial), C.A. No. 2025-0164-LWW (Del. Ch. Mar. 30, 2026)	Declined	Credit and valuation warnings "are quintessential business risks"; the doctrine "addresses failures to monitor internal corporate compliance with positive law" (Q1)	Dismissed
<i>In re The Boeing Co. Derivative Litig.</i> , Consol. C.A. No. 2024-1210-MTZ (Del. Ch. Aug. 13, 2026)	Assumed	Aerospace Safety Committee; safety discussed at every meeting; management report each time; "does not demand omniscience"; "yellow flags" only	Dismissed

SOURCES AND FURTHER READING

Center for AI Oversight: *The Caremark AI Liability Roadmap* (June 2026), which argued that AI is the next mission-critical risk, a claim this paper refines use by use; *The Oversight Layer of AI Governance* (Sept. 2026), the seven elements of an Oversight Program; *The Missing Instrument* (forthcoming), the regulator's role. At cfaio.org/insights.

Foundational authority: Regulation S-K Items 105 and 407(h) (17 C.F.R. §§ 229.105, 229.407(h)) and Exchange Act Rule 10b-5 (17 C.F.R. § 240.10b-5); *In re Caremark Int'l Inc. Derivative Litig.*, 698 A.2d 959 (Del. Ch. 1996); *Stone v. Ritter*, 911 A.2d 362 (Del. 2006); *In re Citigroup Inc. S'holder Derivative Litig.*, 964 A.2d 106 (Del. Ch. 2009); *United Food & Commercial Workers Union v. Zuckerberg*, 262 A.3d 1034 (Del. 2021); *Aronson v. Lewis*, 473 A.2d 805 (Del. 1984); *In re The Walt Disney Co. Derivative Litig.*, 906 A.2d 27 (Del. 2006); *Malone v. Brincat*, 722 A.2d 5 (Del. 1998); *North American Catholic Educational Programming Foundation, Inc. v. Gheewalla*, 930 A.2d 92 (Del. 2007); *Kanter v. Reed*, 92 Cal. App. 5th 191 (2023); 8 Del. C. §§ 102(b)(7), 141(e), 220. The decisions since *Marchand* are in Appendix C.

Other courts and pending matters: *In re Block Inc. Shareholder Derivative Litig.*, No. 25-cv-01262-NW (N.D. Cal. Jan. 6, 2026) (order denying motions to dismiss); *McCollum v. Witty* (UnitedHealth), No. 0:24-cv-02643 (D. Minn.) (verified amended shareholder derivative complaint, redacted public version, Aug. 7, 2026), and *Estate of Lokken v. UnitedHealth Group*, No. 0:23-cv-03514 (D. Minn.) (motion to dismiss granted in part and denied in part, Feb. 13, 2025); *SEIU Pension Plan Master Trust v. Narayan* (Adobe), No. 3:26-cv-03521 (N.D. Cal., filed Apr. 24, 2026); *Anderson v. Nadella* (Microsoft), No. 2:26-cv-02281 (W.D. Wash., filed June 30, 2026), since consolidated as *In re Microsoft Corp. Stockholder Derivative Litig.*; *Berliner v. Huang* (Nvidia), No. 1:26-cv-09153 (N.D. Ill., filed July 31, 2026); *Nixon v. CVS Health Corp.* (Louisiana Sheriffs' Pension & Relief Fund, lead plaintiff), No. 1:24-cv-05303-MMG (S.D.N.Y. Aug. 27, 2026) (securities; motion to dismiss denied in substantial part; reconsideration motion filed Sept. 10, 2026); *In re Evolv Technologies Holdings, Inc. Stockholder Derivative Litig.*, No. 1:24-cv-12822-ADB (D. Mass.) (settlement preliminarily approved Aug. 6, 2026; final hearing Oct. 21, 2026); *Morand v. Tesla, Inc.*, No. 1:25-cv-01213-RP (W.D. Tex., filed Aug. 4, 2025); the consolidated CrowdStrike derivative action, No. 1:24-cv-01031 (W.D. Tex.) (voluntarily dismissed Apr. 15, 2026, after the companion securities action was dismissed); *Barrows v. Oracle Corp.*, No. 1:26-cv-00127 (D. Del., filed Feb. 2026); SEC settled orders against Delphia (USA) Inc. and Global Predictions Inc. (Mar. 18, 2024) (AI-washing). Pending matters are described as alleged.

Rules referred to in Parts II and III: Colorado SB 26-189 (signed May 14, 2026; effective Jan. 1, 2027), repealing and re-enacting the 2024 Colorado AI Act after a federal court entered a stipulated order staying its enforcement in *X.AI LLC v. Weiser*, No. 1:26-cv-01515 (D. Colo.), in which the Department of Justice intervened; Colorado SB 21-169; Illinois HB 3773 (P.A. 103-0804; effective Jan. 1, 2026); Texas HB 149, the Texas Responsible Artificial Intelligence Governance Act (effective Jan. 1, 2026); California SB 1120 (2024); New York City Local Law 144 of 2021 (enforced from July 5, 2023); NAIC Model Bulletin on the Use of Artificial Intelligence Systems by Insurers (Dec. 4, 2023) (25 states and D.C. per the NAIC's map as of Aug. 31, 2026); New York Insurance Circular Letter No. 7 (2024) (July 11, 2024); Regulation (EU) 2024/1689, as amended by Regulation (EU) 2026/1744 (in force July 27, 2026); Executive Order 14365, Ensuring a National Policy Framework for Artificial Intelligence (Dec. 11, 2025); TSA Security Directives Pipeline-2021-01 and -02 (issued May 27 and July 19, 2021); 42 C.F.R. § 422.101 and CMS, FAQs on coverage criteria and utilization management requirements in CMS-4201-F (Feb. 6, 2024); ECOA, 15 U.S.C. § 1691(d), and Regulation B, 12 C.F.R. § 1002.9 (CFPB Circulars 2022-03 and 2023-03

withdrawn May 12, 2025); FDA, Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence-Enabled Device Software Functions (final guidance, Dec. 2024), and Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations (draft guidance, Jan. 2025); NHTSA Standing General Order 2021-01 (as amended); Board of Governors of the Federal Reserve System, OCC and FDIC, Revised Guidance on Model Risk Management (Apr. 17, 2026) (Fed SR 26-2; OCC Bulletin 2026-13), superseding SR 11-7 and OCC Bulletin 2011-12; NIST AI Risk Management Framework (AI RMF 1.0) (Jan. 2023); ISO/IEC 42001:2023.

Commentary quoted or relied on: Roy Shapira, *Mission Critical ESG and the Scope of Director Oversight Duties*, 2022 Colum. Bus. L. Rev. 732; Jennifer Arlen, *Directors' Caremark Liability for Fraudulent Disclosures to Customers about the Company's Cybersecurity: SolarWinds Reconsidered*, 50 J. Corp. L. 1142 (2025); Stephen M. Bainbridge, *Don't Compound the Caremark Mistake by Extending It to ESG Oversight*, 77 Bus. Law. 651 (2022); Pierluigi Matera, *From Red Flags to Black Boxes: Corporate Oversight in the Age of Artificial Intelligence* (SSRN, 2026), summarized as *Corporate Oversight in the Age of Artificial Intelligence*, CLS Blue Sky Blog (Mar. 10, 2026) ("secondary red flags"); Michael A. Rosenblum & Monica J. Roth, *Why "Mission-Critical" Risks Create Liability for Corporate Directors*, Thompson Coburn (Jan. 9, 2026); Ropes & Gray, *AI in the Boardroom, Shareholders in the Courtroom: When AI Strategy Becomes Fiduciary Risk* (Sept. 14, 2026); Yehonatan Shiman, *Oversizing Caremark: An Empirical Analysis* (SSRN 2026), summarized as *How Caremark Has Become Oversized*, CLS Blue Sky Blog (Sept. 2, 2026).

Statistics: The Conference Board, *From Principles to Practice: Governing AI in the Corporation* (Apr. 22, 2026) (83% of S&P 500 companies disclosing AI as a risk, 2025 data); ISS-Corporate, *Artificial Intelligence and Governance: Is 2026 a Tipping Point for Turning Awareness into Action?* (May 20, 2026) (22% of S&P 500 disclosed board oversight of AI, 2025 data); Cornerstone Research, securities class action filings reports (Jan. 29, 2025; Jan. 28, 2026; July 29, 2026) (15, 16 and 15 AI-related filings in 2024, 2025 and the first half of 2026; the total of 46 is the Center's sum).

Cite as: Center for AI Oversight, *Is AI Mission Critical? The Caremark Question for Directors and Officers* (Sept. 2026), cfaio.org. The Center for AI Oversight is a 501(c)(3) nonprofit educational institution defining oversight governance for AI in regulated industries. Center materials are educational and do not constitute legal advice. The Center takes no positions on matters involving specific institutions; pending cases are described as alleged.