

Operational AI Washing: Fortifying The Disclosure Record

By **Donnie King, Ryan Roman and Reginald Janvier** (May 26, 2026, 5:39 PM EDT)

In Part 1 of this five-part series, we identified a new and growing risk: operational AI washing, the theory that a company's public artificial intelligence transformation narrative is used to obscure traditional financial or cost-driven decisions.

Part 2 showed how plaintiffs will seek to obtain your books and records, and how the March 2025 Section 220 amendments to the Delaware General Corporation Law could turn those records into either your best defense or their best evidence at trial.

In this third part of the series, you have cleaned up the board decks. But the risk does not end there. Now your communications team is about to write its Form 8-K, draft the earnings call script and prepare the investor day deck. Any one of these is where it can all fall apart.

One sentence in a press release or a quote from the CEO on an earnings call is all it takes to hand a plaintiffs attorney the factual hook they need to survive a motion to dismiss. The Private Securities Litigation Reform Act safe harbor exists to prevent exactly that, but only if the people drafting your public disclosures understand how to use it.

In an operational AI washing case, that hook takes a specific form: an allegation that the company's stated justification for a workforce reduction — AI-driven transformation, automation, technological efficiency — was materially inconsistent with the operational reality reflected in the board record. If they find it, your company is plunged into the asymmetric, multimillion-dollar reality of federal discovery.

Under Rule 12(b)(6) of the Federal Rules of Civil Procedure, a federal judge cannot resolve factual disputes. They are not software engineers, and they are not going to referee competing narratives about how autonomous a company's AI systems really were at the time of a workforce reduction.

If a plaintiff plausibly alleges that the company invoked AI transformation as the justification for a reduction in force that was, in reality, driven by traditional cost pressures reflected in the board record, the court cannot credit a post hoc explanation from the chief technology officer. It must accept the plaintiff's allegations as true.



Donnie King



Ryan Roman



Reginald Janvier

You have to bypass the technical dispute entirely. The most effective way to do that is by utilizing the PSLRA safe harbor, a statutory defense that mandates the dismissal of claims based on protected forward-looking statements.

The Legal Weapon: Title 15 of the U.S. Code, Section 78u-5

The PSLRA safe harbor provides a statutory shield against civil liability for forward-looking financial and operational projections. If your disclosure qualifies, the federal judge has a statutory mandate to dismiss the claim.

The technical accuracy of your underlying AI algorithm becomes legally irrelevant at the pleading stage. The court does not need to understand the technology. The court only needs to apply the statute.

But there is a catch. The safe harbor protects projections. It does not protect history. If you make a statement about a present or historical fact, you are exposed. That limitation creates a dangerous drafting hazard.

The Danger Zone: The Mixed Statement Doctrine

A mixed statement is a single disclosure that combines a statement of present or historical fact with a forward-looking projection. When corporate communications teams announce an AI-driven workforce reduction, they naturally want to explain the "why" alongside the "what." They blend a past fact with a future projection in a single breath.

Consider this hypothetical Form 8-K disclosure: "Because our AI-driven automation initiative has reached operational maturity, we are implementing a 10% workforce reduction effective immediately, which we expect will generate approximately \$50 million in annualized cost savings."

To a public relations executive, that is a strong narrative. To a securities litigator, it is a target. It is a mixed statement, and it invites potential liability in many federal circuits.

The majority of federal circuits, including the U.S. Courts of Appeals for the Second, Third, Fifth and Seventh Circuits, take a surgical approach.^[1] They sever the sentence. They strip the safe harbor protection from the historical fact that the company's AI-driven automation initiative had already reached operational maturity while protecting the future projection.

Your core justification for the restructuring is left completely exposed. In an operational AI washing case, that justification is the claim that AI maturity, rather than margin pressure or headcount optimization, was the real driver of the decision.

The U.S. Court of Appeals for the Ninth Circuit goes even further. In its 2017 decision in *In re: Quality Systems Inc. Securities Litigation*, the court held that a materially false non-forward-looking component could effectively contaminate the entire statement.

In that jurisdiction, where a cost-saving projection is premised on a materially false statement about the company's current AI maturity or deployment status, the false present-fact component contaminates the entire disclosure and strips the projection of safe harbor protection.

You cannot gamble a multibillion-dollar market capitalization on drawing a favorable jurisdiction. You

have to draft for the strictest possible venue.

The Disclosure Playbook: Getting Every Channel Right

Winning a federal motion to dismiss starts months before a lawsuit is ever filed. The legal department has to control the pen across every public-facing channel.

The three principles below apply with equal force to U.S. Securities and Exchange Commission filings, earnings call scripts, press releases, and investor presentations. The plaintiffs attorney will read all of them.

1. Bifurcate fact from projection.

Never blend your present operational reality with your future financial hopes in the same sentence. If your restructuring requires human-in-the-loop labor to bridge the gap during an ongoing AI rollout, that present operational fact must be clearly disclosed.

State the historical fact plainly. Then, start a new sentence using opinion-based linguistic cues to introduce the projection.

Further, state that the company "anticipates," "expects" or "projects" that the new AI architecture will eventually allow operations to scale. That isolates your corporate optimism from your verifiable present facts. It prevents a judge from severing your sentence or allowing a factual error to contaminate your projections.

2. Draft cautionary language.

The PSLRA safe harbor gives defendants two independent paths to dismissal. The first, and most reliable, turns entirely on the quality of the cautionary language accompanying a forward-looking statement.

"Meaningful cautionary language," under the PSLRA, is language that tells the market, in concrete operational terms, why an AI-driven projection may fail, even if management hopes or believes it will succeed. When that standard is met, it is a powerful shield. The court never reaches questions of intent or internal belief. You win the motion to dismiss even if the plaintiffs allege that management internally doubted the projection.

This structure matters because the safe harbor is disjunctive. First, a forward-looking statement accompanied by meaningful cautionary language is protected as a matter of law. Second, even if meaningful cautionary language is absent, the statement is still protected unless the plaintiff can plead facts showing that the speaker had actual knowledge that the statement was false or misleading when made.

Reliance on the second prong is inherently risky because it invites discovery fights over internal deliberations and state of mind, the very risk the safe harbor was designed to prevent.

To avoid this fight, companies should not assume that generic, hypothetical warnings will suffice under the first prong. Boilerplate cautionary language, particularly when it describes known problems as hypothetical risks, frequently fails the meaningfulness test.

As the U.S. Court of Appeals for the District of Columbia Circuit explained in its 2015 decision in *In re: Harman International Industries Inc. Securities Litigation*, cautionary statements are not meaningful if they warn investors about risks that have already materialized.

In an operational AI washing context, if any public disclosure attributes a workforce reduction to AI-driven automation, you should consider identifying the specific operational contingencies that could cause that initiative to fall short today. That could mean warning investors about the specific, present-day operational constraints that could prevent the AI initiative from delivering the projected results: data-governance limitations, reliance on third-party vendors, integration delays or the continued need for human oversight during the transition period.

The more precisely those warnings track the risks actually reflected in management's board-level discussions, the more powerful they become, both under the safe harbor and under the Section 220 incorporation doctrine, which permits plaintiffs to plead from board-level books and records obtained in a Delaware inspection, as discussed in Part 2.

3. Align the total mix of information.

A highly strategic Form 8-K can be instantly undermined by a reckless press release. Judges evaluate the total mix of information available to the market. Plaintiffs will actively search for a definitive, marketing-heavy quote from your CEO and juxtapose it against the cautious disclosures in your SEC filing to allege a calculated half-truth.

In an AI-driven workforce reduction, the most dangerous version of this scenario is a CEO who tells the press the workforce reduction reflects "the company operating at a new level of AI maturity," while the Form 8-K quietly hedges that full deployment remains in progress. The legal department must ensure that every public-facing document uses the same linguistic cues, bifurcated sentence structures and cautious framing as the Form 8-K itself.

That obligation does not stop at the SEC filing. It runs through every channel where the company speaks to the market.

The Channels That Are Often Missed

The Form 8-K is the most obvious target, but it is rarely where a plaintiff finds their best material. The most damaging mixed statements in AI washing cases tend to surface in three places that receive less legal scrutiny before they go out the door.

First, earnings calls are the highest-risk venue in the disclosure record. They occur in real-time, may depart from a script and are heavily quoted by analysts and the financial press.

A CEO who departs from prepared remarks to characterize a workforce reduction as evidence that the company's AI "is now running the operation" has made a present-fact statement, without cautionary language, in front of the entire market. That statement is transcribed, indexed and available to plaintiffs counsel the same afternoon.

The fix is not to restrict what executives can say; it is to ensure that earnings call scripts and Q&A preparation reflect the same bifurcated structure and hedging language as the Form 8-K, and that

executives understand why the precision matters before they take the call.

Second, the management's discussion and analysis section of Forms 10-K and 10-Q deserves the same level of legal review as any Form 8-K disclosure, and in practice, it often receives less.

MD&A is where companies describe operational trends in narrative form, and AI transformation initiatives are increasingly prominent there. A description of AI deployment progress that overstates current maturity or omits known integration constraints is just as actionable in a Form 10-K as it would be in a stand-alone filing.

Because Forms 10-K and 10-Q disclosures are annual and quarterly rather than event-driven, legal teams sometimes treat them as routine. In the current environment, the AI-related sections of an MD&A should be treated as anything but.

Third, investor day presentations and analyst briefings are where companies tend to make their boldest AI claims — and they are also where legal review is most likely to be minimal or absent.

Slides prepared by the strategy or product teams and presented to institutional investors carry the same exposure as any other public communication. A projection about AI-driven margin expansion presented at an investor day without accompanying cautionary language is a forward-looking statement without a safe harbor shield.

Plaintiffs attorneys know how to use investor day transcripts. Legal counsel should be in the room, or at a minimum, reviewing the materials, before those events take place.

Looking Ahead

A strategically managed disclosure record isolates historical facts and builds a protective wall around operational projections across every channel. It also completes the internal alignment that makes the Section 220 incorporation doctrine a genuine weapon: When the board deck, the cautionary language and the public narrative all reflect the actual status of the AI deployment, a plaintiff who forces production under Section 220 may find that they have inadvertently enabled the company to use its own records as powerful pleading-stage evidence supporting dismissal.

Once incorporated by reference into the complaint, those documents can be used to demonstrate the company's good faith transparency and undermine the plaintiff's theory at the motion-to-dismiss stage.

Plaintiffs will still try to plead around the safe harbor. They will attack the adequacy of the cautionary language or argue that executives knew internal facts that directly contradicted the public narrative.

But once a plaintiff bypasses the boardroom and files suit, the litigation landscape shifts entirely. The same AI-driven workforce narratives that once appeared in earnings calls and Form 8-Ks become the raw material for Rule 10b-5 claims; the company's first and best opportunity to win is at the motion to dismiss.

In Part 4, we will turn to the courtroom. We will examine how defense counsel can leverage the PSLRA's heightened pleading standards to attack scienter before discovery extracts its toll, sever loss causation where macroeconomic noise or predisposed risk obscures any clean causal line, and neutralize aspirational AI statements under the puffery doctrine.

We will also close the loop: A Form 8-K drafted with the precision described in Part 3 is not merely a disclosure document; it becomes, in litigation, a shield. The same specificity that satisfies the safe harbor is what makes allegations of scienter implausible at the pleading stage.

The defense against operational AI washing is not reactive. It is built long before any complaint is filed.

Donnie M. King and Ryan Roman are partners, and Reginald Janvier is an associate, at Akerman LLP.

The opinions expressed are those of the author(s) and do not necessarily reflect the views of their employer, its clients, or Portfolio Media Inc., or any of its or their respective affiliates. This article is for general information purposes and is not intended to be and should not be taken as legal advice.

[1] See *In re: Quality Sys. Inc. Sec. Litig.*, 865 F.3d 1130, 1141–42 (9th Cir. 2017) (collecting cases and observing that the Second, Third, Fifth, and Seventh Circuits have each concluded that where a defendant makes a mixed statement containing both non-forward-looking and forward-looking components, the safe harbor does not extend to the non-forward-looking portion); see also *In re: Vivendi SA Sec. Litig.*, 838 F.3d 223, 246-49 (2d Cir. 2016) (holding that present-fact representations embedded in otherwise forward-looking statements are not shielded by the PSLRA safe harbor).