ML STRATEGIES

Charting the Future of AI Governance: California's SB 53 Sets a National Precedent — AI: The Washington Report

October 07, 2025 | Article | By **Bruce Sokler**, **Alexander Hecht**, **Christian Tamotsu Fjeld**, Nicole Y. Teo

VIEWPOINT TOPICS

- Artificial Intelligence
- AI: The Washington Report

- On September 29, Governor Gavin Newsom signed into law Senate Bill 53 (SB 53), the **Transparency in Frontier Artificial Intelligence Act (TFAIA),** making California the first state in the US to impose targeted AI regulations on major industry players with advanced AI systems. The law requires the public disclosure of safety standards by AI developers, establishes a consortium to develop a framework for creating a public computing cluster, creates a formal mechanism for reporting safety incidents, protects whistleblowers who raise concerns about critical risks, and requires annual recommendations for updates to the law.
- California houses 32 of the world's top 50 AI companies, making SB 53 especially impactful as it establishes the first legal framework in the US focused on the safety of frontier AI models.
- Beyond its immediate impact in California, TFAIA is poised to serve as a regulatory blueprint for other states to follow, and could also shape potential federal AI legislation.
- The passage of SB 53 reflects broader tensions between current state and federal approaches to AI regulation. TFAIA introduces a new regulatory floor and framework for AI regulation after Congress has failed to pass either substantive AI legislation over the last few years, or the 10-year moratorium on state-level AI laws as part of the One Big Beautiful Bill Act this summer, as **we've previously covered**.
- The new California law also sharply diverges from the deregulatory stance favored by the Trump administration, which has advocated for minimal federal oversight of AI technologies to accelerate national AI competitiveness and infrastructure, as **we've covered** in our overview of the White House AI Action Plan. It is unclear whether the California law may reinvigorate attempts to pass legislation establishing the primacy of the federal approach.

Background of the California Legislation

On September 29, Governor Gavin Newsom signed into law Senate Bill 53 (SB 53), the **Transparency in Frontier Artificial Intelligence Act (TFAIA),** making California the first state in the US to impose targeted AI regulations on major industry players with advanced AI systems. The law requires them "to fulfill transparency requirements and report AI-related safety incidents." TFAIA was shaped by a March report commissioned by Governor Newsom, which brought together a working group of leading AI scholars and experts to develop workable guardrails based on empirical research into "the capabilities and attendant risks of frontier models."

TFAIA introduces a comprehensive framework for AI oversight in California. It requires the public disclosure of safety standards by AI developers, establishes a consortium to develop a framework for creating a public computing cluster, creates a formal mechanism for reporting safety incidents, protects whistleblowers who raise concerns about critical risks, and requires annual recommendations for updates to the law.

The new California law imposes its most stringent requirements on "large frontier developers," a category within frontier AI developers that meet specific thresholds for both model compute capacity and annual revenue.[1]

An earlier version of the measure, SB 1047, was vetoed by Governor Newsom almost exactly a year ago due to concerns that its broad scope and strict enforcement, including a "kill switch" for AI systems that were misused or went rogue, along with a coverage of startups for liability, could stifle innovation. Governor Newsom argued that SB 1047's "requirements amounted to 'stringent' regulations [that] could burden the state's leading artificial intelligence companies, as Silicon Valley competes in the global AI race." In contrast, SB 53 narrowed its scope to large frontier developers (those with over $500 million in annual revenue), exempting smaller startups and easing industry concerns.

Focusing on frontier AI models, the most advanced and computationally intensive systems, SB 53 introduces the following key provisions:

- **Transparency Requirements:** Large frontier developers must publicly release a *Frontier AI Framework* detailing how they assess and mitigate catastrophic risks. The legislation also requires developers to disclose on their websites how the company has "incorporated national standards, international standards, and industry-consensus best practices into its frontier AI framework." The transparency report must be published before or concurrently with deploying a new frontier model or a substantially modified version of an existing model.
- **Safety and Incident Reporting:** Developers must report critical safety incidents to California's Office of Emergency Services within 15 days and notify law enforcement within 24 hours if imminent harm is identified.
- **Whistleblower Protections:** Employees who report concerns about critical risks or misleading safety claims of the models are protected from retaliation. These protections are enforceable by the California Attorney General's Office.
- **CalCompute Initiative:** The bill establishes a state-run public cloud computing cluster housed at the University of California to support safe and equitable AI research. The initiative aims to "advance the development and deployment of artificial intelligence that is safe, ethical, equitable, and sustainable by fostering research and innovation." CalCompute will offer free and low-cost access to startups and academic researchers and may accept private donations to support its implementation.
- **Enforcement:** Violations of SB 53 can result in civil penalties up to $1 million, enforced by the California attorney general.

California Continues to Lead State AI Regulatory Efforts

California houses 32 of the world's top 50 AI companies, making SB 53 especially impactful as it establishes the first legal framework in the US focused on the safety of frontier AI models. As Governor Gavin Newsom asserted in the **press release**, this legislation enables California to "protect our communities while also ensuring that the growing AI industry continues to thrive." He further affirmed California's leadership in the field, declaring that "AI is the new frontier in innovation, and California is not only here for it — but stands strong as a national leader by enacting the first-in-the-nation frontier AI safety legislation."

Industry-wise, leading AI labs did not publicly oppose SB 53's passage, while others have gone so far as to backing the measure in the state legislative session.

While SB 53 is the first US law to explicitly address the safety of frontier AI models, other states have also begun to explore AI regulations:

New York: New York's **Responsible AI Safety and Education (RAISE) Act**, currently awaiting the governor's signature, includes safety protocols and incident reporting for frontier models and imposes fines up to $30 million, enforceable by the attorney general. However, unlike the Californian law, there are no whistleblower protections that are the cornerstones of TFAIA.

Texas: **Texas's Responsible Artificial Intelligence Governance Act (TRAIGA)**, signed into law in June and effective January 2026, bans social scoring, biometric data capture without consent, and AI tools promoting self-harm, and allows the attorney general to fine violators up to $10,000. Where the California legislation focuses on frontier model safety and targets major players in the industry, Texas's TRAIGA is focused on broader AI governance, consumer protection, and misuse.

Colorado: Colorado has acted on state-level AI regulation through the **Colorado Artificial Intelligence Act (CAIA),** also known as Senate Bill 24-205. Signed into law in May 2024 and set to take effect on February 1, 2026, CAIA targets developers and deployers of high-risk AI systems — those used in critical decisions like employment, housing, health care, and education. Its primary goal is to prevent algorithmic discrimination, requiring developers to report to the attorney general within 90 days if their systems cause or are likely to cause such harm. Like California's SB 53, CAIA includes transparency and disclosure requirements, but its focus is broader, emphasizing fairness and accountability in everyday AI applications.

State vs. Federal Approach to AI Governance

The passage of SB 53 also reflects the broader tension between state and federal approaches to AI regulation. TFAIA introduces a new regulatory floor and framework for AI regulation after Congress has failed to pass substantive legislation, or the 10-year moratorium on state-level AI laws as part of the One Big Beautiful Bill Act this summer, as **we've previously covered**. Senator Ted Cruz (R-TX), a key proponent of the 10-year moratorium, as of this publication has yet to comment on Governor Newsom's signing of TFAIA, but pointed out earlier in September at the Axios AI Summit that "there is no way for AI to develop reasonably, and for us to win the race to beat China, if we end up with 50 contradictory standards in 50 states — and not just 50 states because cities and municipalities will do this too." His criticism of the fragmented approach to AI governance highlights the tension between state and federal approaches to AI regulation.

SB 53 also diverges from the deregulatory stance favored by the Trump administration, which has advocated for minimal federal oversight of AI technologies to accelerate national AI competitiveness and infrastructure, as **we've covered** in our overview of the White House AI Action Plan.

Where California's SB 53 focuses on oversight and transparency of frontier AI models, the White House's AI Action Plan, released this summer, is a federal policy blueprint emphasizing national dominance in AI innovation, infrastructure, and diplomacy for the global AI race through deregulation. These approaches differ significantly: SB 53 introduces commonsense guardrails on frontier AI models, which are large-scale, high-performance AI systems capable of highly advanced tasks, while the AI Action Plan focuses on removing regulatory barriers to accelerate AI adoption across sectors, including over 90 federal actions.

The contrast in safety and risk management is also notable. SB 53 centers on catastrophic risk mitigation, requiring developers to assess and disclose risks that could result in mass harm or $1 billion in damages. The White House's AI Action Plan, on the other hand, supports research into AI interpretability and robustness but does not mandate safety disclosures or incident reporting for private developers. This reflects two competing visions for AI governance in the US: one led by states prioritizing accountability and public safety, and one by the federal government focused on strategic dominance and innovation through deregulation.

Beyond its immediate impact in California, TFAIA is poised to serve as a regulatory blueprint for other states and potentially for future federal legislation. In the absence of a unified federal AI policy, California's leadership in tech and its concentration of major AI firms gives its laws outsized influence. SB 53 fills the regulatory gap left by federal inaction while still being designed to align with future federal standards — "but only if they maintain or exceed the protections of the bill," said Governor Newsom. TFAIA's structured approach, requiring large AI developers to publish safety frameworks, report critical incidents, and protect whistleblowers, offers a scalable and adaptable model for responsible AI oversight. Its annual update mechanism, led by the California Department of Technology, ensures the law evolves alongside technological advancements and international norms. The Californian legislation demonstrates that meaningful regulation can coexist with innovation, and it challenges federal lawmakers to either match California's leadership or risk falling behind in shaping the future of AI governance.

This set of issues will continue to evolve, perhaps even quickly. We will continue to monitor, analyze, and issue reports on these developments. Please feel free to contact us if you have questions about current practices or how to proceed.


Endnotes

1 Frontier AI models are large-scale, high-performance systems capable of advanced tasks, including language generation, coding, and image synthesis. They are trained on massive datasets and often involve billions or trillions of parameters, making them among the most powerful and complex models that currently exist. Large frontier developers are those that have an annual gross revenue exceeding $500 million in the prior calendar year.

Authors

Bruce Sokler

Bruce D. Sokler is a Mintz antitrust attorney. His antitrust experience includes litigation, class actions, government merger reviews and investigations, and cartel-related issues. Bruce focuses on the health care, communications, and retail industries, from start-ups to Fortune 100 companies.

Alexander Hecht, Executive Vice President & Director of Operations
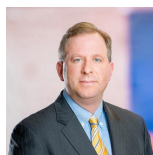
Alex Hecht is a trusted attorney and policy strategist with over 20 years of experience advising clients across a broad range of industries on how to navigate complex policy environments. His strategic insight and hands-on experience in both legislative and regulatory arenas empower clients to advance their priorities with clarity and confidence in an evolving policy landscape.

Christian Tamotsu Fjeld, Senior Vice President
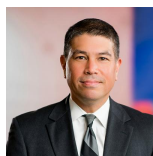
Christian Tamotsu Fjeld draws on two decades of Capitol Hill experience to support clients in building relationships, shaping policy, and engaging effectively with the federal government. His experience working with Congress and his insights help clients anticipate federal developments and advance their priorities with clarity and confidence.

Nicole Y. Teo

Nicole Y. Teo is a Mintz Senior Project Analyst based in Washington, DC.