Cooley

# Illinois Mandates Independent AI Audits: What Developers Should Know

July 7, 2026

## I. Illinois SB 315 signals next phase of AI regulation – from transparency to verification

Over the last several years, lawmakers in the United States and around the world have increasingly focused on regulating AI systems through transparency, documentation and internal risk management requirements.

Recent frameworks, such as California's Transparency in Frontier Artificial Intelligence Act (TFAIA), New York's amended Responsible AI Safety and Education (RAISE) Act and portions of the European Union's AI Act, generally require developers to assess and disclose how they identify, evaluate and manage AI-related risks. Common obligations include transparency reports, system/model cards, risk assessments, governance frameworks and incident reporting.

Illinois' recently enacted Artificial Intelligence Safety Measures Act (AISMA) builds on these existing frameworks by introducing a significant new requirement: independent verification. Rather than relying solely on developer-created documentation and self-reported compliance measures, AISMA requires covered large frontier model developers to undergo audits by independent third parties. This move reflects a broader shift in regulatory efforts from requiring companies to **document** how they manage AI risk to requiring them to **demonstrate** that those processes are actually operating as intended. This is a significant change from self-reported compliance, mirroring a trend in third-party audit requirements in content regimes like the EU's Digital Services Act and South Carolina's Age-Appropriate Code Design.

AISMA may represent the next phase of AI regulation – one focused not only on disclosure, but also on third-party verification.

## II. Key elements of the law

**Who does Illinois' law apply to?**

Developers responsible for the most advanced foundation models.

The law regulates "frontier models" (models trained using more than $10^{26}$ floating-point or integer operations) and imposes obligations on frontier developers broadly. However, its most significant requirements fall on "large frontier developers" – those with annual gross revenues exceeding $500 million.

**When does it go into effect?**

AISMA's effective date is January 1, 2027. However, certain provisions are phased in the following year, including the Frontier AI Framework (January 1, 2028) and the independent audit requirement (January 1, 2028, or 90 days after first qualifying as a large frontier developer, whichever is later).

**What does the law require?**

**Mandatory framework:** Large frontier developers must establish, implement, comply with and publicly publish a Frontier AI Framework that:

- Describes how the developer incorporates national and international standards and industry best practices.
- Defines and assesses catastrophic risk thresholds.
- Applies mitigation measures to address potential catastrophic risks.
- Reviews risk assessments and mitigations before deployment and significant internal use.
- Uses third-party evaluators.
- Updates and maintains its framework over time.
- Protects unreleased model weights through cybersecurity controls.
- Identifies and responds to critical safety incidents.
- Implements internal governance processes.
- Assesses catastrophic risks arising from internal use of frontier models, including risks associated with models circumventing oversight mechanisms.

**Transparency report:** Before deploying a new frontier model, or a substantially modified version of an existing model, a frontier developer must publish, among other things, the model's release date, supported languages, output modalities, intended uses, applicable use restrictions and contact information for the developer.

Large frontier developers must also disclose summaries of catastrophic risk assessments, assessment results, involvement of third-party evaluators and other measures taken to comply with their Frontier AI Framework.

Developers may satisfy many of these disclosure requirements through existing system cards or model cards.

**Ongoing reporting to regulators:** Large frontier developers must provide the Illinois Emergency Management Agency and Office of Homeland Security (Agency) every three months (or on another reasonable schedule) with summaries of assessments regarding catastrophic risks arising from internal use of frontier models.

In addition, frontier developers must report any "critical safety incident" to the Agency and the Illinois attorney general within 72 hours after learning facts sufficient to establish a reasonable belief that such an incident has occurred, or within 24 hours to an appropriate authority where the incident "poses an imminent risk of death or serious physical injury."

**Independent audits:** Developers must annually retain an independent third party to audit compliance with AISMA, which:

- Evaluates whether the developer has substantially complied with AISMA.
- Assesses the developer's internal controls and governance processes.
- Identifies any material deviations from statutory requirements.
- Provides recommendations for improvement where appropriate.

Auditors must possess appropriate expertise, operate free from specified conflicts of interest and conduct their reviews in accordance with generally accepted auditing standards and best practices.

Within 30 days of receiving the report, the developer must publish a high-level summary of the audit findings, publish a redacted version of the audit report and provide the audit report to the Agency and the Illinois attorney general.

# III. Illinois compared with California and New York

**What do all three state laws have in common?**

Illinois joins a growing number of states seeking to regulate the development and deployment of frontier AI models. Before Illinois enacted AISMA, both California and New York had enacted regulatory frameworks for frontier model developers. Although the details differ, California's TFAIA and New York's RAISE Act impose a common set of obligations, including AI framework requirements, transparency and reporting obligations, catastrophic risk assessments, critical safety incident reporting and enforcement by the state attorney general. Together, these laws reflect a broader trend toward requiring frontier model developers to document and disclose how they identify, assess and manage catastrophic AI risks.

Like California's law, AISMA includes whistleblower protections and internal reporting mechanisms intended to surface AI safety concerns before they develop into critical incidents.

Like New York's law, AISMA requires large frontier developers to make registration-style disclosures, identify responsible contacts and pay assessments supporting administration of the regulatory regime.

**What ultimately sets Illinois' law apart?**

Against this shared backdrop, what distinguishes Illinois from both states is its audit requirement. Neither California's TFAIA nor New York's RAISE Act require covered developers to undergo independent audits of their compliance programs. Illinois moves beyond transparency toward independent auditing. The statute reflects the view that AI governance programs should not only be self-reported, but also undergo external verification.

# IV. AI audits in the global context

Although Illinois is the first US state to require annual independent audits of frontier model developers, the concept of independent review and ongoing audits is not unique to AISMA. Similar themes are increasingly appearing in AI regulatory frameworks around the world. For example:

- **EU AI Act**: Providers of certain high-risk AI systems must satisfy conformity assessment requirements and maintain technical documentation, risk management procedures and post-market monitoring processes – reflecting a similar push for documented and verifiable compliance measures.
- **EU Digital Services Act (DSA)**: Very large online platforms and search engines must conduct systemic risk assessments and undergo independent audits. Though not AI-specific, the DSA reflects the same regulatory shift toward requiring organizations to demonstrate governance effectiveness through independent, external review.
- **Vietnam's AI law**: Vietnam's AI law takes a risk-based framework tied to particular AI systems based on their risk classification. High-risk AI systems must undergo conformity assessments, audits and independent testing before deployment and following significant changes. Medium- and low-risk systems are subject to key obligations, such as transparency and incident reporting. Both the Illinois and Vietnam frameworks reflect a similar underlying interest in independent review of AI systems.

Taken together, AISMA's audit requirement may be less of an outlier than it initially appears. Instead, it represents a growing trend toward companies not only maintaining governance programs, but also programmatically demonstrating that those programs are operating effectively.

# V. How AI audits differ from audits clients already know – and why that matters

Most companies are already familiar with financial, cybersecurity and privacy audits. While there are some common elements, AI audits are different in several important ways.

Unlike traditional compliance exercises, AI audits require organizations to evaluate and substantiate complex judgments regarding:

- Evaluating and substantiating judgments about model safety and catastrophic risks.
- Assessing internal governance processes and deployment decisions.
- Demonstrating and verifying the actual effectiveness of risk mitigation measures.

This expanded evaluation scope creates both strategic advantages and potential legal vulnerabilities for frontier developers.

Opportunities include:

- Independent audits can help organizations concretely demonstrate compliance with evolving AI governance and regulatory

obligations – mitigating the risk of regulatory inquiries.
- External verification increases confidence in model safety and security among regulators, customers, investors and the public.
- Rigorous audits can identify weaknesses in internal risk management programs before they escalate into enforcement or litigation issues.

Risks include:

- Audit reports may inadvertently become roadmaps for regulators by exposing governance deficiencies, unresolved risks or gaps between documented policies and actual practices.
- Although these audits can improve governance, the findings may also become relevant evidence in regulatory investigations, enforcement actions or litigation.

# VI. Practical steps companies should consider now

Although AISMA's audit requirement does not take effect until January 1, 2028, or 90 days after an organization first qualifies as a large frontier developer, companies should begin their preparations well before the first audit cycle arrives.

Frontier labs looking to prepare for the audit should consider:

- Assessing whether current or anticipated AI development activities could trigger audit requirements.
- Building and operationalizing audit-ready governance structures, which can take a long time to design and launch.
- Reviewing documentation practices with a view to maintaining consistent model evaluations, safety testing, risk assessments and incident response records.
- Reviewing the role of legal privilege in audit processes.

While independent third-party audits represent a new frontier for AI regulation, navigating first-of-their-kind statutory audit frameworks is not new territory for Cooley. Combining market-leading AI legal acumen with proven, practical experience guiding clients through novel external audit regimes around the globe, Cooley serves as a trusted strategic advisor to technology companies on their most complex digital regulation challenges.

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


| | |
|---|---|
| Janet H. Kim<br>Washington, DC | janetkim@cooley.com<br>+1 202 728 7060 |

| | |
|---|---|
| Sean Quinn<br>New York | squinn@cooley.com<br>+1 202 728 7075 |
| Tristan Lockwood<br>London | tlockwood@cooley.com<br>+44 20 7556 4115 |
| Adam Silow<br>New York | asilow@cooley.com<br>+1 212 479 6163 |
| Isamar Vaquero<br>Washington, DC | ivaquero@cooley.com<br>+1 202 776 2217 |


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.