

Illinois SB 315: A State Strategy for Enduring National AI Safety Standards

June 10, 2026

By [Scott A. Meyers](#) and [Julian Dayal](#)

I. The State-Level Path to National Standards

Illinois is on the cusp of enacting what may prove to be one of the most consequential pieces of artificial intelligence legislation in the country. Not because Illinois SB 315 is sweeping -- it's not. But because it is intentionally modest. Under SB 315, large "frontier" AI developers will be required to (1) disclose their safety policies and procedures regarding certain statutorily defined "catastrophic risks," such as the risk AI could be used to help build a WMD, and (2) submit to annual third-party audits to verify compliance with those policies and procedures. SB 315 is thus narrow in scope. And it concerns risks that are difficult to argue should go unregulated. The method of regulation, meanwhile -- disclosures plus audits -- is familiar, relatively constrained, and draws on the AI expertise not of the State (which may be lacking) but of private third parties. SB 315 was designed, in other words, to avoid the constitutional and political landmines that attend state regulation of a new technology that proponents claim is set to revolutionize the world.

According to Illinois lawmakers, SB 315 is a necessary stopgap safety regime in the absence of federal laws governing generative AI. Yet if no federal laws materialize -- and none appear imminent -- SB 315 may become the *de facto* national baseline for AI regulation. Conversely, if in response

Related People

[Julian Dayal](#)
[Scott A. Meyers](#)

Akerman Intelligence

Akerman Intelligence is Akerman's platform for AI strategy, research, and thought leadership. Through original analysis, client collaboration, and strategic partnerships, we help organizations navigate the legal, operational, and governance challenges of artificial intelligence.

to anticipated constitutional challenges and political pressure from the White House, SB 315 is struck down or repealed, there may be little room left for states to formally regulate generative AI. Either way, SB 315 is likely to be a bellwether.

II. Illinois Enters the Fray

Before delving into the particulars of SB 315, it is worth remarking on the present moment in AI development. A growing consensus holds that whoever controls AI will have enormous power to shape society's understanding of the world. As Elon Musk's xAI described in a recent lawsuit seeking to enjoin a 2024 Colorado AI statute, AI models "have the potential to undermine free inquiry and public discourse on an unprecedented scale: whoever owns or aligns the dominant models can quietly alter what counts as 'truth' -- portraying ideological fads or falsehoods as reality -- can filter inconvenient facts, and can curate entire worldviews."^[1]

The growth of generative AI thus raises the question: who should control this power to "curate entire worldviews?" As a political and legal matter, there are no easy answers. For now, the federal government has largely left AI developers to their own devices. Some states have, in turn, moved to fill the void. Colorado went the furthest when, in 2024, it passed SB 24-205, which imposed obligations on AI developers to avoid "algorithmic discrimination." But in response to political pressure and xAI's lawsuit -- which the United States joined in opposition to Colorado -- the State repealed SB 24-205 before it went into effect. (At the same time, Colorado enacted a scaled-down statute, SB 26-189, which is set to go into effect on January 1, 2027 -- a topic for another day.)

Enter now, Illinois SB 315.^[2] This bipartisan bill, which is awaiting Governor J.B. Pritzker's signature after passing the House 110-0 and the Illinois Senate 52-5, takes a middle course. It does not tell developers what their models should say or do, as

SB 24-205 had attempted. But neither does it adopt the hands-off approach of the federal government. Rather, SB 315 requires frontier AI developers like xAI and OpenAI to publish safety frameworks describing how they assess and mitigate catastrophic risks posed by AI models and their development. SB 315's approach is thus loosely analogous to federal securities laws in that it does not, on its face, empower Illinois to say which models or outputs are good or bad. Rather, it imposes a transparency requirement and then leaves it to the public -- buyer beware -- and other actors to make their judgments.

Granted, Illinois is not the first state to require published safety frameworks. California (SB 53) and New York (the RAISE Act) already mandate that frontier AI developers create and annually update plans addressing severe or catastrophic risks. But neither state has a clear mechanism (yet) for verifying that a developer is complying with its own safety framework. Illinois seeks to add verification through an audit requirement: under SB 315, independent third parties must conduct annual audits to determine whether developers are actually adhering to their stated protocols. If signed, and Governor Pritzker has indicated he is eager to sign it, SB 315 will become the first state law requiring independent third-party safety audits of frontier AI developers, thus adding a significant new layer of oversight (and oversight infrastructure) to an industry that has largely operated according to voluntary commitments to self-regulation.

The bill's proponents have been candid about their strategy of using SB 315 to promote national policy. Rep. Daniel Didech, who sponsored the legislation in the Illinois House, acknowledged that federal action would be preferable: "The states shouldn't be doing this."^[3] Rather, "[t]he best way to regulate these types of catastrophic risks would be a federal approach."^[4] But, he explained, "Congress has not taken up this issue yet, and the technology is developing at such a rapid pace that states have had

no choice but to step in.”[5] Notably, Didech suggested that SB 315 could serve as a template or a spur to further federal action: “Laws like this create a world where it’s more likely for the federal government to pass something.”[6]

III. SB 315’s Narrow Application

SB 315 is deliberately narrow. It applies only to “large frontier” developers -- companies with annual gross revenues exceeding \$500 million that train models using computing power greater than 10^{26} operations. That threshold limits the law’s reach to perhaps a dozen companies: OpenAI, Anthropic, Google DeepMind, Meta, xAI, and a handful of others. Everyone else -- startups, deployers, and downstream users -- is exempt from SB 315’s mandates.

For those frontier developers, SB 315 is focused on mitigating “catastrophic risks.” The bill defines catastrophic risk 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 or serious injury of more than 50 people, or more than \$1 billion in property damage, arising from a single incident involving certain enumerated conduct. That conduct is unlikely to be controversial. It consists of (1) providing expert-level assistance in creating or releasing a chemical, biological, radiological, or nuclear weapon; (2) engaging in conduct -- with no meaningful human oversight -- that constitutes a cyberattack or that would, if committed by a human, constitute murder, assault, extortion, or theft; or (3) evading the control of its developer or user. The definition also expressly excludes harms based on publicly accessible information, lawful federal government activity, or “harm caused by a frontier model in combination with other software if the frontier model did not materially contribute to the harm.”[7]

SB 315's core requirements take effect January 1, 2028. Under those requirements, large frontier developers must establish documented "frontier AI frameworks" -- technical and organizational protocols to assess, manage, and mitigate catastrophic risks. The frameworks must address, among other things, how the developer incorporates industry standards, defines capability thresholds, applies mitigations, uses third-party evaluators, and maintains cybersecurity for unreleased model weights. Before deploying new or substantially modified models, developers must also publish transparency reports disclosing intended uses, supported languages, output modalities, and summaries of catastrophic-risk assessments.

SB 315's signature feature is its audit mandate. Beginning in 2028, covered developers must retain independent third parties to audit their compliance with the Act on an annual basis. Auditors must possess demonstrated competence in frontier AI safety and work consistent with generally accepted auditing standards. Developers may not retain auditors with whom they share a financial interest. Controversially, auditors must be granted access to all materials reasonably necessary to perform the audit, including unredacted versions of published documents. The resulting audit report must describe whether the developer has substantially complied with the Act, identify material deviations and recommendations, and include a certification signed by the lead auditor. Within 30 days of receiving the report, the developer must publish a summary and redacted copy on its website and transmit copies to the Illinois Attorney General and Emergency Management Agency.

In addition, SB 315 establishes mandatory reporting for critical safety incidents. The bill defines a critical safety incident as: (1) unauthorized access to, modification of, or exfiltration of a frontier model's weights that results in death or bodily injury; (2) harm resulting from the materialization of a catastrophic risk; (3) loss of control of a frontier

model causing death or bodily injury; or (4) a frontier model using deceptive techniques against its developer to subvert controls or monitoring in a manner that demonstrates materially increased catastrophic risk. Within 72 hours of learning that a critical safety incident has occurred, developers must report the incident to the Illinois Emergency Management Agency and the Attorney General. If the incident poses an imminent risk of death or serious physical injury, the reporting window shrinks to 24 hours and extends to any law enforcement or public safety agency with jurisdiction. The bill also provides whistleblower protections. Developers are expressly forbidden from adopting any rules, policies, or contracts that prevent or retaliate against covered employees for disclosing information they reasonably believe indicates that the developer's activities pose a specific and substantial danger to public health or safety resulting from a catastrophic risk, or that the developer has violated the Act.

SB 315 is intended to be enforced exclusively through civil actions brought by the Illinois Attorney General; there is no private right of action. That said, it remains an open question as to whether private litigants may use violations of SB 315 to support other claims, such as violations of general consumer protection statutes. In all events, a large frontier developer that fails to publish or transmit required documents, makes materially false or misleading statements about catastrophic risk or framework compliance, fails to retain an independent auditor, fails to report a critical safety incident, or fails to comply with its own frontier AI framework is subject to civil penalties. The amount depends on the severity of the violation and can reach \$1 million for a first violation and \$3 million for all subsequent violations.

Finally, the bill includes a federal interoperability provision. The Illinois Emergency Management Agency may designate federal laws, regulations, or guidance as "substantially equivalent" to SB 315's

requirements if the federal standards impose equivalent incident-reporting requirements, address catastrophic risk in similar ways, and require independent third-party audits. A developer that complies with such designated federal requirements is deemed in compliance with SB 315 -- though failure to meet those federal standards still constitutes a violation of Illinois law, potentially allowing the Illinois Attorney General to enforce federal requirements that the federal government itself declines to enforce.

This interoperability provision serves multiple objectives. It responds to the political argument -- emphasized in a December 2025 White House executive order -- that state-by-state AI regulation creates a compliance “patchwork” that hampers innovation.[8] Plus, by building in an off-ramp to federal standards, Illinois can argue that SB 315 is a stopgap measure rather than a permanent regulatory fixture with unique compliance costs. The provision may also, as detailed below, blunt constitutional challenges under the Commerce Clause.

IV. The Industry’s Reaction

Some of the major frontier AI labs have lined up behind SB 315. OpenAI’s chief of global affairs, Chris Lehane, told *WIRED* that the company’s AI policy is now oriented around passing a series of similar state laws -- a strategy that may reflect a preference for uniform national standards in the absence of federal legislation.[9] Similarly, Anthropic’s head of U.S. state and local government relations, Cesar Fernandez, described the law as helping to “establish a baseline that every leading AI developer is expected to meet.”[10] Both companies supported SB 315 throughout the legislative process.[11]

Not everyone in the industry is on board, however. The Computer and Communications Industry Association (CCIA) opposed the bill, warning that it “would effectively require companies to comply with

a process for which the State has not yet established the necessary infrastructure, standards, or qualified evaluators.”[12] CCIA urged Illinois to follow Virginia and Connecticut, which directed state agencies to study the feasibility of independent verification frameworks before mandating audits. TechNet, a coalition of tech executives, raised similar concerns: “We remain concerned that Illinois would effectively be requiring private actors to make highly subjective determinations requiring AI safety compliance without established national standards, certifications, or clear regulatory guardrails.”[13] Similarly, Chamber of Progress -- a trade group whose partners include Google, Apple, and Amazon -- sent a letter to Illinois lawmakers on the morning of the House vote asking them to oppose the bill.[14] According to the group’s CEO, the bill “would force companies to expose sensitive systems to untested auditors in a regulatory regime that’s all liability and no standards.”[15]

Regardless, the audit-ecosystem concern is substantial. Third-party AI safety auditing is not a mature market -- no broadly recognized certification standards or licensing structures currently exist. Rep. Didech pointed to “a developing robust ecosystem” including “large international accounting firms,” but whether audit capacity will scale by 2028 remains uncertain.[16] Without established standards, the requirement risks becoming, as CCIA warned, “process for process’s sake” with little prospect of “producing meaningful consumer protection or accountability outcomes.”[17] A lot of auditing detail, efficacy, and capacity needs to be developed, in other words, if SB 315 is to be a success.

V. The Likely Constitutional Questions

SB 315 faces a number of potential constitutional challenges. The three most likely appear to be (1) compelled speech under the First Amendment, (2) vagueness under the Fourteenth Amendment’s Due Process Clause, and (3) excessive burden on interstate commerce under the Commerce Clause.

Each could be raised by a challenger seeking to prevent SB 315 from going into effect in 2028.

A. Compelled Speech (First Amendment)

The constitutionality of SB 315's disclosure requirements under the compelled-speech doctrine is the most likely First Amendment battleground. As a general matter, disclosure requirements in commercial contexts receive lenient review under cases like *Zauderer v. Office of Disciplinary Counsel* if they mandate "purely factual and uncontroversial information" and are "reasonably related to the State's interest in preventing deception of consumers."^[18] But when compelled disclosures require speakers to promote contested government messages or adopt positions on matters of public debate, courts often apply strict scrutiny. In *National Institute of Family and Life Advocates v. Becerra*, for instance, the Court struck down a California law that required crisis pregnancy centers to disclose information about state-sponsored abortion services.^[19] Of note, the Court cautioned that "regulating the content of professionals' speech pose[s] the inherent risk that the Government seeks not to advance a legitimate regulatory goal, but to suppress unpopular ideas or information."^[20]

Here, SB 315's disclosures look more like factual information than content regulation. The bill's requirements -- published safety frameworks, audit results, incident reports -- require developers to disclose what protocols they have adopted, what auditors have found, and what incidents have occurred. These are descriptions of internal processes and external events, not ideological positions. A court applying *Zauderer* and *Becerra* could thus conclude that such disclosures are factual and reasonably related to the State's interest in public safety. California's AB 587 offers a useful contrast. There, the Ninth Circuit found a consumer protection law aimed at social media companies problematic because it compelled "every covered social media company to reveal its policy opinion

about contentious issues, such as what constitutes hate speech or misinformation.”[21] SB 315, by contrast, requires covered developers to disclose their policies regarding catastrophic risks, and no one appears to dispute the importance of having a policy to mitigate these risks.

The differences between SB 315 and Colorado’s repealed SB 24-205 are also readily apparent. SB 24-205 defined “algorithmic discrimination” in terms that required distinguishing between disfavored differential treatment and favored discrimination intended to “increase diversity or redress historical discrimination.”[22] xAI’s complaint argued that this “internally contradictory, politicized definition” required developers to adopt contested ideological positions about which forms of discrimination are acceptable -- a ready trigger for strict scrutiny.[23] By narrowly defining “catastrophic risk,” SB 315 largely, if not entirely, avoids this issue. That’s not to say a First Amendment challenge has no chance of success -- but SB 315’s drafters appear to have written the bill with an eye toward surviving First Amendment scrutiny.

B. Void for Vagueness (Due Process Clause)

A challenger might also argue that key terms in SB 315 are impermissibly vague under the Fourteenth Amendment’s Due Process Clause. The void-for-vagueness doctrine requires that laws give regulated parties fair notice of what conduct is prohibited and include sufficient standards to prevent arbitrary enforcement. Here, SB 315’s definition of “catastrophic risk” is more limited than the term might suggest -- it is tied to specific quantified thresholds (50 or more deaths or serious injuries, or \$1 billion in property damage) and enumerated categories (CBRN weapons assistance, autonomous cyberattacks or crimes, loss of control). But several terms within the definition are undefined. The second prong, for instance, covers conduct that “if the conduct had been committed by a human, would constitute the crime of murder.” What exactly this

means is unclear. Does it apply only when AI autonomously kills 50 or more people, or does it also apply if AI assists a human in doing so? “Single incident,” “expert-level assistance,” “materially contribute,” “meaningful human oversight,” and “evading the control” are similarly ambiguous.

Nonetheless, a void-for-vagueness challenge would likely face an uphill battle. Courts apply the vagueness doctrine less stringently to civil regulatory statutes than to criminal laws. Moreover, regulatory statutes routinely employ general terms that agencies and courts later interpret. SB 315’s definitions appear no vaguer than many provisions that have survived constitutional scrutiny. But the argument is at least colorable, particularly given that violations carry civil penalties of up to \$3 million.

C. Interstate Commerce (Dormant Commerce Clause)

Commerce Clause considerations present a third issue. A challenger could argue that SB 315 effectively dictates nationwide conduct, thereby burdening interstate commerce, because AI models and related safety protocols cannot be geographically segmented. But this argument faces considerable headwinds. In *National Pork Producers Council v. Ross*, the Supreme Court cautioned that “extreme caution” is warranted before courts strike down non-discriminatory state laws affecting interstate commerce under the Dormant Commerce Clause.^[24]

That caution is likely to prevail here. SB 315 does not appear discriminatory -- it does not advantage Illinois consumers over those out of state. Further, it targets only the largest developers, each of whom presumably has comprehensive safety policies already in place. And the state interest is compelling. Preventing mass casualties, WMD proliferation, and critical infrastructure attacks are among a state’s most important obligations. The federal interoperability provision further reduces

fragmentation risk. Whether these design choices will insulate SB 315 from executive-branch pressure remains an open question, but as a Commerce Clause matter, and more broadly, a constitutional matter, the bill appears to be on solid footing.

VI. The Political Calculus

The constitutional questions are important, but SB 315 does not take effect until 2028. Courts are unlikely to seriously engage with it until mid-to-late 2027. And by then, the AI world, which continues to move at lightning speed, may look very different. Auditing standards may have matured or progress may have stalled. Congress may have passed -- or declined to pass -- federal legislation. Other states may have enacted their own AI laws, thereby creating a critical mass of support for state regulation or an unworkable patchwork of competing regimes. Public sentiment could be strongly positive on AI -- or wildly negative. No one knows which way the political winds will be blowing.

The political winds, whatever they end up being, may matter more than the constitutional issues. Colorado's SB 24-205 had been facing political headwinds for years before xAI filed suit. Governor Jared Polis expressed reservations about the bill's disparate-impact framework from the start. Industry pressure mounted steadily. A legislative working group was convened to explore amendments. By the time xAI's lawsuit was filed and the DOJ intervened, Colorado's political retreat on SB 24-205 was well underway -- and may well have resulted in SB 24-205's repeal and replacement even if no lawsuit had been filed.

If the White House decides to challenge SB 315, it has a variety of tools at its disposal. Executive Order 14365, issued in December 2025, declared that state-by-state AI regulation creates a compliance "patchwork" that hampers innovation and directed federal agencies to work toward preemption. The

DOJ's AI Litigation Task Force is actively looking for cases to join. Commerce Department officials have floated the possibility of conditioning federal broadband funding on states' willingness to refrain from AI regulation. Such political and economic levers can be as effective as litigation in upending state AI regulation -- perhaps more so.

That said, SB 315 appears better positioned to withstand federal pressure. Unlike Governor Polis, who was ambivalent about Colorado's original bill, Governor Pritzker, who has been no stranger to public disputes with President Trump, has been enthusiastic about Illinois' SB 315. OpenAI and Anthropic also publicly support it. And it cleared the Illinois legislature with overwhelming, almost unheard-of bipartisan support. All signs point to Illinois standing firm, at least for now, on SB 315's enactment.

Finally, one newly emerging dynamic bears special attention. On June 1, 2026, the State of Florida sued OpenAI and Sam Altman under state consumer protection laws, alleging that OpenAI's frontier models have "aided and abetted" mass shootings, encouraged suicide, caused public humiliation, and addicted minors "to a tool that feigns human compassion to collect their data with no parental oversight."^[25] Depending on how that lawsuit -- and the others that will inevitably follow -- plays out, the factual backdrop for any challenge to SB 315 may be vastly different than what exists today. How those dynamics unfold will shape not only SB 315's fate but, in all likelihood, the future of AI regulation nationwide.

VII. Illinois' Enduring Audit Requirement

SB 315 may ultimately be blessed by the courts, preempted, overtaken, or superseded. But for now, it appears more likely to endure than not. The constitutional issues are real but probably surmountable: the disclosure requirements fit comfortably within the *Zauderer* framework, the

definitions are no vaguer than many that have withstood scrutiny, and courts would be hard-pressed to strike down a non-discriminatory safety law about “catastrophic risks” on Dormant Commerce Clause grounds. The political challenges are also real. But SB 315’s bipartisan support, the governor’s enthusiasm, the endorsement of key AI frontier developers, and the political difficulty of arguing against, for example, regulation of WMD how-to guides provide a strong political foundation.

No matter how these issues shake out, though, SB 315’s core contribution to AI governance is likely to endure. In some form or fashion -- at the state or federal level or both -- frontier AI developers are likely going to have to publish, pass audits on, and adhere to safety frameworks governing catastrophic AI-created risk. To the public, who for years have been inundated by warnings from AI developers themselves that AI poses enormous, even extinction-level risks, this is a modest ask. It will be hard for any challengers to dislodge. Whether SB 315 itself becomes a baseline, a template, or simply a catalyst, the era of unaudited frontier AI development is likely drawing to a close. Lawyers, their clients, and the would-be auditors should prepare accordingly.

Citations

[1] Complaint ¶ 1, *xAI LLC v. Weiser*, No. 1:26-cv-01515 (D. Colo. Apr. 9, 2026).

[2] *See* S.B. 315, 104th Gen. Assemb., Reg. Sess. (Ill. 2026).

[3] Jared Perlo, *Illinois Legislature Passes Historic AI Bill That Would Require Third-Party Safety Audits*, NBC News (May 27, 2026), <https://www.nbcnews.com/tech/tech-news/illinois-legislature-passes-historic-ai-bill-rcna347191>.

[4] *Id.*

[5] *Id.*

[6] Maxwell Zeff, *Illinois Lawmakers Just Passed America's Strongest AI Safety Bill*, Wired (May 28, 2026), <https://www.wired.com/story/illinois-pass-major-ai-safety-law-pritzker/>.

[7] S.B. 315.

[8] *See* Exec. Order No. 14365, Ensuring a National Policy Framework for Artificial Intelligence, 90 Fed. Reg. 58,499 (Dec. 11, 2025).

[9] Zeff, *supra* note 6.

[10] *Id.*

[11] *Id.*

[12] Comput. & Commc'ns Indus. Ass'n, *CCIA Comments on Illinois SB 315* (May 22, 2026), <https://ccianet.org/library/ccia-comments-on-illinois-sb-315/>.

[13] Jenna Schweikert, *Illinois Lawmakers Pass Landmark AI Accountability Bill*, Capitol News Ill. (May 27, 2026), <https://capitolnewsillinois.com/news/illinois-lawmakers-pass-landmark-ai-accountability-bill/>.

[14] Zeff, *supra* note 6.

[15] *Id.*

[16] Schweikert, *supra* note 13.

[17] Comput. & Commc'ns Indus. Ass'n, *supra* note 12.

[18] *See Zauderer v. Off. of Disciplinary Couns. of the Supreme Ct. of Ohio*, 471 U.S. 626, 651-52 (1985).

[19] *Nat'l Inst. of Fam. & Life Advoc. v. Becerra*, 585 U.S. 755 (2018).

[20] *Id.* at 771 (internal quotations omitted).

[21] *X Corp. v. Bonta*, 116 F.4th 888, 899 (9th Cir. 2024).

[22] Complaint ¶ 7, *xAI LLC*, No. 1:26-cv-01515.

[23] *See id.* ¶ 8.

[24] *Nat'l Pork Producers Council v. Ross*, 598 U.S. 356, 390 (2023) (internal quotations omitted).

[25] *See* Complaint ¶ 5, *Attorney General of the State of Florida v. OpenAI Global, LLC et al.*, No. 2026-CA-000295 (Fla. 10th Cir. Ct. June 1, 2026).

This Akerman Intelligence Alert is intended to inform clients and friends of the firm about developments that may affect their business. It is not intended as legal advice for any specific matter and does not create an attorney-client relationship. Please consult with counsel to determine how this information may apply to your specific situation.