AI-Washing & Defensible Disclosures: An Executive Playbook for 10-K & 20-F

A practical framework for CFOs, General Counsels, and Audit Committees

Between 2022 and 2024, the share of S&P 500 companies referencing AI in their 10-Ks rose from 12% to 73%. At the same time, the SEC has signaled that "AI-washing" is a live enforcement theme, treating overstated AI claims like any other misleading disclosure.

For CFOs, GCs, CAIOs, CROs, and audit chairs, that combination creates a simple reality: AI language in the 10-K is now part of the evidence record. The bar is straightforward. If a claim about AI cannot be backed by documentation, testing, or governance artefacts, it does not belong in the filing.

A defensible AI disclosure does three things consistently. It ties statements to specific systems and use cases, it reflects impact only where there is reproducible analysis, and it describes governance in a way that matches how the company actually runs AI. Leading issuers are already moving this way, for example by naming the committees that oversee AI and summarizing how AI risks are reviewed at board level.

Executive Summary

AI-related disclosures in 10-Ks and 20-Fs now carry the same enforcement risk as any other material statement. Between 2022 and 2024, S&P 500 AI references rose from 12% to 73%, while the SEC designated AI-washing as a priority enforcement theme.

Key takeaways:

The risk is immediate.

AI-washing occurs when narrative outpaces systems, analysis, or governance—not through fraud, but through routine drafting where claims cannot be traced to specific models, validated analyses, or documented controls.

Defensible disclosure requires evidence.

Leading issuers tie every AI statement to specific systems, reproducible analysis, and governance artefacts that can be shown to regulators and auditors.

A structured pre-clearance process works.

A 10-day review can map AI claims to systems, assess control gaps against IEEE and NIST standards, and align disclosure language with actual governance before filing.

No-regret moves exist.

Establishing an AI inventory, requiring artefacts for new claims, running cross-functional reviews, and aligning to common frameworks reduce risk immediately without multi-year programs.

This playbook provides a practical framework for CFOs, General Counsels, CAIOs, CROs, and audit committees to ensure AI disclosures are defensible when scrutiny comes.

AI language is now moving markets, not just filling conference calls. Mentions of AI risk and disruption are dominating Q&A, and some issuers have seen their stocks sell off sharply immediately after executives flagged AI exposure on earnings calls. When a few sentences about AI can erase value in a single session, generic or unsupported AI claims in 10-Ks and 20-Fs become a direct disclosure and valuation risk, not a branding choice.

What a "defensible" AI disclosure looks like

A defensible AI disclosure links 10-K statements to production systems and controls.

Specific Systems & Uses

Disclosures should refer to concrete AI models, tools, or workflows, rather than generic terms like "advanced AI."

Clear Material Impact

Claims about efficiency, revenue, or cost must be backed by reproducible analysis.

Visible Governance

Documented processes for AI inventory, testing, and reporting to senior management and the board are essential.

Leading filers name the board committee that oversees AI and summarize how AI risks are reviewed. When those disclosures match internal artefacts, the company is in a much better position if questions arise later.

If you cannot evidence it, you cannot credibly claim it. Strong internal documentation is crucial for defensible disclosures.

Where AI-washing risk shows up in practice

In most organizations, AI-washing risk does not start with fraud. It starts where narrative gets ahead of systems, analysis, or governance. Across recent filings and enforcement commentary, three failure modes recur:

Narrative without systems.

Companies describe "AI-powered" products or "advanced AI platforms" without being able to tie claims to a specific model, vendor, or workflow. When questions arise, teams struggle to identify which systems actually underpin the language in the filing.

Impact without analysis.

Statements about efficiency, accuracy, or returns are made without documented analyses, back-tests, or validation work. In some cases, marketing claims migrate into the 10-K with no clear owner for the underlying numbers.

Governance without artefacts.

Filings reference AI "governance," "ethics review," or "risk oversight" while inventories, minutes, risk registers, and testing records are incomplete or out of date. When regulators or auditors ask for evidence, the governance story cannot be reconstructed from available documents.

Each of these failure modes is reversible, but only if AI claims are treated as control statements that must be linked to systems, owners, and records before the filing is approved.

Turning AI claims into evidence

A Zero Trust Governance approach treats every AI sentence in the 10-K as a control statement that should be backed by artefacts. The ETM Evidence Standard organizes this work into four areas:

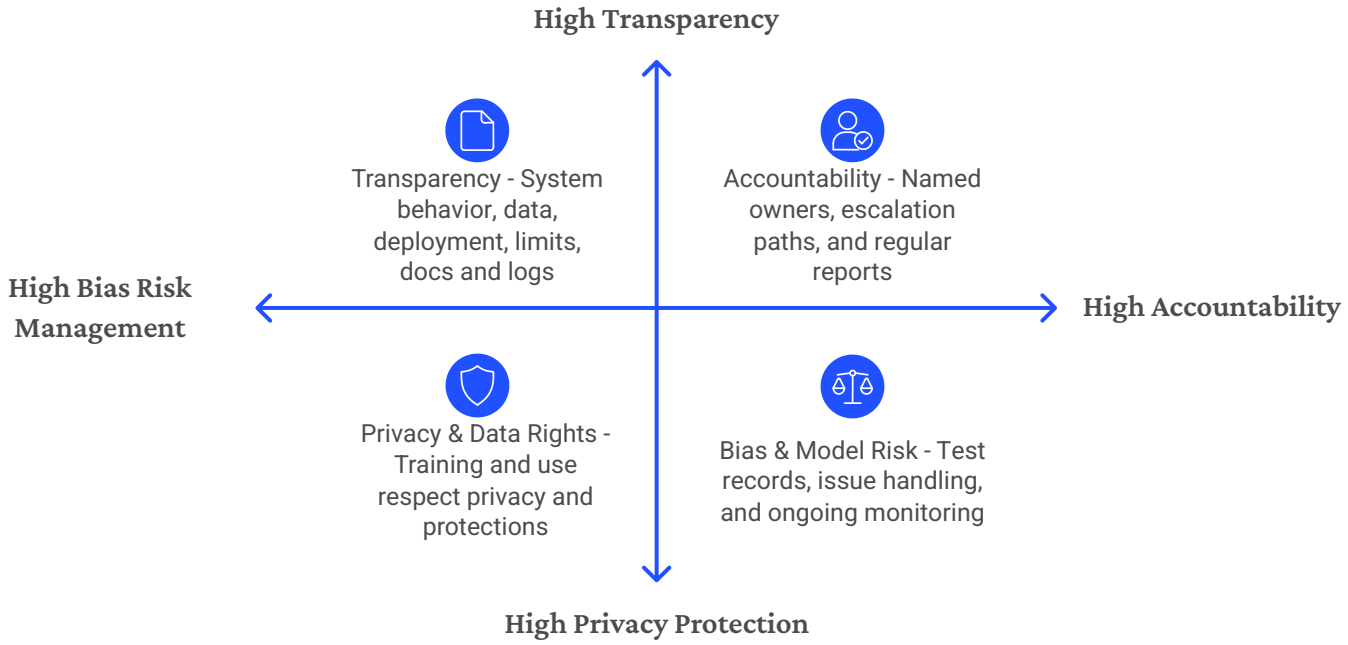

The ETM Evidence Standard is our way of making sure every AI claim in your filings ties back to documentation, testing, and governance records that can be shown to regulators and auditors.

A four-step AI-washing pre-clearance process

You can treat AI-related disclosure like other high-risk areas and run a short, structured review before filing. A practical four-step flow is:

01

Collect all AI claims

Days 1-2: Gather every mention of AI across 10-K/20-F drafts, earnings scripts, investor decks, and ESG reports. Tag each as descriptive, forward-looking, performance-related, or governance-related so reviewers see where the risk lies.

02

Link claims to systems and evidence

Days 3-5: For each statement, identify the actual systems involved, including vendors and human-in-the-loop work. List the artefacts you already have: model documentation, validation work, incidents, committee minutes, and policies.

A four-step AI-washing pre-clearance process (continued)

01

Run an IEEE-aligned controls check

Days 6-8: For material systems, review them against the transparency, accountability, bias, and privacy expectations used in IEEE CertifAIEd assessments and in frameworks like NIST AI RMF. Note where the disclosure language needs to adjust to match the real control posture.

02

Standardize language and reporting

Days 9-10: Develop standard wording for recurring AI use cases and risks, tied to specific thresholds of impact and control maturity. Align this with board reporting so audit and risk committees see the same inventories and dashboards that underpin the filing.

This flow is compact enough to run in the last weeks before filing, and robust enough to feed into a broader AI governance and model-risk program. For most issuers, these steps are no-regret moves: they reduce enforcement and litigation risk while improving the quality of internal AI governance.

No-regret moves for the next filing cycle

Issuers do not need a multi-year program to reduce AI-washing risk before their next 10-K. A small set of moves pays off quickly:

01

Establish a single AI inventory linked to disclosure.

Maintain a list of material AI systems tied to financial statement line items, key metrics, and risk factors, and keep it in sync with what the 10-K describes.

02

Require artefacts for every new AI claim.

Make it standard that any new AI-related statement in the 10-K references specific documentation, tests, or governance records, in the same way non-GAAP measures require support.

03

Run a cross-functional AI disclosure check.

Bring Finance, Legal, CAIO/CIO, CISO, Risk, and Internal Audit into a short pre-clearance review focused only on AI language, rather than spreading responsibility informally.

04

Align evidence with at least one common framework.

Structure artefacts so they can be read against recognized benchmarks such as IEEE CertifAIEd criteria and NIST AI RMF, and reused for SEC, EU AI Act, and internal audit questions.

These steps do not eliminate AI-related disclosure risk, but they make it much more likely that the company can explain and defend its AI narrative when scrutiny comes.

Who owns the controls and what must be shown

Regulators and boards expect clarity on who is responsible for AI-related disclosures and controls. A robust framework clearly defines roles:

CFO/Controller

Owns the accuracy of AI-related financial and operational information for disclosures.

General Counsel

Manages the disclosure process, ensures consistency, and advises on materiality of AI-related information.

CAIO/CIO/Product Leaders

Maintain the AI system inventory and own the technical documentation for AI models and their outputs.

CISO/Data Protection Officer

Owns security and privacy risks tied to AI systems, ensuring compliance with data protection regulations.

CRO/Compliance Officer

Maintains the AI risk register and maps AI-related risks to relevant regulations and internal policies.

Internal Audit

Independently tests AI controls and ensures the reliability of information linked to public statements about AI.

For effective oversight, boards need to see a clear list of material AI systems, verifiable evidence of control testing results, and detailed records of board meeting discussions and decisions regarding AI governance and disclosures.

How Ethical Tech Matters helps

Ethical Tech Matters focuses on AI governance that can be proven, not just described. Our work on defensible filings uses a Zero Trust Governance model and methods aligned with IEEE CertifAIEd assessments so that your AI story is matched by the underlying evidence.

10-day AI-Washing Pre-Clearance Diagnostic

For public and pre-IPO companies, we offer a focused 10-day diagnostic that validates AI claims and identifies gaps before filing.

- **Consolidated map:** We create a comprehensive map of all AI-related statements across your draft disclosures and public materials.
- **Linked evidence & assessment:** Each statement is linked to its supporting systems, responsible owners, and verifiable artifacts, with a gap assessment.
- **Concrete suggestions:** We provide actionable recommendations to tighten risk-factor disclosures, Management's Discussion & Analysis (MD&A), and overall governance language.
- **Board-ready summary:** You receive a concise, board-ready summary slide detailing key findings and proposed next steps for remediation.

Recent pre-IPO client identified 23 unsupported claims across 47 AI references in their disclosure documents using our diagnostic.

The 10-day diagnostic in practice

Our structured 10-day sprint ensures rapid, thorough assessment:

Data Collection
Days 1–3: gather inputs and map systems

Statement Analysis
Days 4–6: evaluate claims and assess risk

Recommendation Development
Days 7–8: craft prioritized remediation

Final Report
Days 9–10: deliver report and presentation

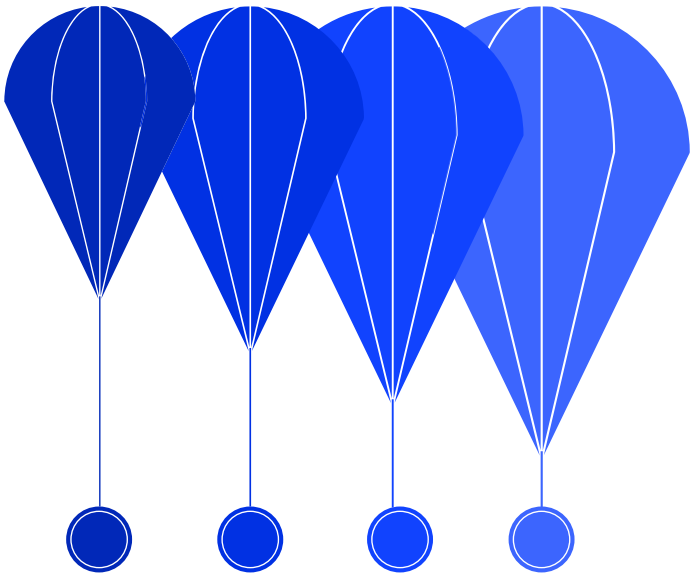

Proof Point: A recent pre-IPO client successfully identified 23 unsupported claims across 47 AI references in their disclosure documents using our diagnostic.

Our Commitment to Independence: Ethical Tech Matters provides objective, unbiased advice. We maintain strict independence from technology vendors, ensuring our recommendations are solely in your best interest and can challenge claims with the same skepticism regulators will apply.

Next Steps: Ensure Your AI Disclosures Stand Up to Scrutiny

If you want a quick internal test, take your top ten AI statements from current drafts and ask a simple question: 'What could we hand the SEC or an audit committee today to back this up?' Any answer that relies on 'we all know this is true' rather than artefacts is a gap.

Do you have readily available artifacts or documentation?

Can you immediately provide tangible artifacts, system logs, or documented processes to substantiate each AI claim, rather than relying on internal assumptions?

Who is the specific owner responsible for this claim?

Identify the individual accountable for the underlying AI system, data governance, and control narratives related to this public statement.

Does this align with current regulatory standards?

Evaluate your claim against the standards set by the SEC, EU AI Act, and NIST. Are there any inconsistencies or potential "AI Washing" vulnerabilities?

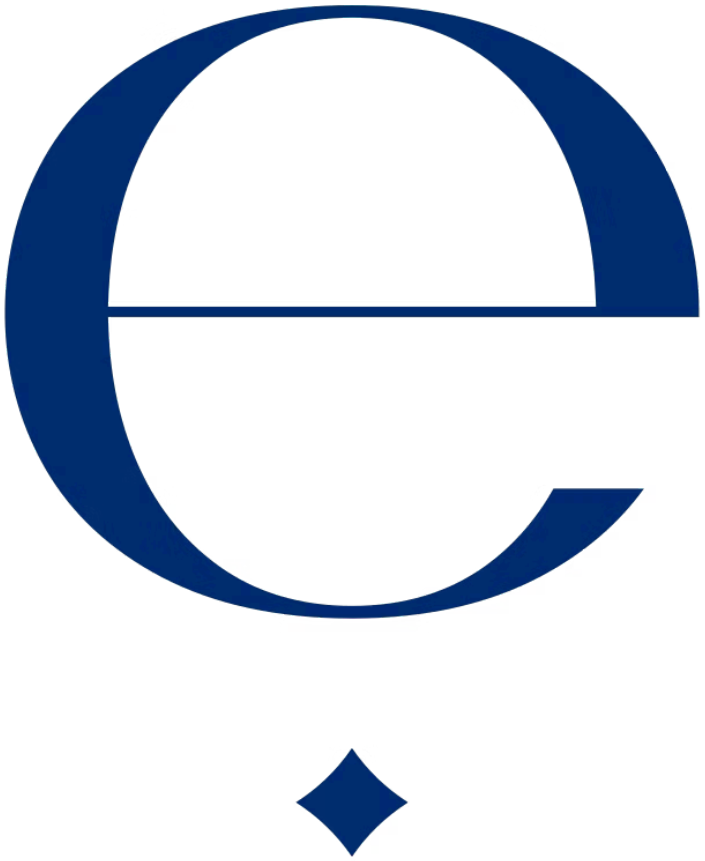

Get Started

Ethical Tech Matters offers a complimentary **AI-Washing 10-K Controls Checklist**. This checklist outlines the key artifacts and control owners that boards and regulators expect to see behind AI disclosures.

To receive the complimentary checklist or schedule a diagnostic consultation, email hello@ethicaltechmatters.com or call (817) 580-4628.

Serving defense contractors, Class I railroads, energy operators, and critical infrastructure providers.

> **AI governance isn't optional. It's operational.**

702 Houston Street, Fort Worth, TX 76102
(817) 580-4628 · hello@ethicaltechmatters.com