wiley

ALERT

New York Finalizes RAISE Act for Frontier AI Models; Law Takes Effect January 1, 2027

April 3, 2026

On March 27, 2026, New York Governor Kathy Hochul signed into law a chapter amendment that is the final version of the Responsible AI Safety and Education Act (RAISE Act), New York's new law regulating frontier AI models. The RAISE Act establishes a range of detailed obligations for developers of frontier models, and grants rulemaking authority to a new office within the New York Department of Financial Services (DFS). While the chapter amendment brings the final version of the RAISE Act more in line with the California Transparency in Frontier Artificial Intelligence Act (TFAIA), including aligning with some key definitions, other provisions of the RAISE Act are a marked departure from TFAIA, including the requirement to disclose incidents within 72 hours (TFAIA has a 15-day timeline). Accordingly, the New York law establishes another unique state standard for AI model transparency and safety, which is in tension with the Trump Administration's efforts to promote comprehensive federal AI legislation in lieu of a fragmented patchwork of state laws.

Below, we provide a high-level summary of the RAISE Act, which will take effect on January 1, 2027.

Key Background

The New York Governor originally signed the RAISE Act into law on December 19, 2025, following the state's legislature passing the bill in June 2025. The Governor negotiated with lawmakers to secure a chapter amendment, which was introduced on January 6, 2026, passed the second chamber of the legislature on March 11, 2026, and was finally signed into the law by the Governor on March 27, 2026.

Authors

Duane C. Pozza
Partner
202.719.4533
dpozza@wiley.law

Kathleen E. Scott
Partner
202.719.7577
kscott@wiley.law

Sydney M. White
Special Counsel
202.719.3425
swhite@wiley.law

Practice Areas

Artificial Intelligence (AI)
Emerging Technologies
Privacy, Cyber & Data Governance
State Privacy Laws
State Regulation
Telecom, Media & Technology
Trump Administration Resource Center

Purpose and Applicability

The purpose of the RAISE Act is to shape AI safety protocols through transparency and safety requirements applicable to frontier developers and frontier models. Specifically, the RAISE Act applies to frontier models "developed, deployed, or operating" in New York and requires large frontier developers to establish and publish safety protocols.

"Frontier developer" is defined as a person who has trained a "frontier model," defined as a foundation model trained using "computing power greater than 10°26 integer or floating-point operations." The Act defines "foundation model" as an AI model that is:

- Trained on a broad data set
- Designed for generality of output
- Adaptable to a range of tasks

"Large frontier developer" is defined as a frontier developer having annual revenue exceeding $500 million. Among the changes made by the chapter amendment are the addition of the computing power and revenue-based thresholds to distinguish the law's applicability to AI models, frontier developers, and large frontier developers. These are the same thresholds as TFAIA.

Similar to TFAIA, the RAISE Act focuses on preventing a frontier developer's development of a model contributing to "catastrophic risk" or a "critical safety incident." Catastrophic risk is defined as "a foreseeable and material risk that a frontier developer's development, storage, use or deployment of a frontier model will materially contribute to the death of, or serious injury to, more than fifty people or more than one billion dollars in damage to or loss of, property arising from a single incident involving a frontier model" that does any of the following:

- Provides expert-level assistance in creating "a chemical, biological, radiological or nuclear weapon."
- Engages in conduct with no human oversight that is a cyberattack or if committed by a human would constitute a murder, assault, or theft.
- Evades control of its developer or user "in a manner that demonstrates materially increased catastrophic risk."

The definition of catastrophic risk excludes foreseeable and material risk from the release of otherwise publicly available information, lawful activity of the federal government, or harm caused by a frontier model and other software "if the model did not materially contribute to the harm."

A "critical safety incident" is defined to include:

- Unauthorized access to, or modification or exfiltration of, the model weights "that results in death or bodily injury."

- "[H]arm resulting from materialization of a catastrophic risk."
- "[L]oss of control of the model causing death or bodily injury."
- Use of deceptive techniques by the model against its developer "in a manner that demonstrates materially increased catastrophic risk."

Key Requirements

Frontier AI Framework. The RAISE Act uses the same language as TFAIA that requires large frontier developers to publish on their website a frontier AI Framework that describes how the developer:

- Incorporates national standards;
- Establishes thresholds and mitigations for catastrophic risk;
- Reviews and assesses mitigations for catastrophic risk;
- Uses third parties to also conduct such assessments ;
- Updates the frontier AI Framework;
- Implements cybersecurity practices to secure unreleased model weights;
- Institutes internal governance practices supporting implementation of the Framework;
- Identifies and responds to critical safety incidents; and
- Assesses and manages catastrophic risk from internal use of frontier models.

Frameworks must be updated annually (or sooner following a material modification).

Transparency Reports. The RAISE Act also borrows transparency report language from TFAIA. Under the Act, before deploying a new frontier model or a modified version of an existing model, the frontier developer is required to publish a transparency report on its website. The report is required to include:

- Contact information for the developer
- Release date of the model
- Languages supported and modalities of output of the model
- Intended use of the model
- Restrictions or conditions on use of the model
- Summaries of assessment of catastrophic risk, extent third-party evaluators were involved in the assessment, and other steps to satisfy the requirement of the frontier AI Framework

A frontier developer that publishes this information in a larger document, such as "a system card or model card, shall be deemed in compliance." Additionally, a frontier developer is permitted to redact information from the transparency report to protect trade secrets, cybersecurity, public safety, or national security, or to comply with state or federal law, subject to justification and recordkeeping requirements.

Reporting Requirements for Catastrophic Risks and Critical Safety Incidents. The Act establishes a new DFS oversight office ("Office"), which will receive various required reports. Specifically:

- Large frontier developers are required to transmit a summary to the Office of "any assessment of catastrophic risk resulting from the internal use of its frontier models every three months" and provide written updates. The Office is required to establish a mechanism to allow large frontier developers to submit summaries confidentially and "shall take all reasonable precautions to limit access to any reports" including to limit access to only authorized personnel and prevent unauthorized access.
- Frontier developers are required to report to the Office within "seventy-two hours from a determination that a critical safety incident involving one of its frontier models has occurred or within seventy-two hours of the frontier developer learning facts sufficient to establish a reasonable belief that a critical safety incident has occurred." If the frontier developer "discovers" additional information about the incident after filing the report, it has the option of filing an amended report.
- If a frontier developer "discovers that a critical safety incident poses an imminent risk of death or serious physical injury, disclosure is required to be made to a "law enforcement or public safety agency with jurisdiction, that is appropriate" within 24 hours. A developer disclosing to law enforcement or a public safety agency is still required to report to the Office within 72 hours.

The Office is authorized to designate federal regulations, laws, or guidance documents as methods of compliance with certain reporting obligations. A developer that plans to comply with the federal regulation, law, or guidance is required to notify the Office.

The Office may, at its discretion and subject to factors outlined in the statute, share the reports or summaries of assessments of catastrophic risk from internal use of frontier models with other governmental agencies including the Office of the Attorney General of New York.

Beginning in January 2028, the Office is required to produce a report for the Governor and Legislature using "anonymized and aggregated" critical safety incident information and include any recommendations on updates to the RAISE Act. The Office is prohibited from including information in the report "that would compromise the trade secrets or cybersecurity of the frontier developer," public safety or national security, "or that would be prohibited by federal or state law."

Disclosure Statement. A large frontier developer is required to file a disclosure statement with the Office, to be renewed every two years or when there is a change in ownership. The disclosure statement is required to:

- Identify the large frontier developer and all names under which it conducts business;
- Include the addresses of the large frontier developer's principal place of business and each New York office; and
- Identify all persons who beneficially own a 5% or greater ownership in the large frontier developer and all persons who beneficially owned a 5% or greater ownership in the past five years, if it is privately held.

Rulemaking Authority

The new Office is granted authority to issue regulations to implement the law, including potentially "additional reporting or publication requirements to facilitate safety and transparency … including … post-critical safety incident information."

Enforcement

The New York Attorney General will enforce the law by imposing civil penalties up to $1 million for a first violation and up to $3 million for subsequent violations. Violations include when a large frontier developer fails to file a required document (including disclosure statement), makes a false or misleading statement in violation of the requirements related to a developer's frontier AI Framework, fails to report an incident, or fails to comply with its own frontier AI Framework. The Act is explicit that it does not establish a private right of action.

Additionally, the Office is authorized to impose civil penalties of $1,000 per day after notice and a hearing for a failure to file a disclosure statement or correct false information.

***

Wiley's Artificial Intelligence Practice counsels clients on AI compliance, risk management, and regulatory and policy approaches, and we engage with key government stakeholders in this quickly moving area. Please reach out to a member of our team with any questions.