

January 6, 2026

If you have any questions regarding the matters discussed in this memorandum, please contact the attorneys on the last page or call your regular Skadden contact.

This memorandum is provided by Skadden, Arps, Slate, Meagher & Flom LLP and its affiliates for educational and informational purposes only and is not intended and should not be construed as legal advice. This memorandum is considered advertising under applicable state laws.

One Manhattan West
New York, NY 10001
212.735.3000

New York Enacts AI Transparency Law on Heels of White House Executive Order Aiming to Curb Such State Laws

Executive Summary

- **What's new:** New York's Responsible AI Safety and Education Act (RAISE Act) was signed into law on December 19, 2025, imposing transparency, compliance, safety and reporting requirements on certain developers of large frontier artificial intelligence (AI) models, as well as penalties for violations of these requirements. The RAISE Act will closely align with California's Transparency in Frontier Artificial Intelligence Act (TFAIA).
- **Why it matters:** Governor Kathy Hochul signed the bill just eight days after President Trump issued an executive order announcing a policy to establish a "minimally burdensome" national standard for AI and directing the Department of Justice to challenge state laws deemed inconsistent with that goal.
- **What to do next:** Although the final version of the RAISE Act has not yet been released, reports indicate it will go into effect on January 1, 2027. This gives developers of models covered by the act one year to come into compliance. However, the RAISE Act could face a challenge from the Department of Justice.

On December 19, 2025, New York Governor Kathy Hochul signed into law a comprehensive AI safety and transparency bill, after reaching an agreement with state legislators to pass several amendments to the act in the January 2026 legislative session to more closely align the law with California's Transparency in Frontier Artificial Intelligence Act (TFAIA). See our October 2, 2025, client alert, "Landmark California AI Safety Legislation May Serve as a Model for Other States in the Absence of Federal Standards." Some AI safety advocates have criticized these amendments as considerably watering down the RAISE Act compared to what the New York State legislature had initially passed.

The act may face opposition from the federal government, however, because a recent White House executive order announced support for a "minimally burdensome" federal AI regulatory regime to head off a patchwork of state AI laws, and called on the Justice Department to challenge state laws that are deemed to conflict with that policy.¹

¹ Texas and Colorado previously enacted AI regulations, although those laws were more focused on the use of AI for unlawful discrimination. See our June 23, 2025, client alert "Texas Charts New Path on AI With Landmark Regulation," and our June 24, 2024, client alert, "Colorado's Landmark AI Act: What Companies Need To Know."

New York Enacts AI Transparency Law on Heels of White House Executive Order Aiming to Curb Such State Laws

Overview of the RAISE Act

Scope and Applicability

The act applies to “large developers” of “frontier” AI models that are developed, deployed or operating in New York.

The final version of the act is reported to define “large developers” as those persons with more than \$500 million in revenue, which aligns with California’s TFAIA, enacted in September 2025.² The legislative version of the act had initially sought to instead define large developers more broadly as those who had spent over \$100 million in compute costs to train frontier models, given that most AI developers do not yet have significant revenue. Colleges and universities engaged in academic research are excluded from the definition.

“Frontier models” are defined as those AI models trained using more than 10^{26} computational operations (*i.e.*, integer or floating point operations; also known as “FLOPs”), or models trained from such systems through “knowledge distillation,” provided that the compute costs of such technique exceeds \$5 million.

“Knowledge distillation” refers to a supervised learning technique where a larger artificial intelligence model or its output is used to train a smaller model with similar or equivalent capabilities.

Transparency and Safety Requirements

The main focus of the RAISE Act is to impose duties on large developers of frontier models with respect to safety protocols, testing and transparency.

- Safety and security protocols. Before deploying a frontier model, large developers must develop and disclose to the public, and then maintain, a written safety and security protocol that details, among other points:

- Protections and procedures that reduce the risk of “critical harm”—defined as the death or serious injury of at least 100 people or at least \$1 billion in damages.
- Reasonable administrative, technical and physical cybersecurity protections that reduce the risk of unauthorized access to, or misuse of, the model leading to critical harm.
- Testing procedures to evaluate if the model poses an unreasonable risk of critical harm or could be used to create another frontier model in a manner that would increase the risk of critical harm.

- Publication and retention. A large developer must maintain an unredacted copy of its safety and security protocol for as

long as the frontier model is deployed, plus an additional five years. A version of the protocol that can be redacted to remove trade secrets, personal information and certain other information must be published and submitted to the New York attorney general and Division of Homeland Security and Emergency Services (DHSES).

- **Testing and safeguards.** Developers must document and retain detailed test results and implement safeguards to prevent unreasonable risk of critical harm.
- **Annual review:** Safety protocols must be reviewed and, if necessary, updated annually to account for any changes in the model’s capabilities and industry best practices.

Incident Reporting

The Act requires large developers to report a “safety incident” relating to a frontier model (*e.g.*, unauthorized access, model misuse or critical control failures) to the New York attorney general and DHSES within 72 hours of discovery, including a description and rationale for why it qualifies as a safety incident.

Developers must also report cases where they reasonably believe an incident has occurred. The reporting obligations are significantly stricter than those under California’s TFAIA, which provides developers with a 15-day reporting window, and only covers cases where there is definitive knowledge of an incident. In her memo approving the Act, Governor Hochul noted that large developers may describe limitations on their knowledge of any safety incidents involving models that have been modified by unaffiliated third parties.

Enforcement and Penalties

The New York attorney general has exclusive authority to enforce the RAISE Act. Governor Hochul’s office reported that the attorney general will be authorized to bring civil actions against large developers for failing to submit the required reporting or making false statements, with penalties up to \$1 million for a first violation and up to \$3 million for subsequent violations. These penalties are substantially lower than the \$10 and \$30 million penalties for first and subsequent violations that the New York legislature had sought, and brings the RAISE Act more in line with California, which has a penalty of up to \$1 million per violation, with no increase for subsequent violations. Courts may also issue injunctive or declaratory relief for violations.

New Oversight Office

The final version of the RAISE Act is reported to create an oversight office within the Department of Financial Services to assess fees on large developers, address enforcement, issue rules and regulations, and publish an annual report on AI safety.

² The final amendments to the RAISE Act have not yet been released, so the following summary is based on reports on what Governor Hochul negotiated with state legislators.

New York Enacts AI Transparency Law on Heels of White House Executive Order Aiming to Curb Such State Laws

Potential Federal Challenge: The Trump Administration's Executive Order

The RAISE Act may face federal opposition following a [December 11, 2025, executive order](#), which seeks to ensure there is a unified, minimally burdensome national AI regulatory framework. See our December 15, 2025, client alert "[White House Launches National Framework Seeking To Preempt State AI Regulation](#)."

We expect that the Department of Commerce will — as empowered by the executive order — consider the RAISE Act “burdensome” and in conflict with the order’s stated goals, asserting that the New York and California laws represent first steps in the “patchwork” of state laws that the executive order says could stifle AI innovation.

Whether the U.S. attorney general then brings a lawsuit challenging the RAISE Act (as well as California’s TFAIA) remains to be seen, but we expect that is likely. Under the executive order, New York could also be threatened with a loss of certain federal funding if the Department of Commerce identifies the RAISE Act as onerous.

Looking Ahead

Although we believe a federal challenge to the New York law is likely, companies that meet the definition of large developers of frontier models should review the requirements of the RAISE Act, and start to put into place policies and procedures to meet its reporting, safety and security protocol requirements.

Contacts

Stuart D. Levi

Partner / New York
212.735.2750
stuart.levi@skadden.com

William E. Ridgway

Partner / Chicago
312.407.0449
william.ridgway@skadden.com

Maria Cruz Melendez

Partner / New York
212.735.2320
maria.cruzmelendez@skadden.com

Mana Ghaemmaghami

Associate / New York
212.735.2594
mana.ghaemmaghami@skadden.com

Brittany E. Libson

Associate / New York
212.735.2548
brittany.libson@skadden.com