

AI Regulatory Landscape and Brand Representation

Research memo, Phase 1

Prepared for: GOSH.ai

Author: Reese Crawford

Executive Summary

This memo covers the five regimes that matter most to AI brand governance: the EU AI Act, NYC Local Law 144, state level legislation, the FTC, and the SEC. A single gap runs through all of them; Existing rules were written to protect consumers and workers from AI. The use case at the center of our work at GOSH, AI systems describing, recommending, or omitting brands when answering a consumer's question, sits outside every one of them. Algorithmic accountability rules point at the wrong party and output accountability rules were written for human publishers.

AI regulation in mid-2026 is fragmented and moving fast: The EU AI Act is in phased enforcement, FTC has brought at least a dozen Section 5 cases, SEC has filed its first AI-washing charges against public companies. States passed 145 AI laws in 2025 and introduced 1,561 more in the first quarter of 2026. Not one of these regimes was built to protect brands harmed by AI.

In the eyes of a risk or legal actor, you cannot comply with rules that haven't been written. The exposure is already here: defamation risk, securities disclosure duties, FTC scrutiny of a brand's own AI claims, and state laws fragmented approach. Brands that build governance ahead of the rules will be ready when the rules arrive. This regulatory gap is the space Sentinel occupies.

1. The Definitional Problem

Every survey of AI regulation hits the same wall on page one: regulators have not agreed on what AI is. The EU AI Act took four years to pass partly because regulators could not land on the definition and was disputed until the final months. The United States has no federal definition, states have turned to writing their own. This incoherence pushes the question further downstream. Before you can ask whether AI is regulated, you must ask: regulated as what, by whom, under whose definition? In mid-2026 the honest answer is that it sits in a gap. The frameworks exist, but nobody has focused on brand representation.

To navigate the landscape, every regulatory framework in this memo has been broken into three structural questions. *How is risk handled*, by tiers (the EU model), by high-level principles (the UK model), or by specific procedural rules (NYC Local Law 144)? *Where does it apply*, horizontally across all sectors or vertically through existing sectoral regulators? *Who is the regulated party*, the AI system and its developer (algorithmic accountability), or the deployer (output accountability)?

That last question is the one that matters most for brand governance. Rules built around the algorithmic accountability system points at the model provider, not the brand a model just misdescribed. While

rules built around output accountability were written with a human publisher in mind, and they fit AI-generated content poorly. That mismatch is why brand representation has fallen through the cracks.

2. EU AI Act

This is the world's first comprehensive AI law, in force since August 2024, with obligations rolling through 2028. Every other major regime gets measured against it. It also reaches well beyond Europe. A provider or deployer outside the EU is covered the moment its system's output is used inside the EU. When ChatGPT generates content about a U.S. brand for an EU consumer, OpenAI's obligations under the Act apply.

The Act splits AI systems into four risk tiers. *Unacceptable-risk* systems are banned. *High-risk* systems (employment, credit, healthcare, critical infrastructure) face conformity assessments, risk management, data governance, and human oversight requirements. *Limited-risk/Minimal-risk* systems require transparency only, including AI-generated content disclosure under Article 50. A parallel framework applies to general purpose AI providers (OpenAI, Anthropic, Google, Meta) with obligations including technical documentation, copyright compliance, and training data summaries. Penalties reach €35 million or 7 percent of global turnover.

Brand recommendation never made the risk list. Those categories were drawn around decisions that hurt people, like a denied loan or a rejected job application, not specific to brand image. The brand itself is a bystander. And nothing in the Act gives a brand a right to be represented accurately.

Sentinel relevance. *The EU AI Act establishes the regulatory vocabulary (risk tiers, conformity assessment, systemic risk) that future regulation of AI brand harm will likely borrow. Brands that establish governance practices now will be positioned when those tools are extended.*

3. NYC Local Law 144 (AEDT Bias Audits)

Local Law 144 was the first city-level AI rule in the United States with real implications. If an employer uses an automated tool to hire or promote for a role based in New York City, it has to run an independent bias audit, post the results, and tell candidates at least ten business days before the tool sizes them up. The trigger is where the job sits, not where the company or the applicant sits. The law was enacted in 2021 and took effect in January 2023.

Then came the reckoning. On December 2, 2025, the New York State Comptroller published an audit of how the Department of Consumer and Worker Protection (DCWP) had actually enforced the law from 2023 through mid-2025. The picture was not flattering. Weak complaint handling, minimal public education, almost no proactive hunt for violators. The Comptroller's own spot check turned up at least 17 likely violations across 32 companies. The DCWP had found one. The violating agency has since promised tougher enforcement, and employment lawyers are warning clients that 2026 will look different.

None of this touches AI brand representation directly. The law is about one use-case, hiring; and one harm: discrimination, against protected classes. What makes it worth studying is the shape.

Independent audit, public disclosure, notice to the people affected, compliance before deployment. That template is already being copied across the country, and it could just as easily be pointed at other AI uses down the line.

Sentinel relevance. *LL144's enforcement demonstrates that laws which look toothless on paper can ramp up quickly. The infrastructure required to satisfy a future brand audit law is a similar infrastructure we're working on with Sentinel.*

4. State-Level AI Laws

In the absence of a comprehensive U.S. federal AI law, states have moved aggressively. 145 AI-related bills became state law in 2025, a 47 percent increase over 2024. In Q1 2026 alone, lawmakers in 45 states introduced 1,561 new AI bills. The National Conference of State Legislatures calls this pattern "Fifty Little AI Acts." The landscape is too fragmented to track law-by-law; more durable framing is by approach. There seems to be four archetypes that have been explored.

The first is comprehensive and risk-based, a micro version of the EU model. It sorts systems by risk and puts obligations on both developers and deployers of the high-risk ones. Colorado is the lead case, covering consequential decisions across employment, education, finance, healthcare, housing, insurance, and government. Connecticut, Virginia, and Illinois are working on similar paths.

The second is sectoral. Rather than regulate AI in general, it picks specific high-risk uses and runs them through the regulators who are familiar with those domains. NYC Local Law 144 covers hiring. The Illinois AI Video Interview Act covers recorded interviews. California's AB 3030 covers healthcare. Tennessee's SB 1580 bans marketing an AI as a licensed mental health professional, and Georgia's SB 444 reaches insurance. Stack them up and these laws now cover most high-risk uses, just unevenly from state to state.

The third leans on disclosure instead of restriction. It does not tell anyone what they can or can't build. It tells them to say when AI is in the loop. California's AI Transparency Act (SB 942) is the anchor, requiring providers to disclose AI-generated content and to support content provenance, with duties phasing through 2028.

The fourth aims at the biggest models and biggest harms. California's TFAIA (SB 53) puts duties on frontier developers. Texas, through TRAIGA (HB 149, effective January 2026), takes an intent-based approach that bans systems built to cause specific harms.

Underneath the patchwork, the same six demands keep surfacing: disclosure, bias prevention, impact assessments, a path to human intervention, data provenance, and accountability that flows through to vendors. And almost none of these speak to AI brand representation. The comprehensive laws chase consequential decisions. The sectoral laws skip past brand discovery. The frontier laws aim at developers. The transparency laws come closest, yet no regulator has gone after a model provider over how it portrays a brand.

Sentinel relevance. *For a multi-location brand operating across all 50 states, state-by-state compliance is practically impossible. The compliance answer is not to follow each state law*

individually; it is to build a defensible internal governance posture that addresses the harm categories underlying all state laws. The fragmentation itself creates demand for a managed-service approach.

5. FTC Guidance and Enforcement

The FTC has no AI statute and Congress has not passed one. So the agency reaches for the authority it already has, Section 5, the ban on Unfair or Deceptive Acts or Practices (UDAP), and applies it to AI conduct. The result is the most active AI enforcement program in the country, run entirely on a law that predates the technology. Operation AI Comply, launched in September 2024, has already produced at least a dozen major actions, and it kept going straight through the change in administration with bipartisan support.

Four working standards have emerged from the cases. A claim about what your AI can do has to be backed by real evidence the moment you make it. What you trained the model on has to match how you market it; that was the heart of the Workado case in April 2025. "AI-powered" has to mean the AI actually does the work, which sank Growth Cave in January 2026. And selling a tool built to deceive is itself a violation. A December 2025 executive order told the FTC to publish a policy statement clarifying how Section 5 applies to AI by March 11, 2026. That statement is overdue, and when it lands it will be one of the most important compliance document U.S. companies have on AI claims.

Here is where the FTC's reach runs out. It has clear footing when an AI provider or a measurement platform oversells what it can do, and when a brand uses AI to lie about itself. It has no clean theory for the harms that dominate this space: an AI getting a third-party brand wrong. *When ChatGPT misdescribes Brand X to a shopper, who committed the deceptive act?* OpenAI is not selling Brand X's products. Brand X never signed anything with OpenAI. Section 5 was built for sellers making claims about their own goods, not for a platform making claims about a brand that is not even its customer. That unanswered question is the gap, stated plainly.

Sentinel relevance. *The substantiation standard cuts in two directions. AI providers and Frontier model/GEO platforms face FTC exposure when they make claims they cannot back up. Sentinel's methodology is our sector's regulatory defensibility.*

6. SEC Posture on AI Disclosure

The SEC has landed in roughly the same place as the FTC. No AI statute, but a hard line on "AI washing" through the antifraud rules it already enforces, Section 10(b) and Rule 10b-5, aimed at public companies and investment advisers that oversell their AI. Under Chair Paul Atkins and the renamed Cybersecurity and Emerging Technologies Unit, AI washing stayed a priority even as the agency talked up a broader "return to basics."

Three standards have settled out. AI claims in anything investor-facing have to be accurate and backed up. They have to stay consistent across the annual report, the website, the press releases, and the marketing deck. And the company needs real internal controls built for AI claims, with management and the board paying attention. The case that makes it concrete: Presto Automation, in January 2025, was

the first AI-washing action against a public company. The SEC said Presto oversold its Presto Voice product and hid that the underlying speech recognition routed to a third party and leaned heavily on human intervention. Nate Inc., that April, became the first joint SEC and DOJ AI-washing case under the new administration. Securities class actions over AI claims doubled from 2023 to 2024.

The SEC does not regulate AI brand representation as such. It regulates how a public company discloses what is material to its business. Once AI brand harm grows large enough to matter to the business, the disclosure rules switch on. A consumer brand watching its AI visibility erode, or staring at real reputational exposure from how AI portrays it, is edging toward a disclosure obligation under the risk-factor rules that already exist.

Sentinel relevance. *For public companies, AI brand representation is shifting from a concern in marketing to a disclosure concern. When the harm becomes material, it becomes disclosable, and inaccurate or absent disclosure becomes its own securities law exposure. This is the clearest path by which AI brand governance becomes a risk and legal function rather than a marketing function.*

7. The Shared Gap

Five regimes, all active, and not one of them aims squarely at the harm at the center of AI brand representation. The EU AI Act builds the vocabulary but leaves brand recommendation off the high-risk list. New York's Local Law 144 is the model for algorithmic audits, but it only covers hiring. The state-level legislations reach most high-risk decisions, unevenly, and never the moment a consumer asks an AI what to buy. The FTC has the most flexible authority of any U.S. regulator and still no theory for third-party AI harming a brand that is not its customer. The SEC intervenes only after the damage to a public company has materialized.

These rules were built to protect consumers and workers from AI. They were never built to protect brands harmed by AI. As a regulatory problem, this is a novel and under-discovered. Regulation is lagging the risk. Which leaves brands in an odd spot. Waiting to comply with rules that do not exist is not a plan, and the exposure is already real: defamation, securities disclosure, FTC scrutiny of a brand's own AI claims. The brands that build governance before the rules arrive will be the ones standing ready when they do. The work it takes to manage the risk today is the same work whatever framework eventually shows up will demand. Getting there first, governance ahead of regulation instead of compliance after it, is the space Sentinel occupies.

Sources

Primary AI Legislation

1. Regulation (EU) 2024/1689 (EU AI Act), Official Journal of the European Union.
eur-lex.europa.eu/eli/reg/2024/1689/oj
2. NYC Local Law 144 of 2021 (Automated Employment Decision Tool Bias Audit Law). rules.cityofnewyork.us
3. Colorado AI Act (SB 24-205), as amended. leg.colorado.gov
4. California AI Transparency Act (SB 942, amended by AB 853). leginfo.legislature.ca.gov
5. California Transparency in Frontier Artificial Intelligence Act (SB 53). leginfo.legislature.ca.gov
6. Texas Responsible AI Governance Act (TRAIGA, HB 149). capitol.texas.gov
7. Tennessee SB 1580; Georgia SB 444; Illinois Artificial Intelligence Video Interview Act. State legislative records

Regulatory Enforcement and Guidance

8. FTC, "FTC Announces Crackdown on Deceptive AI Claims and Schemes" (Operation AI Comply launch), September 25, 2024. ftc.gov
9. FTC enforcement actions against DoNotPay, IntelliVision, Workado, Air AI, and Growth Cave (2025-2026). ftc.gov press releases
10. SEC v. Presto Automation, Inc. (January 2025); SEC v. Saniger / Nate Inc. (April 2025). sec.gov
11. New York State Comptroller, audit of DCWP Local Law 144 enforcement record, December 2, 2025. osc.ny.gov
12. SEC Division of Examinations, 2026 examination priorities, November 2025. sec.gov

GOSH AI Materials

13. GOSH AI, "The Authority Gap: State of AI Visibility, Vol. 1." mygosh.ai/the-authority-gap
14. GOSH AI, ARDI Platform Overview. mygosh.ai/gosh-ai-platform

Research and Industry Sources

15. Aggarwal, P. et al., "GEO: Generative Engine Optimization," Princeton and Georgia Tech, 2024.
arxiv.org/pdf/2311.09735
16. National Conference of State Legislatures, AI Legislation Tracker (2025-2026). ncsl.org

Legal Commentary

17. DLA Piper, "SEC emphasizes focus on AI washing despite perceived enforcement slowdown," May 2025.
18. Norton Rose Fulbright, "SEC heightens enforcement for AI related disclosures," 2025.
19. Holland & Knight, "2025 Cybersecurity and AI Year in Review," December 2025.
20. Benesch, "One Year In, FTC's Operation AI Comply Continues Under New Administration," October 2025.
21. White & Case LLP, "Key considerations for updating 2025 annual report risk factors."
22. D&O Diary, "AI, the SEC, and the 2026 Reporting Season," February 2026.