Insights: Alerts

From Transparency to Oversight: New York's RAISE Act Raises the Bar for Frontier AI Developers

January 12, 2026

Written by **Gregory P. Silberman** and **Anthony D. Glosson**

Background

On December 19, 2025, New York Governor Kathy Hochul signed **Assembly Bill A6453-A, the Responsible AI Safety and Education Act (RAISE Act)**, into law.The statute establishes a targeted framework governing the development and deployment of the most advanced "frontier" artificial intelligence models, with a focus on AI safety, transparency, and the prevention and reporting of incidents involving catastrophic harm.The RAISE Act does not include a staged phase-in of obligations. Instead, **the statute becomes fully effective on July 1, 2027**, with certain requirements operating on an ongoing or annual basis once in force.

The RAISE Act imposes affirmative safety, documentation, audit, and incident-reporting obligations on **large developers of frontier AI models**, enforced by the New York Attorney General. It does not regulate AI deployers or ordinary users as a separate class.

In announcing the legislation, Governor Hochul indicated that the enacted bill reflects agreed-upon chapter amendments intended to further align New York's approach with California's **SB 53, theTransparency in Frontier Artificial Intelligence Act (TAFAIA)**. For example, the final statute narrows its focus to the most capable frontier models and concentrates compliance obligations on large developers, rather than regulating AI systems or deployers more broadly. At the same time, the RAISE Act incorporates New York-specific features, including enhanced oversight mechanisms and a shorter safety-incident reporting timeline.

The RAISE Act should be read alongside SB 53, which we previously analyzed in **California Enacts the Transparency in Frontier Artificial Intelligence Act (SB 53) (October 7, 2025)**. Like SB 53, New York's statute is calibrated around catastrophic or critical harm risks associated with advanced general-purpose AI models.

At the federal level, these state AI laws operate against the backdrop of a December 11, 2025, White House Executive Order directing federal agencies to pursue a national policy framework for artificial intelligence and to evaluate whether state AI laws conflict with federal objectives or interstate commerce principles.

As a result, statutes such as the RAISE Act, which impose prescriptive requirements on a narrow but economically significant class of AI developers, could become subjects of federal review or litigation, particularly if they are framed as impeding interstate commerce or conflicting with emerging federal policy. That said, the Executive Order does not preempt state law, and any challenge would require further agency action or judicial proceedings. In the near term, companies should assume that both California and New York requirements remain fully operative, even as federal scrutiny and litigation risk increase.

Why does the RAISE Act matter?

The RAISE Act is significant for three reasons. First, it moves beyond purely transparency-based obligations and **requires large frontier developers to implement, document, and audit safety and security protocols** as a condition of deployment. Second, it establishes a **short, mandatory safety-incident reporting clock (72 hours)**, which is materially more aggressive than California's default 15-day reporting window under SB 53. Third, it signals increasing **convergence among large states around frontier model governance**, even as the precise compliance mechanics differ.

For organizations already preparing for SB 53, the RAISE Act raises the bar on documentation discipline, audit readiness, and incident-response escalation, despite applying to a narrower set of developers.

Who does the RAISE Act apply to?

The RAISE Act applies to **large developers of frontier AI models** that are developed, deployed, or operate in whole or in part in New York State.

A “frontier model” is defined by reference to extremely large-scale training or knowledge distillation, and heightened obligations attach only to developers that meet specified compute-cost thresholds, rather than to AI deployers or users generally. Accredited colleges and universities are excluded to the extent they are engaged in academic research.

While this tiered structure is conceptually similar to California SB 53, the coverage mechanics differ. TAFAIA anchors applicability primarily to an explicit technical threshold (**10^{26} FLOPs**) and a consolidated revenue threshold. By contrast, the RAISE Act combines a high technical benchmark with **dollar-denominated per-model and aggregate training spend thresholds**, using New York-specific statutory definitions.

As a result, organizations should expect **functional alignment rather than identical triggers** between the two regimes. Companies that are in scope or near threshold under SB 53 should plan for a separate jurisdiction-specific assessment for New York.

Importantly, the RAISE Act is **deliberately narrow** and is unlikely to be triggered inadvertently. Unlike SB 53, which is triggered primarily by crossing a high technical capability threshold, RAISE applies only when that capability is paired with very large training investments, effectively limiting its reach to the largest frontier AI developers.

What large developers of frontier AI models need to do?

Covered large developers must adopt, implement, and maintain a written AI safety and security protocol addressing risks associated with the development and deployment of frontier AI models. A redacted version of the protocol must be made public, while an unredacted version must be retained and made available to the State upon request.

Developers must also document testing procedures, results, and safeguards used to evaluate and mitigate the risk of critical harm. These materials must be retained in sufficient detail to support regulatory review and independent audit. This documentation obligation is more prescriptive than California's SB 53 approach, which emphasizes truthful public summaries and internal retention but does not mandate a general audit regime.

The RAISE Act further requires an annual independent third-party audit of compliance with its safety and security requirements, with a redacted report made public and unredacted materials retained for government review. This is a major point of divergence from SB 53, which intentionally avoided mandatory audit requirements. Although the statute as signed includes a requirement for annual independent third-party audits, it remains unclear whether the agreed-upon chapter amendments will retain, modify, or eliminate this obligation, and neither the Governor's office nor the bill sponsors have publicly clarified the status of this provision.

Covered developers must report qualifying safety incidents to the State within 72 hours after learning of the incident or obtaining information sufficient to form a reasonable belief that such an incident has occurred. A safety incident is an event involving a frontier AI model that provides demonstrable evidence of an increased risk of critical harm, including catastrophic loss of life or major economic damage caused or materially enabled by autonomous behavior, compromise or release of model weights, failure of safeguards, or improper use. By contrast, SB 53 generally allows15 days for reporting, with a 24-hour requirement only where there is an imminent risk of death or serious physical injury. In practice, this difference may drive companies to design incident-response workflows to the New York standard.

The RAISE Act prohibits materially false or misleading statements in required disclosures and documentation. Enforcement authority rests with the New York Attorney General, who may seek civil penalties and injunctive relief for violations of the statute. Civil penalties are expected to be capped at up to $1 million for a first violation and up to $3 million for subsequent violations. Although the statute does not create a general private right of action for third parties, whistleblowers may seek judicial relief for retaliation under the statute's employment-protection provisions.

How does RAISE compare to California SB 53?

Both statutes target developers of the most capable general-purpose AI models and focus on preventing catastrophic or critical harm rather than regulating routine or sector-specific AI uses. Each relies on a

combination of public disclosures, internal governance requirements, and enforcement by the state Attorney General, rather than private rights of action.

The principal differences lie in **how prescriptive the obligations become once a developer is in scope**. New York's RAISE Act imposes mandatory safety and security protocols, annual independent audits, and a compressed 72-hour safety-incident reporting timeline. California's SB 53, by contrast, places greater emphasis on standardized transparency reports, framework-based governance, whistleblower protections, and longer default reporting periods.

Unlike California's SB 53, which relies primarily on existing state agencies and Attorney General enforcement, the RAISE Act is expected to establish a **dedicated oversight function within the New York State Department of Financial Services**. Under the agreed-upon chapter amendments, this new office will: (a) assess large frontier developers on an ongoing basis; (b) issue rules and regulations interpreting the statute's requirements; (c) assess fees on covered developers to fund the office's operations; and (d) publish annual reports on AI safety, among other things.

This represents a meaningful departure from California's approach and creates additional compliance infrastructure that developers should plan for. The fee-funded model means that covered developers will bear the direct costs of New York's oversight regime, and the office's rulemaking authority suggests that additional compliance guidance (and potentially additional obligations) will likely emerge after the statute takes effect.

SB 53 is more likely to function as the primary regulatory framework for frontier-model governance, with RAISE imposing more demanding obligations on a narrower set of developers once its higher thresholds are met. California SB 53 is likely to apply to most globally recognized frontier model developers, while New York's RAISE Act is calibrated to reach only the very largest and most capital-intensive developers and, under the Act's spend-based thresholds, plausibly a single-digit number worldwide.

Federal overlay: the White House Executive Order

On December 11, 2025, the White House issued an Executive Order directing federal agencies to advance a national policy framework for artificial intelligence and instructing the Department of Justice to evaluate whether state AI laws conflict with federal objectives or impose undue burdens on interstate commerce.

The Executive Order does not preempt or suspend New York or California law. It does, however, increase the likelihood of litigation, federal review, and policy pressure around state regulation of frontier AI models. Companies should plan for a period in which state compliance obligations remain in force while federal alignment and preemption theories are tested.

How does this impact deployers?

Although the RAISE Act does not directly regulate AI deployers, deployers of highly capable models should expect indirect compliance and contractual impact and take practical steps to manage risk:

- **Confirm deployer versus developer status.** Assess whether fine-tuning, retraining, distillation, or other modifications could constitute “substantial modification” and trigger developer obligations under New York or California law.
- **Review AI vendor contracts.** Expect developers to impose additional use restrictions, incident-notification requirements, and audit-cooperation provisions. Ensure these obligations are operationally feasible and aligned with internal processes.
- **Prepare for incident escalation.** Deployers are often first to detect anomalous behavior or misuse. Confirm the ability to escalate potential safety incidents promptly to developers, particularly given New York's 72-hour reporting timeline.
- **Review internal AI controls.** Evaluate access controls, monitoring, and misuse-prevention measures, which may be scrutinized indirectly as part of a developer's compliance or audit process.

Bottom line

For organizations developing or substantially modifying highly capable AI models, the RAISE Act represents a meaningful escalation beyond California's SB 53, even though it applies to a narrower group of developers. Companies already preparing for SB 53 should expect additional obligations in New York, particularly around audit readiness and accelerated incident reporting.

For in-house legal teams, the lowest-regret strategy is to harmonize internal governance artifacts so that a single control framework can support both SB 53 transparency obligations and RAISE Act safety, audit, and incident-response requirements, while continuing to monitor federal developments that may affect the longer-term regulatory landscape.

Related People

Gregory P. Silberman

t 650.462.5310

gsilberman@ktslaw.com

Anthony D. Glosson

t 202.508.5842

tglosson@ktslaw.com