LAW360

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: Dismantling Claims Before Discovery

By **Donnie King, Ryan Roman and Cherly Lucien** (June 2, 2026, 1:49 PM EDT)

The first three parts of this series traced the arc of an operational artificial intelligence washing claim from inception to the boardroom door: the half-truth doctrine, the Section 220 demand and the public disclosure exposure. Part 4 begins where those leave off.

Donnie King

When a plaintiff bypasses negotiation and files an action under Rule 10b-5 of the Securities Exchange Act, the company's most important decision is not how to litigate; it is how to quickly end the case. The Private Securities Litigation Reform Act was designed for precisely this moment, and defense counsel who understand its architecture can exploit heightened pleading standards to dispatch the complaint before discovery begins to extract its real cost.

Ryan Roman

What follows is not a checklist of defenses. Each argument addressed below is strongest when it rests on the governance and disclosure work the company has already done and weakest when it contradicts it. When those layers speak from the same documentary record, the defense is as strong as it can be. When they don't, the defense collapses, regardless of how well individual arguments are framed.

The First Line of Defense: Attacking Scienter

No element of a PSLRA claim is more plaintiff-hostile than scienter. A complaint cannot survive a motion to dismiss on speculation. It must plead specific facts giving rise to a strong inference, not merely a plausible one, that the defendants either knew their statements were false when made or acted with a degree of recklessness so extreme it approaches intentional misconduct. In the AI context, that bar is particularly difficult to clear.

Cherly Lucien

The reason is straightforward: AI integration is genuinely uncertain. Model performance projections, efficiency timelines and workforce impact estimates involve layers of engineering, organizational and market assumptions that reasonable actors can get wrong. Defense counsel should anchor the scienter argument there. If leadership relied on documented internal assessments — technical road maps, vendor representations, data science team projections — the failure of an AI initiative to deliver on those projections is a business execution problem, not a disclosure crime. Courts have consistently drawn that line, and the defense should draw it early.

Plaintiffs often try to bridge the scienter gap by pleading insider trading allegations in an attempt to show motive. The trading record should therefore be examined early and carefully. Where executives held, accumulated or sold shares pursuant to routine compensation practices or preexisting trading plans during the class period, any inferences of fraudulent intent may quickly collapse. Plaintiffs who cannot point to suspicious sales at suspicious times or in suspicious amounts are left arguing with the implausible theory that management intentionally deceived the market for no discernible personal benefit.

A third scienter pressure point arises from internal documents. Plaintiffs frequently surface engineering memos, internal Slack threads or risk committee minutes where employees raised concerns about AI readiness, then attempt to treat those documents as smoking guns.

The answer for the defense is to contextualize the documents. Companies that surface, debate and document technical concerns in real time are not committing fraud; they are operating as functioning enterprises are supposed to. Internal dissent is not evidence of a cover-up. It is evidence of a due diligence culture, and it should be framed that way from the opening brief.

Severing Loss Causation

Even when plaintiffs clear the scienter hurdle, they face a separate evidentiary problem: They must trace the stock decline directly to the alleged AI disclosure, not to everything else that was happening in the market at the same time.

In practice, AI-related corrective disclosures almost never occur in isolation. They coincide with earnings misses, sector downturns, interest rate movements or broader technology sell-offs. The task of the defense's economic expert is to disaggregate those factors and demonstrate that the plaintiff's claimed disclosure did not cause, or did not exclusively cause, the price decline. Where the causal picture is genuinely ambiguous, courts have been willing to find loss causation unpled.

There is a related argument that deserves equal attention: the no-new-information doctrine.

If the company had previously disclosed that AI integration carried material risks, that implementation timelines were uncertain or that efficiency gains depended on factors outside management's control, then a later admission that those risks materialized is not a corrective disclosure at all. It is the realization of a publicly acknowledged risk, already priced into the stock. That argument should be developed carefully against the actual disclosure record, rather than asserted generically.

Neutralizing Aspirational Statements: The Puffery Doctrine

A significant portion of what companies actually say about AI in public communications is not, legally speaking, a representation of fact. Characterizations like "AI-first," "transformative technology" or "next-generation efficiency" represent the kind of generalized corporate optimism that courts have long held to be nonactionable puffery. Reasonable investors do not make investment decisions based on superlatives, and they are not entitled to sue when superlatives fail to materialize.

The key distinction is between concrete, falsifiable operational statements — such as quantified efficiency gains, AI-attributed headcount reductions or AI-based margin forecasts — and generalized, aspirational characterizations that do not purport to describe verifiable facts.

The former may be actionable if false. The latter generally are not. Defense counsel should map every contested statement against that framework before the motion to dismiss is filed and categorize aggressively. The puffery defense is particularly well-suited to AI capability statements, which tend toward aspiration rather than verifiable fact.

That said, the puffery defense is not a cure-all. It will not protect a statement that pairs aspirational language with hard metrics or one made in close temporal proximity to data suggesting the opposite was true. The argument has to be built on the actual record, not deployed as a generic shield.

The Integrated Defense

The individual arguments above matter most when they are deployed as part of a coordinated strategy rather than stand-alone defenses. A successful defense against an operational AI washing claim is not assembled at the courthouse; it is built in the boardroom and the disclosure function before litigation ever begins. The litigation arguments discussed here in Part 4 are strongest precisely because they rest on the governance and disclosure work covered in Parts 1 through 3.

Consider how each defense layer connects to the work already done.

Defeating the scienter argument at the motion to dismiss stage substantially undermines any subsequent Section 220 demand, because the rationale for that demand — that inspection of books and records is necessary to investigate potential wrongdoing — becomes far harder to sustain after a court has found the wrongdoing insufficiently pled. A strong substantive win on Rule 10b-5 reinforces the procedural defenses erected against the records demand and keeps internal communications out of the plaintiff's hands at the most sensitive stage.

The puffery and loss causation arguments, in turn, are validated by the Form 8-K architecture described in Part 3. By expressly identifying AI-specific risks — performance degradation over time, uncertain implementation timelines, variability in training costs and dependencies on successful system integration — the disclosure both renders subsequent statements nonactionable puffery and defeats any claim that later disclosures revealed "new" information. It satisfies the PSLRA safe harbor for forward-looking statements and makes it nearly impossible for a plaintiff to later claim that management concealed what it had, in fact, told the market in plain terms.

Finally, there is the directors and officers insurance dimension, which practitioners sometimes underweigh. A swift dismissal is not only a litigation victory; it is a policy management event. Prolonged AI-related litigation invites carriers to examine whether the professional services exclusion applies or whether application-stage representations about the company's risk posture support a rescission argument. Disposing of the case early, on substantive grounds, preserves limits and closes the window for those arguments.

The thread running through all of it is consistency. The same factual record that defeats scienter should be the record reflected in the company's public disclosures. The risk factors that win the safe harbor argument should be the same risks the board was monitoring in real time.

When governance, disclosure and litigation strategy speak from the same documentary record, the defense is as strong as it can be. When they contradict each other, when disclosures are generic but internal records show specific known risks, the defense collapses regardless of how the individual arguments are framed.

The best AI washing defense, ultimately, is not one that will be built after the complaint is filed. It is one that was built correctly from the beginning.

Looking Ahead

The litigation and disclosure frameworks developed throughout this series ultimately answer to a higher-order question: What do boards owe when AI becomes mission-critical?

In the final installment, we'll close the loop at the Delaware level. As AI systems move from the margins of corporate operations to their center, the Caremark duty of oversight follows. Part 5 examines what that duty requires of directors — in governance architecture, reporting systems and engagement with management — and why the path to fiduciary protection in the age of AI is less novel than it first appears.

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