

# New York's Frontier AI Law Gets a California Makeover – With Some Key Differences

April 1, 2026

New York is joining California in regulating the most advanced forms of artificial intelligence – frontier models – and creating a state-led US AI regulatory approach in the absence of federal legislation. On March 27, 2026, Gov. Kathy Hochul signed an amended version of the Responsible AI Safety and Education (RAISE) Act, substantially overhauling the original law she signed in December 2025. The result closely tracks California's Transparency in Frontier Artificial Intelligence Act (TFAIA) in many respects but diverges in a few important ways.

Hochul signed the original RAISE Act with an approval memo flagging concerns and conditioning her signature on an agreement with lawmakers to adopt chapter amendments in 2026. Those amendments are now law. Most changes align the RAISE Act with the TFAIA, which took effect January 1, 2026, as the first state law requiring standardized safety and transparency disclosures from frontier model developers. But the amended RAISE Act also makes important additions related to incident response, ownership disclosures, penalties and regulatory oversight.

These state laws are emerging as the federal government considers its own AI regulation. As [we discussed on March 25](#), the White House is urging Congress to preempt state AI laws, which could have an impact on the RAISE Act and the TFAIA. However, until Congress or courts act, both laws remain in force, and frontier AI companies should ensure their compliance programs are prepared.

Below, we outline the most notable changes in the amended RAISE Act and where those changes align with or diverge from the TFAIA.

## Notable RAISE Act changes aligning with the TFAIA

### **New framework and transparency requirements**

The amended RAISE Act replaces the original law's requirement to publish and maintain a "safety and security protocol," with a new obligation to publish a "frontier AI framework" – the same term used in the TFAIA. Under the original RAISE Act, the safety and security protocol had to specify protections to reduce the risk of "critical harm," describe cybersecurity measures, detail testing procedures, and designate senior compliance personnel. The amended law instead requires large frontier developers to publish a frontier AI framework describing how the developer handles "catastrophic risk" thresholds, mitigations, third-party evaluations, cybersecurity practices, critical safety incidents, and internal governance. These framework requirements, including the defined term "catastrophic risk," are copied from the TFAIA.

The amended RAISE Act introduces a transparency report requirement that has no counterpart in the original law and mirrors the TFAIA. Before or concurrently with deploying a new frontier model or a substantially modified version of an existing one, frontier developers must publish a report disclosing the model's release date, supported languages, output modalities, intended uses, and any restrictions on use. Large frontier developers must additionally include summaries of catastrophic risk assessments and the extent of third-party evaluator involvement. In the industry, these reports are often referred to as model cards.

### **Dropped audit requirement**

The original RAISE Act required large developers to annually retain a third party to perform an independent compliance audit and publish a redacted version of the resulting report. The amended law dropped this audit requirement entirely, aligning with the TFAIA, which contains no audit mandate.

## **Deployment restriction removed**

The original RAISE Act prohibited large developers from deploying a frontier model if doing so would “create an unreasonable risk of critical harm.” The amended law removes this prohibition. Like the TFAIA, the amended RAISE Act focuses on transparency and reporting rather than imposing prescriptive deployment restrictions.

## **Revised definitions for key terms**

The amended RAISE Act overhauls several key definitions. The original law’s “critical harm” threshold meant death or serious injury of 100 or more people or at least \$1 billion in damages arising from specific types of frontier model conduct related to chemical, biological, radiological or nuclear (CBRN) weapons or autonomous conduct. In the amended law, this term is replaced by “catastrophic risk,” which is defined as a foreseeable and material risk of death or serious injury to more than 50 people or more than \$1 billion in damage. In addition, the “catastrophic risk” definition relies on the same general categories of model conduct but revises the language to specify harm from a single incident involving a frontier model providing expert-level CBRN weapon assistance, engaging in specified criminal conduct without meaningful human oversight, or evading developer or user control. This definition of “catastrophic risk” is also found in the TFAIA.

The amended RAISE Act also revises which companies are covered. The original law defined “large developers” by compute cost, meaning those spending more than \$5 million on a single frontier model and more than \$100 million in aggregate. The amended version instead covers “large frontier developers,” defined – as in the TFAIA – that, together with their affiliates, had annual gross revenues exceeding \$500 million in the preceding calendar year.

The amended RAISE Act also changed the definition of “frontier model,” tracking California’s approach. The original RAISE Act captured any AI model trained using more than  $10^{26}$  computational operations and costing more than \$100 million in compute, or any model derived through knowledge distillation from a frontier model. The amended law removed both the dollar-cost threshold and the knowledge distillation prong, leaving in its place the compute threshold for the most advanced, emerging models. Under the amended law, a “frontier model” is now a “foundation model” trained using computing power of more than  $10^{26}$  operations, which includes “any subsequent fine-tuning, reinforcement learning, or other material modifications.” That definition is lifted nearly verbatim from California’s TFAIA. For developers, the explicit inclusion of fine-tuning and reinforcement learning runs in the compute count means post-training modifications could push a model across the threshold even if the original pre-training run did not.

## **Catastrophic risk assessment submissions**

The amended RAISE Act adds a requirement for large frontier developers to transmit summaries of catastrophic risk assessments from internal use of their frontier models to the regulator every three months or on another reasonable schedule agreed upon with the regulator. This matches the TFAIA’s requirements.

## **Federal reciprocity for incident reporting**

The amended RAISE Act introduces a federal reciprocity provision permitting frontier developers to satisfy the law’s “critical safety incident” reporting requirements by complying with federal laws, regulations, or guidance documents that impose “substantially equivalent” or stricter standards, as designated by the New York regulator. Developers electing the reciprocity path must send copies of any federal incident reports to the New York regulator concurrently. This mechanism mirrors the TFAIA’s approach.

# **Notable RAISE Act changes diverging from the TFAIA**

## **Reduction of monetary penalties**

The amended RAISE Act reduces maximum civil penalties from the original law’s \$10 million for a first violation and \$30 million for subsequent violations to \$1 million and \$3 million, respectively. While this is a significant reduction, the RAISE Act still permits higher penalties than the TFAIA, which caps all violations at \$1 million per violation, regardless of whether it is a first or subsequent offense.

## **Removal of frontier-specific whistleblower and employee protections**

The original RAISE Act contained a dedicated section protecting employees, including contractors, subcontractors, and corporate officers, from retaliation for reporting safety concerns to the developer or the New York attorney general. The amended law removes this section entirely. This is a notable departure from the TFAIA, which includes frontier-specific whistleblower protections, prohibiting frontier developers from retaliating against covered employees who disclose information about catastrophic risks or TFAIA violations, and requiring large frontier developers to maintain an anonymous internal reporting process.

### **Shorter incident reporting timelines**

The amended RAISE Act retains the original version's 72-hour window for reporting critical safety incidents but adds a 24-hour reporting requirement for incidents posing an "imminent risk of death or serious physical injury." The RAISE Act's 72-hour baseline is significantly shorter than the TFAIA's, which provides a longer baseline window of 15 days for critical safety incident reports. However, the TFAIA also requires 24-hour disclosure for incidents posing imminent risk of death or serious physical injury.

### **New large frontier developer disclosure requirement**

The amended RAISE Act adds a "large frontier developer disclosure" provision with no parallel in the TFAIA. Under the RAISE Act, large frontier developers may not develop, deploy or operate a frontier model in New York without filing a current disclosure statement with the regulator and paying a required assessment. Disclosure statements must be renewed at least every two years and identify the developer's corporate names, New York addresses, certain beneficial owners, and designated points of contact. Large frontier developers will also be required to pay a fee to defray the regulator's operating expenses. The regulator is required to publish a list of developers who file disclosure statements.

### **New effective date**

The amended RAISE Act takes effect on January 1, 2027, which is one year after the TFAIA's effective date of January 1, 2026. This gives frontier developers additional time to prepare for compliance in New York, though companies already subject to the TFAIA will have had a year of experience with substantially similar requirements.

### **Oversight shift and rulemaking authority**

The amended RAISE Act shifts regulatory oversight from the New York Division of Homeland Security and Emergency Services to a new office within the New York Department of Financial Services (DFS Office). The DFS Office will receive disclosures, incident reports and catastrophic risk assessment summaries, with authority to share reports with other governmental entities, including the New York attorney general. The amended law also grants the DFS Office broad rulemaking authority to implement the law, including the power to consider "additional reporting or publication requirements." DFS already oversees a range of cybersecurity matters, including experience with cyber incident reporting for covered entities, and is an active regulator. The TFAIA does not have a similar grant of rulemaking authority to the administering agency.

### **Territorial scope**

The RAISE Act is explicitly limited to frontier models "developed, deployed, or operating in whole or in part in New York state." The TFAIA contains no express territorial limitation. California courts generally apply a presumption against extraterritoriality of state law.

### **Academic exemptions**

The RAISE Act exempts accredited colleges and universities in New York engaged in academic research, as well as the Empire AI consortium, a public-private AI research partnership. The TFAIA contains no comparable exemptions for academic institutions.

## **Key takeaways**

The amended RAISE Act signals a meaningful convergence with California on how to regulate frontier AI, aligning two of the country's largest and most influential state economies in the absence of federal legislation. Companies developing or deploying frontier models may consider several practical implications.

- First, the substantial alignment between the RAISE Act and the TFAIA suggests that a harmonized compliance approach is feasible, but developers must account for key differences. In particular, the RAISE Act includes shorter incident reporting timelines, higher penalty ceilings for repeat violations, and new ownership disclosure requirements, and the DFS Office's rulemaking authority may introduce further distinctions.
- Second, while the RAISE Act primarily targets developers, companies that fine-tune, retrain or otherwise materially modify frontier models should assess whether those activities bring them within the law's scope.
- Finally, with the White House urging AI legislation to preempt state laws, companies should maintain compliance with existing state laws while closely monitoring developments in Washington that could reshape the regulatory landscape.

This content is provided for general informational purposes only, and your access or use of the content does not create an attorney-client relationship between you or your organization and Cooley LLP, Cooley (UK) LLP, or any other affiliated practice or entity (collectively referred to as "Cooley"). By accessing this content, you agree that the information provided does not constitute legal or other professional advice. This content is not a substitute for obtaining legal advice from a qualified attorney licensed in your jurisdiction, and you should not act or refrain from acting based on this content. This content may be changed without notice. It is not guaranteed to be complete, correct or up to date, and it may not reflect the most current legal developments. Prior results do not guarantee a similar outcome. Do not send any confidential information to Cooley, as we do not have any duty to keep any information you provide to us confidential. When advising companies, our attorney-client relationship is with the company, not with any individual. This content may have been generated with the assistance of artificial intelligence (AI) in accordance with our AI Principles, may be considered Attorney Advertising and is subject to our [legal notices](#).

## Key Contacts

<b>Sean Quinn</b> New York	<b>squinn@cooley.com</b> +1 202 728 7075
<b>Michael Egan</b> Washington, DC	<b>megan@cooley.com</b> +1 202 776 2249
<b>Adam Silow</b> New York	<b>asilow@cooley.com</b> +1 212 479 6163

This information is a general description of the law; it is not intended to provide specific legal advice nor is it intended to create an attorney-client relationship with Cooley LLP. Before taking any action on this information you should seek professional counsel.

Copyright © 2023 Cooley LLP, 3175 Hanover Street, Palo Alto, CA 94304; Cooley (UK) LLP, 22 Bishopsgate, London, UK EC2N 4BQ. Permission is granted to make and redistribute, without charge, copies of this entire document provided that such copies are complete and unaltered and identify Cooley LLP as the author. All other rights reserved.