CLIENT ALERT | July 15, 2026

Illinois Joins Growing State-Level Effort to Regulate Frontier AI With New Safety Measures Act

Illinois' SB 315 is the first law in the US to mandate third-party audits for large AI developers.

Key Points:

- Illinois Senate Bill 315 largely tracks similar frontier AI laws in California and New York, requiring frontier model developers to adopt and publish an AI framework, issue transparency reports, report critical safety incidents, and register with state regulators.
- In a first, the law mandates that frontier model developers that earned more than $500 million in revenue in the preceding year must retain an independent third party to perform an annual compliance audit, a significant new obligation with no analog in California or New York law.
- Violations are enforceable exclusively by the Illinois Attorney General, with civil penalties of up to $1 million for a first violation and up to $3 million for subsequent violations.
- The law will take effect on January 1, 2027, but certain obligations (including the third-party audit requirement) will not phase in until January 1, 2028.

On July 7, 2026, Illinois Governor JB Pritzker signed Senate Bill 315, known as the Artificial Intelligence Safety Measures Act (the Act), into law, making Illinois the latest in a growing group of states to enact legislation governing frontier artificial intelligence.

The Act is modeled substantially after California's Transparency in Frontier AI Act (the TFAIA) and New York's Responsible AI Safety and Education Act (the RAISE Act), borrowing many of its core obligations from those statutes. Like its California and New York predecessors, the Act requires large frontier developers to adopt a frontier AI framework, publish transparency reports, report critical safety incidents to state regulators, and file disclosure statements with the state.

But the Act also adds a notable new wrinkle: It is the first law in the US to require large frontier developers to undergo annual independent third-party compliance audits. As such, the Act represents another significant measure that frontier developers must account for when building out their compliance programs.

This Client Alert analyzes the key requirements of the Act, compares its provisions to those of the TFAIA and the RAISE Act, and identifies the practical compliance considerations that frontier developers should be evaluating now.

Illinois Senate Bill 315 – Artificial Intelligence Safety Measures Act

Scope and Application

The Act broadly applies to developers of "frontier models," defined as any foundation model trained using a quantity of computing power greater than 10^{26} integer or floating-point operations (FLOPs), inclusive of the original training run and any subsequent fine-tuning, reinforcement learning, or other material modifications. However, a number of the Act's more onerous requirements apply only to "large frontier developers" — that is, developers of frontier models that, together with their affiliates, collectively had annual gross revenues in excess of $500 million in the preceding calendar year.

These threshold definitions are materially identical to those in the TFAIA and the RAISE Act, meaning that any developer currently in scope under California or New York law will almost certainly be in scope under the Illinois Act as well, provided it develops, deploys, or operates a frontier model in whole or in part within Illinois.

Key Risk Concept: Catastrophic Risk

The Act's obligations are organized around the concept of "catastrophic risk," which is defined consistently with the TFAIA and the RAISE Act as a foreseeable and material risk that a developer's development, storage, use, or deployment of a frontier model will materially contribute to the death of, or serious injury to, more than 50 people, or result in more than $1 billion in damage to, or loss of, property arising from a single incident involving a frontier model that: (i) provides expert-level assistance in the creation or release of a chemical, biological, radiological, or nuclear weapon; (ii) engages in conduct with no meaningful human oversight, intervention, or supervision that is either a cyberattack or, if committed by a human, would constitute murder, assault, extortion, or theft; or (iii) evades the control of its frontier developer or user.

The Act includes the same carveouts found in the TFAIA and the RAISE Act: Catastrophic risk does not include risk from information that is already publicly accessible in substantially similar form, lawful activity of the federal government, or harm caused by a frontier model in combination with other software where the model did not materially contribute to the harm.

Independent Third-Party Audits

The Act's most significant departure from the TFAIA and the RAISE Act is its mandate that large frontier developers retain an independent third party to perform an annual compliance audit. Beginning on January 1, 2028 (or 90 days after a developer first qualifies as a large frontier developer, whichever is later), each large frontier developer must engage an independent auditor to assess whether the developer has substantially complied with the Act's requirements.

While both the TFAIA and the RAISE Act require large frontier developers to describe how they "use third parties to assess the potential for catastrophic risks and the effectiveness of mitigations" as part of their frontier AI framework, this is a self-described policy commitment rather than an affirmative audit mandate.

Auditor independence and qualifications. The third-party auditor must possess "demonstrated competence" to perform the audit, including experience employing or contracting with individuals who have technical expertise in the safety of frontier models. The Act prohibits a developer from retaining an auditor if either the developer or the auditor has a financial interest in the other party. The developer may compensate the auditor for its services but may not condition any payment on the results of the audit.

Scope and methodology of the audit. The third-party auditor must be granted access to all materials reasonably necessary to perform the audit, including unredacted versions of all materials published under the Act. To protect trade secrets, confidential business information, and cybersecurity interests, the developer may impose security protocols on the auditor, such as restrictions on note-taking, copying, or removing materials; requirements for on-premise review; and confidentiality obligations.

The auditor's report must include: (i) a description of whether the developer has substantially complied with the Act; (ii) a description of any material deviations, an explanation and rationale for each deviation, and recommendations for improving compliance; (iii) a detailed assessment of the developer's internal controls, including its designation and empowerment of senior personnel responsible for implementation; (iv) a list of the personnel involved in the audit; (v) the auditor's procedures for managing conflicts of interest and any conflicts of interest of audit personnel; (vi) the methodology of the audit and the nature of the information reviewed; and (vii) the signature of the lead auditor certifying the results.

Publication and retention requirements. The developer must retain an unredacted copy of the audit report for as long as a frontier model is deployed, plus five years. No later than 30 days after receiving the report, the developer must publish on its website a high-level summary of the audit findings and a copy of the report with appropriate redactions, and the developer must transmit a redacted copy of the report to both the Illinois Emergency Management Agency and Office of Homeland Security (the Agency) as well as the Illinois Attorney General.

Other Obligations for Frontier Developers

Most of the Act's remaining obligations overlap with requirements that already arise under the TFAIA and/or the RAISE Act, including:

Frontier AI framework. Beginning January 1, 2028, large frontier developers must publish and comply with a frontier AI framework that describes how the developer incorporates national, international, and industry consensus best practices; defines and assesses catastrophic risk thresholds; applies and reviews mitigations before deployment; uses third-party risk assessments; maintains cybersecurity for

unreleased model weights; identifies and responds to critical safety incidents; and institutes internal governance.

The framework must be reviewed and, as appropriate, updated annually. Material modifications to the framework must be published with justification within 30 days.

Transparency reporting. Before deploying a new or substantially modified frontier model, all frontier developers must publish a transparency report including the model's release date, supported languages, output modalities, intended uses, usage restrictions, and developer contact information.

Large frontier developers must also disclose summaries of catastrophic risk assessments, the results of those assessments, third-party evaluator involvement, and other framework compliance steps.

Quarterly risk reporting. Large frontier developers must submit summaries of catastrophic risk assessments from internal model use to the Agency once per quarter, or on another schedule agreed upon with the Agency.

Critical safety incident reporting. All frontier developers must report critical safety incidents[1] to the Agency and Illinois Attorney General within 72 hours. Incidents posing imminent risk of death or serious injury must be disclosed within 24 hours to an appropriate authority, including law enforcement or public safety agencies. This 72-hour window is consistent with the RAISE Act but significantly shorter than the TFAIA's 15-day window.

Disclosure statement and fees. Beginning January 1, 2027, large frontier developers must file an annual disclosure statement with the Agency that includes the developer's identity, principal place of business, Illinois addresses, ownership information, and designated contacts. The Agency will also assess pro rata fees on large frontier developers to cover administration costs.

Whistleblower protections. Developers may not retaliate against covered employees[2] who disclose to the Agency or other authorities either a violation of the Act or another specific and substantial danger to public safety arising from catastrophic risk.

Large frontier developers must provide an anonymous internal reporting process and monthly updates to the discloser regarding the status of the developer's internal investigation.

Enforcement and Penalties

Violations of the Act are enforceable exclusively by the Illinois Attorney General through civil actions. The Act does not create a private right of action. Civil penalties may not exceed $1 million for a first violation and $3 million for subsequent violations, with the amount dependent upon the severity of the violation.

Comparison to the TFAIA and the RAISE Act

This chart summarizes the key similarities and differences among Illinois SB 315, the California TFAIA, and the New York RAISE Act:

Provision	Illinois SB 315	California TFAIA	New York RAISE Act
Effective date	Jan. 1, 2027 (disclosure); Jan. 1, 2028 (framework/audit)	Jan. 1, 2026	Jan. 1, 2027
Frontier model threshold	$>10^{26}$ FLOPs	$>10^{26}$ FLOPs	$>10^{26}$ FLOPs
Large frontier developer revenue threshold	>$500M annual gross revenue	>$500M annual gross revenue	>$500M annual gross revenue
Frontier AI framework	Required (large frontier developers)	Required (large frontier developers)	Required (large frontier developers)
Transparency reports	Required (all frontier developers); machine-readable format for large developers	Required (all frontier developers)	Required (all frontier developers)
Critical safety incident reporting	72 hours (24 hours if imminent risk)	15 days (24 hours if imminent risk)	72 hours (24 hours if imminent risk)
Independent third-party audit	Mandatory annual audit with detailed reporting	Not required	Not required
Disclosure/registration	Annual renewal	Not required	Renewal every two years
Penalties	Up to $1 million for first violation; up to $3 million for subsequent violations	Up to $1 million	Up to $1 million for first violation; up to $3 million for subsequent violations

Strategic Implications for Frontier Developers

The good news for frontier developers that have already built compliance infrastructure to satisfy the TFAIA and the RAISE Act is that Illinois' AI Safety Measures Act will not require any significant new outlay of resources for more than a year. As described above, the framework, transparency reporting, and incident reporting obligations are sufficiently similar across all three states that existing compliance programs should accommodate Illinois' requirements with only modest adjustments.

While the Act's third-party audit requirement represents a new addition to frontier AI law in the US, the obligation does not phase in until 2028, giving covered developers more than a year to plan their approach. Developers can make good use of that time by doing the following:

Identify and Engage Qualified Auditors Early

The Act requires that the third-party auditor possess "demonstrated competence" to perform the audit, including experience employing or contracting with individuals who have "technical expertise in the safety of frontier models." The auditor must also be financially independent of the developer. These requirements will likely constrain the pool of qualified candidates significantly in the short term. Traditional financial auditors and compliance consultants may lack the requisite AI safety expertise, while AI safety researchers and technical consultants may lack formal audit experience or the institutional independence the Act demands.

Large frontier developers should begin identifying potential audit firms now, evaluating candidates against the statute's competence, independence, and conflict-of-interest requirements. Developers may also wish to consider whether firms with experience in adjacent domains (such as cybersecurity auditing, SOC 2 compliance, or algorithmic auditing under other regulatory frameworks) could be positioned to develop the necessary frontier-model expertise in time for the 2028 compliance date.

Establish Audit Access and Security Protocols

The Act grants auditors access to all materials "reasonably necessary" to comply with their obligations, including unredacted versions of all materials published under the Act. At the same time, the Act expressly permits developers to impose security protocols on the auditor to protect trade secrets, confidential business information, cybersecurity, national security, and public safety. This includes restrictions on note-taking, copying, retaining, or removing materials; requirements for on-premise review; and confidentiality agreements.

Developers should begin considering which of these measures are necessary given their business model and, where necessary, start building out protocols and incorporating them into auditor engagement agreements well in advance. In some cases, developers should be able to build on security frameworks that already exist within the organization — for example, protocols governing external access to proprietary model weights, red-team evaluation procedures, or cybersecurity assessment workflows can be relatively easily adapted or extended to auditors.

Prepare Internal Controls Documentation

The audit report must include a detailed assessment of the developer's internal controls, including its "designation and empowerment of senior personnel responsible for such implementation by the large frontier developer, its employees, and its contractors." In practice, this means auditors will be evaluating whether the developer has established clear lines of accountability for AI safety compliance and whether the personnel charged with implementation have the authority and resources to carry out their

responsibilities effectively. Such internal controls could include, among other things, formal designation of a chief AI safety officer or equivalent role; board- or committee-level oversight of catastrophic risk management; documented escalation procedures for critical safety incidents; periodic internal reviews of compliance with the frontier AI framework; and change-management protocols for updates to the frontier AI framework.

Developers that have already implemented governance structures in response to the TFAIA or the RAISE Act should map those structures against the Act's requirements to identify any gaps. Developers that have not yet formalized their internal controls should treat the lead time before the January 1, 2028, effective date as an opportunity to design a governance infrastructure that will withstand an auditor's scrutiny.

Monitor the Federal Interoperability Mechanism

The Act contains a federal interoperability provision that allows a frontier developer to be deemed in compliance with Illinois law by instead complying with a federal law, regulation, or guidance document that the Agency has designated as imposing standards substantially equivalent to those required by the Act. While the Trump administration has been hesitant to impose prescriptive AI regulation, it has shown an increasing willingness to engage on frontier AI governance. In June, President Trump signed an executive order addressing frontier AI safety and implementing a voluntary governmental review process for new frontier models. If the administration goes further by adopting a federal framework that includes audit, incident reporting, and risk-assessment requirements for frontier developers, those requirements could potentially satisfy the Act's interoperability criteria.

Track Your Company's Growth

Finally, even developers of AI systems that are below the Act's model-scale or revenue thresholds should monitor their growth trajectories. Like the TFAIA and the RAISE Act, the obligations under Illinois' law are triggered once a company or its model crosses defined thresholds, and developers that fail to comply in a timely manner with the Act's requirements once those thresholds are met could face substantial fines. Developers approaching frontier scale should begin building compliance infrastructure well in advance.

Contacts

Michael H. Rubin
michael.rubin@lw.com
+1.415.395.8154
San Francisco

Andrew Gass
andrew.gass@lw.com
+1.415.395.8806
San Francisco

Ghaith Mahmood
ghaith.mahmood@lw.com
+1.213.891.8375
Los Angeles

Sy Damle
sy.damle@lw.com
+1.202.637.3332
+1.212.906.1659
Washington, D.C. / New York

Fiona Maclean
fiona.maclean@lw.com
+44.20.7710.1822
London

Endnotes

[1] A "critical safety incident" includes unauthorized access to, modification of, or exfiltration of model weights resulting in death or injury; harm from a materialized catastrophic risk; loss of model control causing death or injury; or a model using deceptive techniques to subvert developer controls, demonstrating materially increased catastrophic risk.

[2] A "covered employee" is any employee responsible for assessing, managing, or addressing the risk of critical safety incidents.