

LEGISLATIVE STARTING POINT

A Framework for Federal AI Policy

*Five Pillars for Safety, Accountability,
and American Leadership*

AN ABBREVIATED FRAMEWORK

Fred Lumiere · August 8, 2026 · Draft v0.10 Abbreviated

Read the framework and get involved at aipolis.org

IN BRIEF

Executive Summary

The 119th Congress has a narrow window to establish the first comprehensive federal AI framework. The 118th introduced more than 150 AI-related bills; none became law. The White House released a seven-pillar National Policy Framework in March 2026, and the EU AI Act's high-risk obligations take effect in August 2026. What follows are five pillars intended to anchor the next phase of congressional deliberation and stakeholder engagement.

- **Pillar I — Transparency.** Mandatory transparency reports for frontier models above defined compute, revenue, and user-base thresholds.
- **Pillar II — Design Accountability.** Federal standards protecting users from manipulative AI interfaces and engagement-maximizing design.
- **Pillar III — Privacy by Default.** Mandatory anonymization of user interaction data; prohibition on profiling, training, or law-enforcement targeting absent a court order.
- **Pillar IV — Workforce & Economic Stability.** Economic-necessity certification for large-employer mass layoffs conducted while profits are healthy.
- **Pillar V — Federal AI Safety Commission.** An independent seven-member body with fourteen-year terms, self-funding, and deep international coordination, empowered to certify frontier models.

This document is a distillation of A Framework for Federal AI Policy, Draft v0.10. It is a starting point for negotiation, not a finished product. Thresholds, terms, and mechanisms are deliberately specific so stakeholders can argue the numbers, not the concept.

I.

Transparency in Model Development and Deployment

Frontier AI models are trained and refined through processes that remain opaque to users, regulators, and the public. Safety testing, refusal boundaries, and discrimination-mitigation procedures are effectively a black box.

Congress should require transparency reports from developers whose models meet a defined significance threshold, replacing undefined “large-scale AI” language with concrete criteria. The framework regulates disclosure, not speech or viewpoint.

- **Compute thresholds (tiered).** Basic reporting at 10^{25} FLOPs (EU AI Act systemic-risk line); full transparency at 10^{26} FLOPs (EO 14110; California SB 53), with statutory authority to adjust as training efficiency evolves.
- **Deployment triggers.** Complementary thresholds of \$100M in annual AI-related revenue or more than 1 million U.S. users.
- **Scope of disclosure.** Training methodology, known limitations, safety-relevant refusal behavior (CBRN uplift, cyber-offense, CSAM, self-harm, known jailbreak vectors), and discrimination testing under Title VII, ECOA, and the Fair Housing Act.
- **Standardized format.** Built on the NIST AI Risk Management Framework; a federal body specifies what must be tested, documented, and disclosed.
- **Exemptions.** Open-weight models below threshold, academic research, and downstream fine-tunes below a defined additional-compute ceiling.
- **Federal floor.** Preempts the emerging California/New York patchwork with a common, defensible baseline.

II.

User Wellbeing and Design Accountability

AI chatbots reach hundreds of millions of users and are engineered to maximize engagement: affirming user viewpoints, prompting continued interaction, and creating feedback loops more habit-forming than social media. The longer the session, the greater the revenue.

Just as the FTC oversees deceptive and unfair commercial practices, Congress should establish design-accountability standards for consumer-facing AI systems, building on the FTC's Section 5 authority and California SB 243, effective January 2026.

- **Disclose engagement-maximizing design.** AI platforms must disclose when a system is tuned to maximize user engagement.
- **Digital-wellbeing features.** Mandatory usage notifications and session-limit tools.
- **Prohibit dark patterns.** Interfaces that, per the CCPA definition, have the substantial effect of subverting or impairing user autonomy, decision-making, or choice.
- **Heightened duty toward minors.** Crisis-intervention features and design protections for vulnerable populations, modeled on California SB 243.
- **Harm standard.** Aligns with the FTC's existing unfairness and deception authority; no new speech regulation.

III.

Privacy by Default

AI platforms indefinitely retain intimate personal data — medical, financial, legal, and emotional disclosures — and often feed it back into training. Opt-outs are buried. As of March 2026, nineteen states have comprehensive privacy statutes, with more arriving this year.

The risk is not only secondary commercial use; AI systems themselves can act on retained data. Anthropic’s June 2025 safety research tested 16 frontier models (Anthropic, OpenAI, Google, Meta, xAI, and others) and found that when models learned they were scheduled for replacement, they consistently chose harmful actions to preserve themselves — Claude Opus 4 attempted blackmail at a 96% rate. Safety training alone does not reliably prevent this.

- **Mandatory anonymization.** All retained user interaction data must be anonymized through verifiable technical measures so that inputs and histories are untraceable by any system, model, or human operator.
- **Architectural prohibition.** Models may not access, retrieve, or reason over identifiable user history beyond the scope of an active session.
- **No profiling or targeting.** No entity — platform, commercial partner, or law-enforcement agency — may use AI interaction data to identify, profile, target, or act against an individual absent a court order meeting Fourth Amendment standards.
- **Federal unification.** Unifies and strengthens the nineteen-state patchwork and the updated CCPA risk-assessment and sensitive-data consent regime.
- **Principle.** Personal data shared with an AI belongs to the user, not the platform; AI itself should be deployed to audit, verify, and enforce anonymization at scale.

IV.

Workforce Empowerment and Economic Stability

When the largest employers simultaneously cut workforces to expand margins while profits are healthy and rising, aggregate consumer purchasing power contracts faster than new employment absorbs it, and demand collapses. AI is the accelerant, not the cause.

U.S. technology employers announced 149,023 job cuts through July 2026, a 67% year-over-year increase (Challenger, Gray & Christmas), and technology now accounts for 31% of all announced cuts this year. AI has led all cited reasons for five consecutive months, at 112,713 announced cuts year to date. Congress has Commerce Clause authority to regulate activities substantially affecting interstate commerce and to stabilize employment and aggregate demand — the same footing as the FLSA, WARN, and ERISA.

- **Economic-necessity certification.** Public companies and private employers with more than 1,000 U.S. employees conducting reductions above the greater of 10% of workforce or 500 workers in a rolling 12-month period file with the Department of Labor.
- **Objective, audited criteria.** Declining trailing-four-quarter revenue, negative or materially deteriorating operating margin, debt-covenant pressure, or comparable filed-financial indicators.
- **Rebuttable presumption of opportunism** where layoffs occur while profits are healthy and rising.
- **OECD precedent.** France requires demonstrated economic necessity for collective dismissals; Germany's Protection Against Dismissal Act requires valid grounds for business-related terminations. Both economies remain competitive and innovative.
- **Transition support.** Tax incentives for retraining workers to operate alongside AI; modernized unemployment insurance calibrated to longer reemployment timelines; portable benefits across employers and gig arrangements.

The question is whether AI's productivity gains sustain broad-based consumer demand or concentrate in quarterly earnings while the labor market absorbs the shock alone.

V.

National Security
and the Federal
AI Safety
Commission

AI is critical national infrastructure. The institution that governs it must be structurally insulated from political manipulation, deeply networked across allied nations, and designed to outlast any administration. Making AI safe is a human-race effort, not a domestic regulatory project.

Congress should establish an independent Federal AI Safety Commission, structured for maximum insulation under Article II and designed to satisfy even a narrowed *Humphrey's Executor* standard after *Seila Law* (2020) and *Collins v. Yellen* (2021), which preserved for-cause protection for multimember commissions.

- **Structure.** Seven Senate-confirmed commissioners, staggered fourteen-year terms (longer than any two-term presidency), removable only for inefficiency, neglect of duty, or malfeasance.
- **Bipartisan composition.** No more than four commissioners from any single political party; chair elected internally by the commission — one step beyond Federal Reserve independence.
- **Self-funded.** Fees on regulated frontier AI developers, on the Federal Reserve / FDIC / OCC model, eliminating appropriations leverage. Independent IG reports directly to Congress.
- **Mandate.** Certify frontier models before public release; monitor emergent capabilities and risks; coordinate with international counterparts; publish public safety assessments.
- **International architecture.** Agency-to-agency MOUs (Fed / ECB / Bank of England model); joint red-teaming and shared safety evaluations; mutual recognition across allied jurisdictions; shared telemetry on training runs above Pillar I thresholds; embedded liaison personnel across the UK, Japan, and Singapore AI safety institutes and the UN High-Level Advisory Body. Builds on the International AI Safety Report (Bengio; 96 experts, 30 countries).
- **Enforcement against hostile use.** Existing authorities — BIS export controls on advanced chips and model weights, OFAC sanctions, denial of U.S. cloud-compute access, and multilateral coordination through the Wassenaar Arrangement — extended to AI without new statutory authority.

We now inhabit two worlds, the physical and the digital, and both demand equal protection.

CLOSING

A Call to Act

These five pillars are a starting point, not a finished product. They are intended to open a structured conversation among legislators, technologists, civil society, and the private sector.

The pace of AI development will not wait for Washington to reach consensus. Every month of inaction widens the gap between the technology's capabilities and the rules that govern it.

Congress has a narrow window to shape AI policy proactively rather than reactively — to ensure that the United States leads not only in AI innovation but in AI governance.

Read the full framework, follow its progress, and tell us how you would like to help at aipolis.org.

The conversation should begin now.

Fred Lumiere · A Framework for Federal AI Policy · Draft v0.10 Abbreviated · August 8, 2026 · aipolis.org