Portfolio Media. Inc. | 230 Park Avenue, 7th Floor | New York, NY 10169 | www.law360.com
Phone: +1 646 783 7100 | Fax: +1 646 783 7161 | customerservice@law360.com

Operational AI Washing: A New Securities Class Action

By **Donnie King, Eric Coleman and Reginald Janvier** (May 12, 2026, 4:23 PM EDT)


Donnie King


Eric Coleman

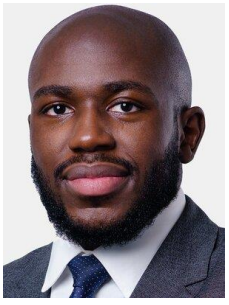
Reginald Janvier

For the past two years, the artificial intelligence securities litigation landscape has been dominated by AI washing, which refers to the practice of exaggerating or misrepresenting the role, sophistication or impact of artificial intelligence to create a misleading impression of technological capability or competitive advantage.

These practices have fueled class actions by allowing plaintiffs to allege that investors were misled by specific, verifiable claims about software performance and automation that proved untrue.

To date, most litigation and enforcement activity has focused on AI washing in the product context, that is, whether a company overstated the role, sophistication or capabilities of AI in the products or services it sells. For clarity, we refer to that theory as "product AI washing."

This article addresses a distinct and emerging risk that has received far less attention: "operational AI washing," where companies invoke AI not to describe a product but to explain internal business decisions — particularly workforce reductions, restructurings and cost-cutting measures — in ways that may obscure underlying financial distress.

The distinction is conceptual rather than doctrinal, but it is increasingly relevant as AI narratives migrate from product marketing into earnings calls, restructuring disclosures and board-level communications.

But as we enter the 2026 restructuring cycle, that focus is shifting in a more consequential direction, with the plaintiffs bar moving beyond product-level representations and toward how companies invoke AI to explain operational decisions.

That shift is driven by the market companies now face, one that rewards AI-driven efficiency but punishes demand-driven contraction. That tension creates an obvious risk: Financial distress may be framed as technological evolution. Plaintiffs may, in turn, frame AI-based explanations for operational changes as a cover for underlying financial distress, even where management believed the explanation to be accurate.

Plaintiffs have noticed and are evolving their legal theories toward operational AI washing, a shift that is still underappreciated. The new target is not the software you sell. It is how you explain the layoffs.

The theory likely to be advanced by plaintiffs is straightforward: that the workforce reduction was not driven by an AI breakthrough, but by underlying financial underperformance.

For a chief legal officer, this is not just a new wrinkle. If a plaintiff can convince a court that a survival-oriented layoff was dressed up as a growth-oriented tech pivot, the exposure extends well beyond reputational damage. You are creating the conditions for a Rule 10b-5 claim.

This article is the first in a five-part series examining how AI narratives are reshaping securities litigation risk, corporate governance exposure and shareholder inspection strategy.

The Signals: Why We Are Issuing This Warning

This theory has not yet triggered a dispositive ruling in the federal courts. However, taken together, the following signals converge on a familiar litigation risk.

The "Ghost Work" Precedent

Plaintiffs already have a well-developed playbook relating to alleged exaggerated AI claims at the product level.

In a series of cases sometimes described as involving so-called ghost work, companies promoted products as powered by proprietary or autonomous AI when, in reality, the core functions were performed by human labor, often offshore. Courts and regulators have shown little patience for that kind of mischaracterization.[1].

When AI turns out to be people behind the curtain, the issue is straightforward: The product was not what investors were told it was.

The operational AI washing theory builds on that same instinct to look past labels, but it targets a different kind of statement. The focus shifts from what a product can do to why a company took a particular business action, and whether AI was genuinely the driver offered to the market.

Klarna as an Illustration

Klarna offers a useful illustration. In late 2023, the fintech company announced a hiring freeze and began attributing its shrinking workforce to AI-driven efficiency. By early 2024, it announced that an AI chatbot built with OpenAI was handling the equivalent workload of 700 customer service agents.

Klarna's CEO reported that its headcount had fallen roughly 40%, from approximately 5,000 to 3,000, with a target below 2,000.[2] The narrative was explicit: This was not a distressed restructuring. It was a technological transformation.

The market listened — and then started checking. By late 2024, analysts and journalists were scrutinizing whether the AI efficiency story held up against the underlying financials. By mid-2025, the CEO publicly acknowledged the company had "gone too far": that replacing customer-facing roles with AI had degraded service quality and eroded customer trust. Klarna began selectively rehiring.[3]

In September 2025, it completed its New York Stock Exchange initial public offering, closing its first day

up roughly 15%, with the CEO confirming the workforce would continue declining toward sub-2,000 by 2030 through the freeze-and-attrition strategy already underway.

Within months, a federal securities class action followed — Nayak v. Klarna Group PLC, brought in the U.S. District Court for the Eastern District of New York on Dec. 22 — though not on the operational AI washing theory this series describes. The complaint alleged that Klarna's IPO registration statement materially understated the risk that credit loss provisions would surge postoffering, a credit risk disclosure claim under the Securities Act rather than a Rule 10b-5 challenge to the AI efficiency narrative itself.

That arc illustrates precisely the conditions that generate operational AI washing exposure, even if the litigation that followed targeted a different disclosure theory. When service quality falls and the board deck tells a different story than the press release, the correction becomes evidence.

Under the half-truth doctrine addressed below, the moment Klarna offered AI efficiency as the explanation for those reductions, it placed the cause squarely at issue. The capital expenditure record, the board presentations and the customer satisfaction data are each either a corroboration or a contradiction. There is no middle ground.

The SEC's New Focus

The U.S. Securities and Exchange Commission's fiscal year 2026 examination priorities direct examiners to assess whether registrants' representations regarding AI capabilities are fair and accurate, and whether technology-driven recommendations are consistent with investors' stated strategies.

While the priorities focus on registered investment advisers and broker-dealers, the underlying principle, that AI-related representations must be substantiated, is migrating into issuer-level disclosure scrutiny through SEC comment letters and the Investor Advisory Committee's December 2025 recommendation for standardized AI disclosures.[4]

When regulators start asking how AI claims map to operational reality, plaintiffs are never far behind.

Together, these signals set the stage for how plaintiffs will frame the case once an AI-based explanation for business decisions makes its way into public disclosures.

The Legal Trap: Weaponizing the Half-Truth

When these complaints are filed, they will lean heavily on what courts have described as the "half-truth" framework, which confines securities fraud claims to affirmative statements rendered misleading by what they omit, rather than simple failures to disclose.

The argument is well known. Companies are not required to disclose general financial headwinds. But once a company attributes a reduction in force to AI integration, it places the rationale for that decision squarely at issue.

Plaintiffs will argue that leaving out a liquidity crunch or restructuring pressure turns the AI explanation into a misleading half-truth.

Proving Scienter: The "Red Flags" Defense Counsel Must Watch

To survive a motion to dismiss under the Private Securities Litigation Reform Act, plaintiffs need to plead scienter, or intent to deceive. Lacking emails, they will try to reverse engineer intent from ordinary disclosures, and they are often surprisingly effective at doing so.

In practice, that effort follows a predictable pattern: Plaintiffs point to gaps, inconsistencies and silences in the public record, and argue that they reveal what the company really knew when it spoke.

The CapEx Gap

Companies should expect plaintiffs to look for the receipts. You cannot credibly claim a corporate restructuring was tech-driven if your financial disclosures show no corresponding investment in AI infrastructure, compute or other AI-enabling operational capacity.

When the numbers do not line up, the explanation itself becomes the target.

Executive Stock Sales

Unless executives are trading pursuant to a rigid, preexisting Rule 10b5-1 plan, any stock sales during a restructuring period will be weaponized. Routine transactions are recast as opportunism.

The Governance Void

If your public filings tout a massive operational transformation, but your board and audit committee minutes are silent on AI integration, complaints will contend that the story was a fabrication.

Plaintiffs will argue that it is implausible for an AI-driven pivot to have occurred without board-level discussion, and that silence turns into circumstantial evidence of intent.

The Pleading-Stage Trap

Plaintiffs will argue that courts may acknowledge the existence of internal documents without crediting their contents as true at the pleading stage, blunting attempts to resolve factual disputes early.

The Enterprise Risk: Not Stopping at Securities Fraud

If a plaintiff successfully tests this theory, the damage will not stay contained to the securities docket. It will spread quickly.

First, the employment class action: Every wrongful termination suit requires the employee to prove that the stated reason for their firing was a pretext. Once one court accepts the premise, others will borrow it. Expect Worker Adjustment and Retraining Notification Act or discrimination class actions to follow.

Second, the director and officer freeze: Insurance carriers dislike allegations of intentional fraud. Coverage fights become part of the case.

How to Reduce Exposure in 2026 Restructuring

This is not an argument for dressing up ordinary restructuring decisions in AI language. It is a warning

that once AI is offered as the reason, that explanation must be true, accurate and supported by the facts driving the decision.

You can protect your company against these claims. But legal, HR and investor relations need to be on the same page before the press release goes out.

Audit the board deck.

If the board understood the cuts as driven by cash flow, the market should not be told they were driven by technology. Public explanations must reflect the reasons actually considered by decision-makers, not a more convenient narrative.

Anticipate the evolving state regulatory landscape.

New state AI laws are adding real compliance deadlines to the operational AI washing risk.

Illinois' amendment to its Human Rights Act, effective Jan. 1, 2026, prohibits AI-assisted employment decisions that have the effect of discriminating against a protected class, and requires employers to notify workers when AI influences those decisions.

The California Consumer Privacy Act's automated decision-making regulations also took effect on Jan. 1, 2026, but the automated decision-making technology obligations that matter most for employment, specifically notice, opt-out rights and access rights for significant employment decisions, do not phase in until Jan. 1, 2027. That is the operative California deadline for employment-decision compliance.

Colorado's AI Act, which requires developers and deployers of high-risk AI systems to use reasonable care to prevent algorithmic discrimination in employment and other consequential decisions, was originally effective Feb. 1, 2026, but was pushed to June 30, following a special legislative session. The 2026 regular legislative session, which runs through May 13, could still produce amendments, but as of now, the law's substantive requirements remain unchanged.

None of these laws yet requires quarterly reporting of automated-versus-traditional job eliminations. But the trajectory is clear: Companies that cannot document which functions were genuinely automated will face increasing regulatory and litigation risk on top of what plaintiffs are already building in federal court.

Verify "meaningful" oversight.

Under emerging governance frameworks — including the European Union AI Act's requirement of "effective" human oversight for high-risk systems, which takes full effect in August; the National Institute of Standards and Technology's AI Risk Management Framework; and the SEC's own examination focus on whether AI supervision is substantive rather than pro forma — a "human-in-the-loop" defense will not survive scrutiny if the human is just rubber-stamping AI outputs.

You need to maintain records proving that your professionals had the time, data and authority to overrule the AI. Speed itself can become evidence of inadequate oversight.

Paper the tech.

If investor relations wants to talk about automation, make sure there is a funded, documented internal strategy that backs it up.

Assertions without support invite discovery. Maintain contemporaneous records supporting every public claim.

If you cannot point to the line item in the budget that paid for efficiency, do not claim it in the press release.

Kill the puffery.

Train the communications teams accordingly. Broad AI statements are not safe. Plaintiffs will treat them as promises.

Coming Up Next

This article is the first part of a five-part series. Understanding the Rule 10b-5 framework is only the first step. Locked out of traditional discovery, plaintiffs will look for alternative paths to the evidence they need.

In the next installment, we will examine how plaintiffs may use Section 220 of the Delaware General Corporation Law to bypass federal pleading barriers and gain access to board-level materials tied to AI-based operational narratives. Subsequent articles will address emerging defenses and governance practices that companies can adopt now to reduce exposure.

Donnie M. King and Eric D. Coleman are partners, and Reginald Janvier is an associate, at Akerman LLP.

[1] See, e.g., Presto Automation, Inc., Exchange Act Release No. 102,177, File No. 3-22413 (Jan. 14, 2025) (settling charges that the company falsely claimed its proprietary AI product eliminated human order-taking while human agents abroad processed the vast majority of orders); United States v. Saniger, No. 25-cr-00157 (S.D.N.Y. Apr. 9, 2025) (charging private startup founder with claiming AI autonomously processed transactions when human contractors manually completed orders); D'Agostino v. Innodata Inc., No. 2:24-cv-00971 (D.N.J. filed Feb. 21, 2024) (alleging the company marketed itself as AI-powered while operations relied primarily on offshore manual labor; motion to dismiss Second Amended Complaint fully briefed and pending).

[2] https://www.cnbc.com/2025/05/14/klarna-ceo-says-ai-helped-company-shrink-workforce-by-40percent.html.

[3] https://www.bloomberg.com/news/articles/2025-05-08/klarna-turns-from-ai-to-real-person-customer-service.

[4] https://www.sec.gov/files/approved-artificial-intelligence-disclosure-recommendation-120425.pdf.