The Center for
AI Oversight

CENTER PAPER NO. 2 · JUNE 2026

The AI Regulatory Enforcement Landscape

The 2026 enforcement environment, and the one obligation stable across all of it. Deadlines move. The obligation does not.

By **Brian J. Allen**, Executive Director

> *This paper moves in three tiers. A **board line** states the governance proposition in a sentence; a **why it matters** note explains the consequence; a shaded **detail** block carries the authority. It takes an analytical posture throughout: it explains the mechanics of contested questions, including federal preemption, without advocating outcomes.*

EXECUTIVE SUMMARY

Volatility is the argument for governing at the oversight layer

KEY TAKEAWAYS

1. In twelve months the EU deferred its high-risk AI rules by sixteen months, the White House moved to preempt state AI laws, and Colorado repealed its flagship statute six weeks before it took effect. Volatility is the system's steady state.
2. An institution that builds AI governance to a statute rebuilds it every legislative season. This year required three rebuilds.
3. Four forces are converging: SEC disclosure, federal preemption, the state patchwork, and sector supervision, the quietest force and the most stable.
4. Every regime, however much it differs on dates and definitions, presumes the same four things: know your AI, name who is accountable, escalate what matters, keep the record.
5. Those four demands are the oversight layer. They did not change through a year in which nearly everything else did. The program is the governance.

This paper is the regulatory companion to the Center's *Caremark AI Liability Roadmap*. Where the Roadmap traces how courts built the oversight standard, this analysis maps the enforcement environment independently demanding the same thing. The conclusion is not to wait for the volatility to resolve. It will not resolve. The conclusion is that regulatory volatility is itself the argument for programmatic oversight governance, because the obligation to govern is the only element stable across every regime surveyed here.

PART I

The SEC framework: the template, and the first enforcement

1.1 The cybersecurity progression as preview

United States technology-risk regulation moves from guidance to mandated disclosure of who oversees, how they are informed, and what happened when it failed. Cybersecurity walked the path AI will follow.

Why it matters. Interpretive guidance became Commission guidance became binding rules in twelve years, and the mandate's content is governance. Two facts keep the template honest in 2026: the SEC declined to mandate board-level expertise disclosure, and the incident-reporting mechanism is under reconsideration. The governance disclosure architecture is not the subject of that debate.

> **RULE**
>
> SEC Release No. 33-11216 (July 26, 2023): Form 8-K Item 1.05 (material incident disclosure within four business days of a materiality determination) and Regulation S-K Item 106 (processes, the responsible board committee, and management's role). The cybersecurity disclosure controls tie to the CEO and CFO certifications under Exchange Act Rules 13a-14 and 15d-14. Item 1.05 is under active reconsideration following rescission petitions and the Commission's January 2026 Regulation S-K reform statement.

1.2 AI-washing: enforcement without an AI statute

Misstatements about AI capability or AI governance are actionable today under existing law. No new statute is required to police the gap between what an institution says about its AI and what is true.

Why it matters. Governance claims are claims. An institution that describes rigorous board oversight of AI in a filing or on a website has made that oversight a disclosure matter, and the documentary record must be able to substantiate the sentence.

> **ENFORCEMENT**
>
> *In re Delphia (USA) Inc.* and *In re Global Predictions, Inc.* (SEC, Mar. 18, 2024; penalties of $225,000 and $175,000 under Advisers Act §§ 206(2), 206(4) and Rules 206(4)-1, 206(4)-7); *In re Presto Automation Inc.* (SEC, Jan. 2025); FTC Operation AI Comply (Sept. 25, 2024); Texas AG Assurance of Voluntary Compliance with Pieces Technologies (Sept. 2024), the first state AG AI action.

QUESTIONS FOR THE BOARD

- Does our AI risk disclosure describe who oversees AI, or only that AI is a risk?
- Can we substantiate every public statement we make about our AI governance, today, from our own records?

PART II

The federal posture: acceleration and preemption

Federal policy is doing two things at once: accelerating AI adoption and moving to centralize the rules under which it occurs. The result is that every state-law obligation now carries litigation exposure alongside its compliance date.

Why it matters. The federal government has answered the accountability question for its own AI in writing, a named officer and a current inventory, while contesting the states' authority to impose their own answers. This paper takes no position on how the constitutional questions resolve. The institutional consequence does not depend on the answer: building governance to any single statute's specifications is a wager on that statute's survival.

> FEDERAL ACTION
>
> America's AI Action Plan (July 2025); OMB Memoranda M-25-21 and M-25-22 (Apr. 3, 2025), requiring agency Chief AI Officers, public use-case inventories, and minimum risk practices, with documentation flowing to vendors through procurement; Executive Order on advanced AI innovation and security (June 2, 2026). Executive Order 14365 (Dec. 11, 2025) directed an AI Litigation Task Force to challenge state AI laws; 36 state attorneys general opposed it in a joint letter (Nov. 25, 2025). Executive orders do not themselves preempt state law; preemption ordinarily requires a federal statute, and Congress has not enacted one.

PART III

The state patchwork: expansion under fire

3.1 Colorado: the sequence that proves the thesis

Colorado enacted the first comprehensive state AI statute, and it never applied to anyone. The institutions that built oversight programs kept everything when it was repealed; the institutions that built compliance artifacts lost them.

Why it matters. Read at the oversight layer, the Colorado sequence is not chaos; it is instruction. The inventory survived the repeal because the replacement still turns on knowing which systems matter; the accountability structure survived because duties remain allocated by role; the documentation discipline survived because notices and records are the replacement's core. And the impact assessments did not die, they migrated, reappearing in California privacy regulation. Obligations outlive the statutes that carry them.

> RULE

Colorado SB 24-205 (2024) was deferred, then its enforcement paused by stipulated court order (D. Colo., Apr. 27, 2026) in a constitutional challenge. SB 26-189 (signed May 14, 2026; effective Jan. 1, 2027) repealed and replaced it with a narrower notice-based framework, eliminating the duty of care, the impact assessments, and the compliance presumption.

3.2 The statutes in force

The in-force state laws disagree about nearly everything operational, intent versus impact, notice formats, audit rules, and agree about the oversight layer. Two of them now make the documented program a legal defense.

Why it matters. Each statute presumes the institution knows its covered systems, has allocated responsibility by role, can act on what comes in, and can produce a record. The variation across them is the strongest practical argument for building governance above the statute rather than rebuilding it state by state.

IN FORCE

Texas Responsible Artificial Intelligence Governance Act, Tex. Bus. & Com. Code chs. 551-552 (eff. Jan. 1, 2026), making documented substantial alignment with the NIST AI RMF a statutory defense; California SB 53, Cal. Bus. & Prof. Code § 22757.10 et seq. (operative Jan. 1, 2026); New York RAISE Act (eff. Jan. 1, 2027; 72-hour incident reporting); NYC Local Law 144 (bias-audit publication since 2023); Illinois H.B. 3773 / Pub. Act 103-0804 (eff. Jan. 1, 2026; unintentional disparate impact; ZIP codes barred as proxies); Connecticut Public Act 26-15 (signed June 2, 2026; documented anti-bias testing credited in mitigation). More than 1,000 state AI bills were introduced in 2025.

QUESTIONS FOR THE BOARD

- If a statute we built our AI compliance around were repealed tomorrow, like Colorado's, what would survive?
- Have we structured our governance to absorb a new state statute, or to be rebuilt by one?

PART IV

Sector supervision: the stable track

The quietest regulatory force is the one this year's upheaval left untouched. While the comprehensive-statute track produced a repeal, a federal challenge, and a deferral, sector supervision moved through 2026 without a single reversal.

Why it matters. Institutions should notice which track their examiners walk. Insurance regulators expect a written AI program in over half the states; banking rewrote its model risk guidance and kept the core principles intact. The stable obligations are the ones least likely to move under preemption pressure.

SECTOR AUTHORITY

NAIC Model Bulletin on the Use of AI Systems by Insurers (Dec. 4, 2023), adopted guidance in 24 states and D.C., with an examiner evaluation tool piloting in twelve states; Colorado Division of Insurance Regulation 10-1-1 (3 CCR 702-10), which survived the repeal of the state's comprehensive statute; NYDFS Insurance Circular Letter No. 7 (July 11, 2024) and cybersecurity Industry Letter (Oct. 16, 2024); ACA Section 1557 final rule, 89 Fed. Reg. 37522 (May 6, 2024).

BANKING

Interagency Guidance on Third-Party Relationships, 88 Fed. Reg. 37920 (June 6, 2023); SR Letter 26-2, Revised Guidance on Model Risk Management (Apr. 17, 2026), superseding SR 11-7 after fifteen years. The revision kept tiered model inventory and the non-transfer principle, moved to a risk-based cadence, and placed generative and agentic AI outside the framework. The principles that survive a fifteen-year rewrite are the durable layer; the systems excluded from it are the open question.

PART V

The European Union: the dates moved, the architecture did not

The EU AI Act remains the most comprehensive AI statute in force. Its high-risk deadlines were deferred by sixteen months in May 2026, but the architecture institutions must build toward did not change at all.

Why it matters. An institution that calendared its EU program to August 2026 has now re-planned twice and may re-plan again at formal adoption. An institution that built the classification and documentation disciplines as oversight elements has nothing to unwind. The deferral is the EU's contribution to this paper's thesis.

RULE

Regulation (EU) 2024/1689. Prohibited-practice provisions apply since February 2025; general-purpose AI obligations (Arts. 51-55) since August 2025. The Digital Omnibus on AI reached provisional agreement on May 7, 2026 (Council endorsement May 13), deferring Annex III high-risk obligations to December 2, 2027 and Annex I embedded systems to August 2, 2028. The agreement is provisional in the precise legal sense: until the amending regulation is published in the Official Journal, expected before August, the original August 2, 2026 date remains the binding text.

PART VI

What the landscape demands

Surveyed end to end, the regimes converge without coordination on four demands: know what AI acts on your behalf, name a human body accountable for it, move material events to that body on defined clocks, and keep the record. Those four are the oversight layer.

Why it matters. The demands fall differently on the market's two positions, a distinction most regimes now draw explicitly. Builders carry the frontier statutes and developer documentation; Buyers carry the deployer obligations and the non-transfer principle. Most regulated institutions are Buyers, which is why the Center develops that position in a companion paper. For both, the strategic conclusion is identical: building to a statute means rebuilding each season; building at the oversight layer means absorbing the season.

Regulatory volatility is not the reason to wait. It is the reason the program exists, and the program is the governance.

QUESTIONS FOR THE BOARD

- Across every regime that reaches us, can we show the same four things: inventory, accountability, escalation, record?
- Are we a Builder, a Buyer, or both, for each AI system, and have we allocated the right obligations to each?

Sources and Further Reading

Federal: Executive Order 14365 (Dec. 11, 2025); Executive Order on advanced AI (June 2, 2026); America's AI Action Plan (July 2025); OMB M-25-21, M-25-22 (Apr. 3, 2025); SEC Release No. 33-11216 (July 26, 2023); SEC Regulation S-K reform statement (Jan. 13, 2026); SEC AI-washing orders (2024-2025); FTC Operation AI Comply (2024).

State: Colorado SB 24-205 and SB 26-189; Colo. Div. of Ins. Reg. 10-1-1; Tex. Bus. & Com. Code chs. 551-552; Cal. Bus. & Prof. Code § 22757.10 et seq.; New York RAISE Act; NYC Local Law 144; Illinois Pub. Act 103-0804; Connecticut Pub. Act 26-15; CPPA regulations (adopted July 24, 2025).

Sector and international: NAIC Model Bulletin (Dec. 4, 2023); NYDFS Circular Letter No. 7 (2024); Interagency Third-Party Guidance, 88 Fed. Reg. 37920 (2023); SR Letter 26-2 (Apr. 17, 2026); ACA Section 1557 rule, 89 Fed. Reg. 37522 (2024); Regulation (EU) 2024/1689; Digital Omnibus on AI provisional agreement (May 7, 2026). The current status of every obligation cited is maintained in the Center's standing publication, *The AI Oversight Obligations Reference*.