Markets Need Referees: An SEC-Model Commission for AI

The **American Artificial Intelligence Leadership and Accountability Act of 2026**, establishing the Artificial Intelligence Oversight Commission (AIOC)

Published jointly by **THE ALIGN AI FOUNDATION** *and* **THE GIE FOUNDATION**

THE MOMENT

- **Voters of both parties already agree.** 80% of Democrats and 84% of Republicans support a new federal AI oversight agency (Public Citizen, Apr. 2026); 84% and 83% say companies should not deploy smarter-than-human systems without demonstrated control (NBC News, 2026).
- **The state patchwork is real — and unstable in both directions.** 40 states enacted AI laws in 2025; California, New York, Colorado, and Texas run four incompatible regimes; Colorado's law was rewritten in May 2026 before it ever took effect. Industry can't plan; the public can't rely.
- **Congress has priced the deal — twice.** The Senate stripped bare preemption 99–1 (2025); it failed again in the NDAA. The bipartisan Obernolte–Trahan draft (June 2026) trades a federal framework for a state-law freeze — but creates *no institution* to administer the bargain.
- **The executive branch concedes the gap.** The June 2, 2026 executive order asks frontier developers for *voluntary* 30-day pre-release access while expressly disclaiming licensing authority; Treasury officials reportedly favor a FINRA-style body; industry leadership proposed a supervised standards body in July 2026. Everyone is circling the same missing referee.

THE PROPOSAL

Do for AI what the Securities Exchange Act of 1934 did for capital markets: create an expert, bipartisan, independent commission — **five members, no more than three from one party, staggered terms, majority fee-funded** — built on the SEC's core insight: **mandated disclosure and accountable process, not government product approval**. The SEC doesn't pick good stocks; the AIOC doesn't pick good models. It makes the developers of the most capable systems tell the truth — in standard formats, on fixed schedules, with real penalties for deception — and puts catastrophic-risk claims in front of someone other than the companies making them. The default is deployment: the agency bears the burden, on the record, on a clock, under expedited judicial review.

WHO IS COVERED

Tier	Who	Obligations
Tier 0 — Exempt	Startups, researchers, open-source contributors, deployers — everyone below Tier 1	None. Express statutory exemption.
Tier 1 — Covered	Training runs >10^{25} FLOP *and* >$100M covered-AI revenue	Registration + standardized annual disclosure (a Form 10-K analog); truthfulness with penalties.
Tier 2 — Frontier	Training runs >10^{26} FLOP and >$500M revenue (aligned with CA SB 53)	Plus: pre-deployment safety case, 72-hour incident reports, third-party audit, examinations.

Thresholds recalibrated biennially by rule, in either direction. Coverage turns on capability and scale — never on open-weight release as such, and never on model viewpoint or content.

THE MACHINERY (NINE BILL TITLES)

- **Registration & disclosure** — annual standardized filings on capabilities, evaluations, safety and security practices (Title II).
- **Frontier safety review** — a pre-deployment Frontier Model Safety Case (an S-1 analog), **deemed effective in 90 days** unless the Commission acts; stop orders only on clear-and-convincing evidence of defined catastrophic risk, with 60-day D.C. Circuit review; codifies uniform pre-release evaluation access (Title III).
- **Incident reporting** — 72-hour critical-incident reports (an 8-K analog) with a good-faith safe harbor; one federal filing satisfies every state regime (Title IV).
- **Exams, audits, enforcement** — accredited third-party audits; cease-and-desist; penalties to the greater of $10M or 4% of covered-AI revenue for knowing misstatements; Dodd-Frank-style whistleblower awards (10–30%) overriding NDAs for safety disclosures; **no new private right of action** (Title V).
- **Supervised self-regulation** — a FINRA-style standards SRO *under* the Commission — rule approval, abrogation, appeals — not instead of it (Title VI).
- **Partial preemption, earned** — registrants in good standing get uniform national rules for the development layer; states keep consumer-protection, civil-rights, tort, child-safety, criminal, and use-level law in full; state AGs may enforce the federal floor; preemption lapses for willful violators (Title VII).
- **Innovation in statute** — regulatory sandbox (SANDBOX Act / TRAIGA model), statutory small-entity exemption, red-team research safe harbor, Office of Innovation with a reporting advocate, biennial GAO review, and a **seven-year sunset** forcing reauthorization on the record (Title VIII).

WHAT EACH SIDE GETS

Conservatives	Progressives
• National uniformity with substance — the preemption two moratorium attempts couldn't deliver • Disclosure, not licensing; deployment by default, agency on a clock • Markets price risk: insurers, customers, investors read standardized filings • Lean, fee-funded, sunsetted; sandbox and small-entity exemptions in statute • Statutory national-security visibility the June 2026 EO can only request	• A real regulator: exams, subpoenas, penalties, stop-order power — not a pledge registry • Whistleblower awards and anti-retaliation protections that override lab NDAs • Mandatory 72-hour incident transparency, publicly summarized • Preemption that is partial, earned, and revocable — states keep every consumer tool • Public-interest seats, public comment on SRO rules, published enforcement outcomes

COST AND TIMELINE

Launch lean: 300–500 staff, $150–250M/year (vs. the SEC's ~4,500 staff for tens of thousands of registrants — the AIOC's launch population is measured in dozens), majority fee-funded by year three. Commissioners confirmed in 180 days; disclosure and incident rules in year one; safety-case filings at 18 months, with preemption attaching developer-by-developer upon compliant registration — **the benefit to industry arrives exactly as fast as the transparency arrives to the public**.

THE BOTTOM LINE

AI is being governed today — by four state legislatures, a DOJ litigation task force, plaintiffs' lawyers, and foreign parliaments. The choice is not regulation or none; it is institution or improvisation. In 1934, Congress gave markets what they needed to grow and the public what it needed to trust them, and American disclosure standards became the world's benchmark for ninety years. The same terms are available now, and four-fifths of both parties' voters are already at the destination. **Congress should take them.**

The Align AI Foundation funds independent research on AI safety and alignment — interpretability, secure evaluation, and deception risk — on a theory-agnostic, evidence-driven basis, with no corporate revenue. **The Geopolitical Insight and Education Foundation** is an independent, nonpartisan Washington think tank, founded in 2024, researching the drivers of instability — technological, environmental, and geopolitical — and their effects on democracies and their institutions.

1717 N Street NW, Suite 1, Washington, DC 20036 · info@giefoundation.net · giefoundation.net · alignai.org

Full 22-page position paper with section-by-section framework, objections and responses, and sources available on request.