



The Center for  
AI Oversight

---

CENTER PAPER NO. 1 · JUNE 2026

# The Caremark AI Liability Roadmap

Board and officer oversight obligations applied to AI governance. Six cases, one standard, and the program that answers it.

By **Brian J. Allen**, Executive Director

Each section of this paper moves in three tiers. A **board line** states the governance proposition in one sentence. A **why it matters** note explains the consequence. A shaded **detail** block carries the case, rule, or record for readers who want the underlying authority. Read the board lines straight down for the argument in brief; drop into the detail where a section bears on your institution. Each Part closes with questions a director can take into the next meeting.

## EXECUTIVE SUMMARY

# The standard, and why it is the program

### KEY TAKEAWAYS

- 1 Delaware's oversight doctrine is procedural. Directors and officers are exposed not when an AI decision goes wrong, but when no **program** existed to make an informed decision possible.
- 2 No court has yet applied the doctrine to AI. That is the reason to build the program now, because the record that decides these cases is written in board minutes before any complaint is filed.
- 3 Oversight liability reaches officers, not only directors, placing the executives closest to AI under personal fiduciary duties within their domains.
- 4 The institution's own filings are assembling the plaintiff's case: 72 percent of the S&P 500 disclose AI as a material risk, while director-level AI expertise sits below 3 percent.
- 5 Each case describes one object seen from a different angle: the oversight program. Building it is the answer to all six. The program is the governance.

Across six cases spanning three decades, *In re Caremark*, *Stone v. Ritter*, *Marchand v. Barnhill*, *In re Clovis Oncology*, *In re Boeing*, and *In re McDonald's*, the Delaware courts sharpened a single question, expanded who must answer it, and described what a satisfactory answer contains. Read together, they do not describe six different duties. They describe one: the duty to maintain a functioning oversight program, viewed from six angles. The doctrine converges on an operating standard this paper names the **Informed Decision Standard**: oversight is demonstrated by the ability to show that informed decisions were possible, not by the outcomes of the decisions themselves.

This paper traces how each case built on its predecessor, extracts the obligations the line establishes, and assembles them into the program that satisfies the standard. Boards that build that program before the first AI-related corporate trauma earn the outcome Marriott's board earned: dismissal, because the record showed a program that was real and actively monitored. Boards that do not are assembling the evidentiary record future plaintiffs will obtain through books-and-records demands under Section 220 of the Delaware General Corporation Law.

---

## PART I

## How six cases built one standard

The Caremark line is not a single moment of legal clarity. It is a thirty-year construction in which each decision tightened the standard, expanded who is covered, or specified what good faith oversight requires. The obligations are cumulative, and they converge on a single object. Each case below is a view of the same thing a board must build: the oversight program.

### 1.1 Caremark (1996): The Affirmative Obligation

The board has an affirmative duty to ensure a reasonable system exists to inform it of the corporation's central risks. The duty attaches to the system, not to any single decision.

**Why it matters.** From the doctrine's first sentence, the unit of analysis has been the program. For AI, the question a court will ask is not whether the board approved the model that failed; it is whether a system existed through which the board could have been informed of the risks that model carried.

#### CASE

*In re Caremark Int'l Inc. Derivative Litig.*, 698 A.2d 959 (Del. Ch. 1996). Chancellor Allen framed the two-prong test: liability where directors (1) utterly failed to implement any reporting system, or (2) consciously failed to monitor one they implemented. The standard is good faith effort, not perfection. The opinion called the theory possibly the most difficult in corporation law on which a plaintiff might win.

### 1.2 Stone v. Ritter (2006): Loyalty, and the End of Exculpation

Oversight failure is a breach of the duty of loyalty, not care. No charter provision can exculpate it, and since 2022 the same is true for officers.

**Why it matters.** Section 102(b)(7) lets a charter exculpate duty-of-care breaches; it cannot reach loyalty breaches. Every director and officer should understand that there is no charter language between them and an oversight claim.

#### RULE

*Stone v. Ritter*, 911 A.2d 362, 370 (Del. 2006). The Delaware Supreme Court adopted the Caremark standard, codified the two-prong test, and held that the good faith requirement is a subsidiary element of the duty of loyalty. The 2022 amendment extending 102(b)(7) to officers preserves the same carve-out for loyalty and bad faith.

### 1.3 Marchand (2019): Mission Critical Risk, and the Duty to Name It

For risks that are mission critical to the institution, the board must rigorously oversee them, and it must first identify which risks those are. A risk the board never named is a risk it never oversaw.

**Why it matters.** A board cannot rigorously oversee a mission critical risk it has not identified, and AI enters the institution through hundreds of channels at once, including vendor software never evaluated as AI. The more

diffuse the risk, the less credible the claim that the board could oversee it without a systematic inventory, the program's first element.

#### CASE

*Marchand v. Barnhill*, 212 A.3d 805 (Del. 2019). The Delaware Supreme Court, en banc (Strine, C.J.), reversed dismissal of a claim against Blue Bell Creameries after a listeria outbreak killed three people. The claim survived at the pleading stage because no committee addressed food safety, no process required management to report it, no schedule existed for board consideration, and red-flag reports never reached the board. The decision endorsed books-and-records inspection under 8 Del. C. § 220 as the plaintiff's tool.

## 1.4 Clovis Oncology (2019): Monitoring Is Not Optional

**A board that receives detailed reports on its mission critical risk and consciously ignores the red flags in them is liable. Reporting is not oversight; acting on what the reports show is.**

**Why it matters.** Clovis is the second-prong case, and it is the more demanding one for AI. The first prong, no system at all, is becoming unavailable to plaintiffs as boards build AI reporting. Future claims will be Clovis claims: the reports existed, reached the board, and the question will be whether the board acted on the red flags they carried. The evidence of governance is the program in operation, not the existence of a reporting line.

#### CASE

*In re Clovis Oncology, Inc. Derivative Litig.*, C.A. No. 2017-0222-JRS (Del. Ch. Oct. 1, 2019) (Slights, V.C.). Clovis, a monoline company whose future rode on the lung-cancer drug Rociletinib, adopted the RECIST protocol for its TIGER-X trial. The board received detailed trial information at every meeting, so the claim failed Caremark's *first* prong. It survived under the *second*: the board consciously disregarded red flags that management was deviating from the protocol and reporting unconfirmed results, placing FDA approval in jeopardy. The disclosure cost roughly \$1 billion in market value.

## 1.5 Boeing (2021): The Anatomy of an Adequate Program

**An adequate oversight program has five parts: a committee that owns the risk by charter, standing agenda time, risk-centric reporting, an independent escalation channel, and a documented response to red flags.**

**Why it matters.** Boeing is the closest thing the doctrine offers to a specification, because the court described the adequate program by cataloguing each element Boeing lacked. These five elements are the anatomy a future AI oversight claim will be measured against, and Part V assembles them into the AI oversight program directly.

#### RECORD

*In re The Boeing Co. Derivative Litig.*, C.A. No. 2019-0907-MTZ (Del. Ch. Sept. 7, 2021) (Zurn, V.C.). Built on more than 44,000 documents produced under Section 220 after the two 737 MAX crashes. No committee charter named airplane safety; reporting carried only good news; the safety reporting process had no link to the board; and the board treated the first crash as a public relations matter rather than investigating. Boeing settled for \$237.5 million, the largest

Caremark cash settlement in Delaware history, funded by D&O insurers, after its motion to dismiss was denied. The court: minimal regulatory compliance is not a reasonable reporting system.

## 1.6 McDonald's (2023): The Program Reaches the Officers

Officers owe the same oversight duties as directors within their domains, including a duty to escalate the red flags they encounter. The program's accountability layer now extends down the org chart.

**Why it matters.** For AI, McDonald's is the front line, developed in Part II. The officers closest to AI activity hold their responsibilities as fiduciary duties; the 2022 amendment does not exculpate them; and an officer who sees a red flag within the domain and sits on it has personally completed the second prong.

### CASE

*In re McDonald's Corp. Stockholder Derivative Litig.*, 289 A.3d 343 (Del. Ch. 2023) (Laster, V.C.). Officers owe oversight duties within their areas of responsibility, with context-driven application: a CEO carries a company-wide remit, while other officers carry the duty within their specific area. The surviving claim ran against the company's chief people officer for failures within the domain he led.

## 1.7 One Standard, One Program

The six cases ask two questions in every season and industry: did a program exist through which the institution could know its central risks, and did the people accountable for it actually use it.

**Why it matters.** The standard is procedural and evidentiary. Courts do not second-guess the decision; they examine whether an informed decision was possible and whether the record proves it. The Center names this the **Informed Decision Standard**. What it requires is not six separate compliance exercises but one functioning program, which is what Part V assembles.

### QUESTIONS FOR THE BOARD

- Does a program exist that is reasonably designed to bring AI risks to our attention, or are we assuming management has it covered?
- Have we determined whether AI is mission critical for this institution, and is the determination documented in our minutes?
- Do all directors and officers understand that an AI oversight failure is a loyalty breach no charter provision exculpates?

---

## PART II

## The officer front line

McDonald's extended the program's accountability layer from the boardroom to the org chart. Every officer who owns a piece of AI now carries a personal fiduciary duty to maintain a system within that domain and to escalate out of it.

**Why it matters.** AI is distributed across more domains than any risk the doctrine has previously reached: risk, security, technology, legal, and the AI function itself. The program's *who* must therefore be answered below the board with the same precision as above it. A program that assigns AI to the board and to no one else has answered half the question and left its officers individually exposed.

Three features of the officer duty deserve precision. It is domain-bounded but not domain-excused: an officer who encounters a red flag of sufficient gravity outside the domain is not free to ignore it. It is loyalty-based, so the 2022 extension of Section 102(b)(7) to officers does not shield it. And it is evidentiary in the same way the director duty is: the officer's protection is the documented program in operation, the reports rendered, the escalations made.

### DISCLOSURE LINK

Regulation S-K Item 106(c)(2) already requires public companies to describe management's role and expertise in managing material cybersecurity risk, so the officer layer of technology-risk governance is a mandated public disclosure in the adjacent domain. When a CEO and CFO certify disclosure controls under Exchange Act Rules 13a-14 and 15d-14, and AI risks are material, the adequacy of AI risk controls sits inside what they are certifying.

The officer corps is where AI oversight will be tested first, because officers see the red flags first. McDonald's means the red flag an officer absorbs is the officer's personal one.

### QUESTIONS FOR THE BOARD

- For each officer who owns a piece of AI, is there a defined system within the domain and a documented duty to escalate?
- Does our disclosure of management's role in technology risk reflect who actually owns AI, or a structure that exists only on paper?

---

## PART III

## The protection proof: what the program earns

The doctrine protects a functioning program as reliably as it punishes the absence of one. Marriott's board survived a 500-million-record breach because the record showed a program that was real and monitored.

**Why it matters.** Marriott is the roadmap's promise and its discipline in one case. The promise: the Informed Decision Standard is satisfiable, and satisfying it produces dismissal even after a catastrophic outcome. The discipline: everything that protected Marriott's board was built and documented before the breach, because under Section 220 the record arrives in the plaintiff's hands as it existed.

#### CASE

*Firemen's Ret. Sys. of St. Louis v. Sorenson*, C.A. No. 2019-0965-LWW (Del. Ch. Oct. 5, 2021) (Will, V.C.). The Marriott/Starwood breach exposed up to 500 million guest records through Starwood's legacy reservation database. Pre-acquisition due-diligence claims were dismissed as time-barred; the oversight claim was dismissed on the merits because the board and audit committee were routinely apprised of cybersecurity risk, received annual enterprise risk assessments evaluating it, engaged outside consultants, and adopted responsive measures when briefed on the Starwood vulnerabilities. Identical doctrine to Boeing, opposite outcome; the difference was the documented program.

For AI, the translation needs no elaboration. The minutes being written this quarter are the exhibits of a case that has not been filed yet, and they will read either like Marriott's or like Boeing's. The difference between the two is whether a program existed and operated.

---

## PART IV

# Why AI is the next mission critical risk

## 4.1 The State of the Law, Stated Plainly

No court has yet applied the Caremark mission critical standard to AI. That is the reason to build the program now, not the reason to wait.

**Why it matters.** The doctrine's structure makes eventual application a matter of facts and time. The question is whether the board will have built the evidentiary record of a good-faith program before the first complaint, because the record cannot be built afterward. Marchand and Boeing were both decided on records assembled through Section 220 before the merits were reached.

## 4.2 The Mission Critical Case, and the Disclosure Gap

For institutions deploying AI in consequential decisions, AI is now mission critical on the institution's own filings, and the filings disclose the risk far faster than the institution builds the program to oversee it.

**Why it matters.** The mission critical designation does not require that AI be the institution's product, any more than food safety was Boeing's product. It requires that the risk be central to operations. The gap between disclosed risk and demonstrated oversight is the plaintiff's opening exhibit, drawn from the institution's own filings.

**DATA**

72% of S&P 500 companies disclosed AI as a material risk in 2025 annual filings, up from 12% in 2023 and 58% in 2024 (The Conference Board / ESGAUGE; the figure reached 83% by the April 2026 update). Over the same period, disclosure of AI expertise among S&P 500 directors rose only from 1.5% to 2.7%. Institutions are naming the risk far faster than they are building the oversight to govern it.

### 4.3 The SEC's Cybersecurity Framework as the Template

**Technology-risk governance migrates from guidance to mandated disclosure of who oversees, how they are informed, and what happened when it failed. Cybersecurity already walked the path AI will follow.**

**Why it matters.** The SEC's final cyber rules map onto the Boeing anatomy with near-mechanical correspondence: committee identification, management's role, process description, and a four-day incident clock are Boeing's elements rendered as disclosure. The 2026 reconsideration concerns the incident mechanism; the governance disclosure architecture is not its subject.

**RULE**

SEC Release No. 33-11216 (July 26, 2023). Regulation S-K Item 106(c)(1) (committee responsible and how the board is informed); Item 106(c)(2) (management's role and expertise); Item 106(b) (processes for managing material risk); Form 8-K Item 1.05 (material incident disclosure within four business days of a materiality determination). Item 1.05 is under active reconsideration following rescission petitions and the Commission's January 2026 Regulation S-K reform statement.

### 4.4 The Disclosure-Side Analogue: AI-Washing

**Misstatements about AI capability or AI governance are actionable today, under existing law, with no AI statute required. A governance claim is itself a claim that the program must substantiate.**

**Why it matters.** AI-washing and the Caremark line are two faces of one standard. The regulator asks whether the statement about governance was true; the court asks whether the governance behind the statement existed. Both are answered by the same documentary record, and the answer the doctrine rewards is not saying less. It is being able to prove what the program does.

**ENFORCEMENT**

*In re Delphia (USA) Inc.* and *In re Global Predictions, Inc.* (SEC, Mar. 18, 2024; civil penalties of \$225,000 and \$175,000 under Advisers Act §§ 206(2), 206(4) and Rules 206(4)-1, 206(4)-7); *In re Presto Automation Inc.* (SEC, Jan. 2025, Exchange Act reporting); FTC Operation AI Comply (Sept. 25, 2024); Texas AG Assurance of Voluntary Compliance with Pieces Technologies (Sept. 2024), the first state AG action of its kind.

### QUESTIONS FOR THE BOARD

- Have we documented our determination of whether AI is mission critical, with the basis recorded in the minutes?
- Can our program substantiate every public statement we make about our AI governance, in a filing, a proxy, or on our website?
- If a plaintiff obtained our board minutes under Section 220 tomorrow, would they show a program, or only a committee?

---

## PART V

# The roadmap: assembling the program

The doctrine specifies the anatomy; this Part assembles it into one program. Each element is drawn from the holding that established it and mapped onto AI at the oversight layer, what must be governed and who is accountable. How each element is implemented is the work of management, counsel, and operators; the roadmap does not cross into it. What follows is the program a court applying the Informed Decision Standard looks for, and the program the Marriott dismissal shows is sufficient.

## 5.1 Identify and Inventory

Determine whether AI is mission critical, record the determination, and maintain a current inventory of the AI acting on the institution's behalf, including vendor-embedded and agentic systems.

**Why it matters.** An institution cannot govern, escalate, or document AI it has not identified. This is the program's foundation; every later element depends on it.

### AUTHORITY

Marchand

## 5.2 Assign the Governing Body, Then Govern

Give a board committee explicit charter responsibility for AI risk, then sustain the activity that makes the assignment real: reports reviewed, questions asked, red flags pursued, management held to account.

**Why it matters.** Boeing teaches that charter silence reads as oversight absence; Clovis teaches that even a well-reported board is liable if it ignores the red flags. The committee is part of the program, not a substitute for it.

**AUTHORITY**

Boeing element 1; Clovis

### 5.3 Set the Cadence

Give AI risk regular, dedicated agenda time on a defined schedule, independent of events.

**Why it matters.** A standing agenda item demonstrates that oversight does not depend on the news cycle. Its absence demonstrated the opposite for Boeing.

**AUTHORITY**

Boeing element 2

### 5.4 Demand Risk-Centric Reporting

Require standing management reporting that is risk-centric: incidents, bias findings, regulatory inquiries, vendor developments, and the state of the inventory, not AI achievements inside strategy updates.

**Why it matters.** Boeing's board received reports that carried only good news, and the court treated the filtered channel as a systemic failure. A reporting protocol that conveys only favorable information is the plaintiff's Exhibit A.

**AUTHORITY**

Boeing element 3

### 5.5 Build the Independent Channel

Establish an escalation path by which AI red flags reach the governing body without depending on management's discretion: defined triggers, named recipients, board access.

**Why it matters.** Triggers are written before deployment, because a trigger invented after the incident is evidence for the plaintiff, and because the statutory environment now attaches clocks to escalation as fast as 72 hours.

**AUTHORITY**

Boeing element 4

### 5.6 Respond to Red Flags, Visibly

When a red flag arrives, respond actively and document it: inquiry beyond management's assurance, responsive action, follow-through to resolution.

**Why it matters.** Boeing's board treated the first crash as an anomaly; the court read that response as bad faith. Clovis turned on the same point. The first AI incident is the institution's Lion Air moment.

#### AUTHORITY

Boeing element 5; Clovis

## 5.7 Name the Officers

Give each AI-relevant domain a named accountable officer who maintains a system within it, reports on the standing protocol, and carries the documented duty to escalate.

**Why it matters.** The program's accountability layer is answered at two levels or it is not answered. A program that stops at the board leaves its officers exposed under McDonald's.

#### AUTHORITY

McDonald's

## 5.8 Keep the Record

Memorialize in minutes and committee materials the discussions held, reports reviewed, questions asked, determinations made, and actions taken.

**Why it matters.** Under the Informed Decision Standard the record is the program's proof of existence, and Texas and Connecticut now make the documented program a statutory defense. The record is no longer only the defense to the oversight claim.

#### AUTHORITY

Caremark through Boeing; 8 Del. C. § 220

Eight elements, one program. The cases describe its absence; the roadmap is its assembly. That is what it means to say the program is the governance.

#### QUESTIONS FOR THE BOARD

- Does a charter name AI, does the agenda carry it on a schedule, does reporting reach us unfiltered, does escalation reach us without management's permission, and did we respond to the last red flag in a way we would defend in print?
- Can we produce, from our own records today, the proof that each of these eight elements exists and operates?

## CONCLUSION

## Build the program before the case

The doctrine has told institutions exactly what it will ask. The program is the answer, and whether the minutes read like Marriott's or Boeing's is the one element of the case entirely within the board's control.

The Caremark line spent thirty years building a standard that is procedural, evidentiary, and now two layers deep in the institution. Applied to AI it asks for one thing assembled from many: a program. A named risk, a governing committee that governs, a current inventory, reporting that arrives unfiltered, escalation that arrives unbidden, officers who own their domains, and a record that proves all of it operated. That is not a compliance posture assembled against a statute that may change. It is the oversight layer, stable across every regulatory reversal of the past year, and it is what lets an institution adopt AI at speed, because clear obligation produces faster, more defensible decisions than ambiguous obligation. The control expires when the trauma occurs. Build the program first. The program is the governance.

## Sources and Further Reading

**Primary authorities:** *In re Caremark Int'l Inc. Derivative Litig.*, 698 A.2d 959 (Del. Ch. 1996); *Stone v. Ritter*, 911 A.2d 362 (Del. 2006); *Marchand v. Barnhill*, 212 A.3d 805 (Del. 2019); *In re Clovis Oncology, Inc. Derivative Litig.*, C.A. No. 2017-0222-JRS (Del. Ch. Oct. 1, 2019); *In re The Boeing Co. Derivative Litig.*, C.A. No. 2019-0907-MTZ (Del. Ch. Sept. 7, 2021); *In re McDonald's Corp. Stockholder Derivative Litig.*, 289 A.3d 343 (Del. Ch. 2023); *Firemen's Ret. Sys. of St. Louis v. Sorenson*, C.A. No. 2019-0965-LWW (Del. Ch. Oct. 5, 2021); 8 Del. C. §§ 102(b)(7), 220; SEC Release No. 33-11216 (July 26, 2023); SEC AI-washing orders (2024-2025); FTC Operation AI Comply (2024); Texas AVC, Pieces Technologies (2024).

**Empirical sources:** The Conference Board and ESGAUGE, *AI Risk Disclosures in the S&P 500* (Oct. 2025; updated Apr. 2026). **Further reading:** commentary on the Caremark line and AI governance is collected at the Harvard Law School Forum on Corporate Governance. The Center's companion paper, *The AI Regulatory Enforcement Landscape*, maps the regulatory environment this analysis presumes; the standing publication, *The AI Oversight Obligations Reference*, catalogues the obligations cited here with quarterly status lines.