

Colorado's Revamped AI Law Charts a New Course, Reaching Both Developers and Deployers Alike

New Media and Technology Law Blog

On August 11, 2026, the Colorado Attorney General proposed rules to “clarify and implement” the Colorado AI Act. For organizations and other deployers of AI tools that have previously considered most state AI regulation as not particularly relevant to their operations, the Colorado law’s hybrid approach should feel different.

Back in May 2026, Colorado Governor Jared Polis signed [SB 26-189](#) into law. SB 26-189 scraps the broader algorithmic discrimination risk management model of the state’s landmark 2024 AI law ([SB 24-205](#)) before its effective date and instead focuses on imposing obligations on both developers and deployers of covered automated decision-making technology (ADMT). Among other things, developers will need to provide deployers with information about their covered ADMT, while deployers will need to retain records, provide disclosures, and allow consumers who experience adverse outcomes to exercise certain rights (including in the employment context).^[1] Notably, while the law does not create a private right of action (transparency and documentation obligations will be enforced by the state attorney general), it does establish how fault may be allocated among developers and deployers in civil actions alleging unlawful discrimination under state law arising from a consequential decision materially influenced by a covered ADMT. Fault allocation rules for civil discrimination cases mean that proper documentation, contracting and vendor/customer management will be central to compliance.^[2] In August 2026, the Colorado Attorney General [released](#) its proposed rules pursuant to the law, which would turn the statute’s high-level disclosure, correction and human-review rights into detailed process requirements for developers, deployers and entities sitting between them in the AI supply chain.

Taking a wider view of the last few years' state AI-related enactments, Colorado's law is noteworthy because it does not fit neatly into the two tracks that have dominated many high-profile state AI laws: upstream frontier model safety laws and targeted laws addressing specific GenAI use cases, such as AI companions, chatbot safety, deepfakes, digital replicas and synthetic content transparency. SB 26-189 instead takes a broader downstream approach, regulating the use of covered ADMT in consequential decision-making and allocating responsibilities between developers and deployers.

The Two Tracks of State AI Regulation

We see passage of SB 26-189 as part of a wider trend of state efforts to regulate AI, which has begun to move from the predictable two-track nature of AI regulation (as Congressional efforts at federal regulation have sputtered).[\[3\]](#)

???**Track 1:** Some states have regulated upstream frontier model safety and transparency. This track targets a small set of developers building the most advanced general-purpose models and focuses on systemic or catastrophic risks, model safety, cybersecurity, transparency, and whistleblower protections. In July 2026, for example, Illinois Governor J.B. Pritzker signed [SB 315](#) (Pub. Act 104-0538), the "Artificial Intelligence Safety Measures Act," which requires certain large frontier developers to establish and publish AI safety frameworks and report critical safety incidents to the state, and, beginning in 2028, undergo annual independent third-party audits. The Illinois law follows the same frontier model regulatory track as California's [SB 53](#) ("Transparency in Frontier Artificial Intelligence Act") and New York's [RAISE Act](#), although the three laws differ in scope, thresholds and mechanics.[\[4\]](#)

???**Track 2:** Other state laws have targeted specific GenAI use cases and harms. Examples include: AI chatbot (or AI companion) safety, including age-related safeguards, disclosures/transparency and response protocols (e.g., [N.Y. GBL Art. 47, § 1700-1704](#); [Colorado HB26-1263](#); [Oregon SB 1546](#); [Georgia SB 540](#); [Nebraska LB525](#)), [digital replicas of performers](#) (e.g., [Cal. Labor Code § 927](#); [NY A8887B](#)), GenAI-generated patient communications involving clinical information (e.g., [Cal. Health & Safety Code § 1339.75](#)), [AI-related digital replicas and impersonation](#) (e.g., [Tenn. Code Ann. §47-25-1101 et seq.](#)), nonconsensual deepfakes and election deepfakes (e.g., [10 DE Code § 7802](#); [Maryland SB 141](#)), and AI tools that provide, advertise or appear to provide therapy or professional mental or behavioral healthcare (e.g., [Ill. Pub. Act 104-0054](#); [Nevada AB 406](#)).

For businesses and organizations that are only deployers of enterprise GenAI tools, Track 1-style frontier model laws, which ask questions like “What safety framework, testing, incident-reporting, cybersecurity or audit obligations apply before or after release?”, generally will not impose direct compliance obligations, although they may affect vendor diligence and contractual representations. Similarly, the Track 2-style laws, which generally take a more targeted, harm-specific approach to certain AI use cases (and often, specific industries like entertainment or social media), are not relevant to many businesses – even if they use AI in their operations, most companies simply do not deploy “companion”-style consumer-facing AI chatbots, use digital replicas, etc.

A Third Track of AI Regulation with a Wider Scope

What makes the revised Colorado AI law’s focus on downstream use of ADMT in specific domains so notable for many businesses and organizations is its hybrid approach, arguably signaling a third track of AI regulation that will matter more to entities developing and deploying AI tools than most existing state AI laws.

Unlike frontier model laws, which generally apply to a small group of companies developing the most advanced AI models, and unlike other targeted laws that are aimed at discrete use cases, Colorado’s SB 26-189 can reach a wide range of ordinary business uses of ADMT and applies to both developers and deployers. This means the law can reach developers that meet the statutory definition, regardless of whether they are developing frontier models or commercializing another computational technology that falls within the statute’s definition, and also businesses that use those ADMT tools to materially influence consequential decisions in certain domains.[\[5\]](#)

Colorado’s revised law is concerned with issues such as notice, explanation, meaningful human review, recordkeeping, and accountability in the decision-making process itself, as well as allocation of liability between developers and deployers in state law anti-discrimination actions arising from the use of a covered ADMT.

Compliance questions will come from both deployers and developers of different stripes:

???Is the technology being used to materially influence a consequential decision?

???What types of documentation must the developer provide to the deployer?

???What pre-use notice must the deployer provide to the consumer affected by the covered decision?

???What human review, recordkeeping or vendor-management processes must the deployer have in place?

???What are the risks of potential liability under state anti-discrimination law for both parties?

Other states have also begun to regulate ADMT in consequential domains, regardless of whether the underlying system is a frontier model or another computational technology. California, for example, will require businesses using ADMT for “significant decisions” to provide pre-use notices and honor applicable access and opt-out rights.[\[6\]](#) Connecticut’s sweeping new AI law ([SB 5](#)) is another hybrid approach, combining downstream automated employment-related decision-technology requirements with frontier model, AI companion, and synthetic content provisions.[\[7\]](#)

Compared with frontier model laws and many narrow use-case laws, Colorado’s approach may have practical significance for a broader set of organizations that may be currently experimenting with, or have already integrated, AI tools into high-impact decision-making processes.

Final Thoughts on the Larger Legal Landscape

California’s Transparency in Frontier Artificial Intelligence Act, New York’s RAISE Act, and Illinois’ Artificial Intelligence Safety Measures Act, to name a few Track 1-style laws, focus on the developers of the most advanced AI models and emphasize frontier AI safety. Other states have regulated what they perceived to be an array of potential harms from GenAI chatbots and GenAI outputs. Colorado’s SB 26-189 stands apart from the growing number of targeted state AI safety laws and separate frontier AI developer laws and instead regulates both developers and deployers surrounding the use of covered ADMT in consequential decisions, regardless of whether the underlying technology is a frontier model or even an LLM at all. In that respect, Colorado’s hybrid-style law is more relevant to a broader range of businesses that may already or may eventually deploy ADMT in covered domains or integrate third-party ADMT into consumer-facing decision processes.

With the new Colorado and Connecticut AI laws on the books, both laws add another layer to state AI regulation – one that is likely to matter more directly to deployers of enterprise AI applications than previous frontier model or targeted GenAI use laws. Moreover, since state legislatures often use other existing state laws as models (and with what looks like little consensus in Congress to pass comprehensive bipartisan federal AI legislation in the near term), other states will likely pick up on what Colorado and Connecticut have passed to fashion similar regulation.[\[8\]](#) Indeed, many existing state AI laws (including these two) have already borrowed concepts from the EU AI Act, including risk-tiering, model transparency obligations and prohibited use frameworks.

For companies operating across jurisdictions, the practical result is that AI governance can no longer be managed as either upstream model governance or isolated end-user compliance – Colorado-style laws force governance across the developer-deployer relationship. Developers and deployers will increasingly need to understand their defined roles in a fragmented regulatory environment, how covered ADMT tools are being used, on top of the contractual and technical controls that allocate responsibility among developer, deployer and user.

The authors would like to thank Aniket C. Mukherji, a Proskauer legal assistant, for his contributions to this post.

Appendix: Breakdown of the Revised Colorado AI Legislation

Colorado AI Legislation Background

In May 2024, Governor Polis signed SB 24-205 into law. The Act, which was originally set to take effect on February 1, 2026, imposed disclosure requirements on developers and deployers of high-risk AI systems and required them to use reasonable care to protect consumers from algorithmic discrimination. Deployers also had to implement risk management policies and complete impact assessments. The original AI law was criticized by some for being overbroad, prompting Governor Polis at the time to suggest revisions before its effective date. Last year, he signed a law extending SB 24-205's effective date to June 30, 2026. This April, the original law faced a legal challenge (including the U.S. Department of Justice intervening in the suit). Now, Colorado's 2024 law has been repealed and replaced by SB 26-189. SB 26-189 largely eliminates the prior law's affirmative reasonable care and emphasis on algorithmic discrimination, replacing it with a more targeted ADMT regime centered on documentation, disclosures, consumer rights and fault allocation.

Relevant Definitions

SB 26-189 contains the following definitions:

???**Automated decision-making technology (ADMT)** means “a technology that processes personal data and uses computation to generate output ... that is used to make, guide, or assist a decision, judgment, or determination concerning an individual.” Excluded technologies include anti-malware technologies, calculators, databases, certain consumer-facing natural language tools, and general internet infrastructure actions.

???**Consequential decision** means a decision “about a consumer that relates to the provision of or a consumer's access to, eligibility for, selection for, or compensation for a covered domain.” The definition also includes certain decisions involving differentiated price, cost sharing, compensation, or other material terms that are reasonably likely to fundamentally alter a consumer's access to a covered domain. Excluded are “low-stakes” or routine decisions that do not materially influence access to an opportunity, data processing functions that do not generate consumer inferences, and advertising, content moderation and cybersecurity technologies.

???**Covered ADMT** means ADMT “used to materially influence a consequential decision.”

???**Covered domain** means education, employment, the lease or purchase of residential real estate, a financial or lending service, insurance, healthcare services, or essential government services and public benefits.

???**Deployer** means a “person doing business in Colorado that deploys a covered ADMT.”

???**Developer** means a “person doing business in Colorado that (i) develops... or otherwise makes commercially available a covered ADMT; (ii) develops a component... designed, marketed...configured, or contracted to be used as part of a covered ADMT; or (iii) intentionally and substantially modifies an ADMT such that it becomes a covered ADMT.”

Developer Obligations – Documentation

The law requires developers to make certain information available to deployers of their covered ADMT in a reasonably understandable manner (while still protecting trade secrets). This includes a general statement describing the covered ADMT’s intended uses and known harmful or inappropriate uses, a description of the categories of training data used to train the covered ADMT (to the extent known), known limitations, instructions for appropriate use, monitoring and meaningful human review (where applicable), and information reasonably necessary for the deployer to comply with disclosure obligations. Developers must also notify deployers of material updates, intentional and substantial modifications, and changes to intended uses, limitations, or risk mitigation for the covered ADMT within a reasonable time (public release notes, with proper notice, are permitted).

Importantly, a developer is subject to the disclosure requirements only for a deployer’s use of a covered ADMT where the covered ADMT was “marketed, advertised, configured, contracted, sold, or licensed to be used to materially influence a consequential decision.”

Deployer Obligations – Disclosure

Before using a covered ADMT to materially influence a consequential decision, a deployer must provide a “clear and conspicuous notice” to consumers that it used or will use such technology, with instructions on how to obtain additional information. A prominent public notice reasonably accessible at points of consumer interaction (including through a link or posting) satisfies this requirement.

If a deployer’s use of a covered ADMT results in an adverse outcome for a consumer, the deployer must provide the following within 30 days:

1. A plain language description of the decision and the role the covered ADMT played in it (with notice of any information that is redacted to protect trade secrets or

information protected from disclosure by law). Disclosures can also be limited for cybersecurity, fraud prevention, AML, or other programs required by law.

???Instructions and a simple-to-follow process to request additional information about the covered ADMT.

???An explanation of the rights created by this law for consumers who experience adverse outcomes (outlined in the next section).

The state attorney general must adopt rules implementing these post-adverse outcome disclosure requirements by January 1, 2027. On August 11, 2026, the Colorado Attorney General's office [released](#) for public comment [proposed ADMT rules](#) under SB 26-189, as well as proposed rules under a separately passed chatbot safety law.

Consumer Rights

Consumers who experience adverse outcomes involving covered ADMTs may request, and deployers must provide:

1. Instructions for requesting personal data and correcting factually incorrect or materially inaccurate personal data used in a consequential decision that used a covered ADMT;[\[9\]](#)

???An opportunity for meaningful human review and reconsideration of the consequential decision, "to the extent commercially reasonable."

Enforcement by Attorney General

The state attorney general may enforce violations through the Colorado Consumer Protection Act; any violation is a "deceptive trade practice." Until January 1, 2030, before bringing an enforcement action, the attorney general must provide notice and a 60-day opportunity to cure if the attorney general deems cure possible (unless the attorney general can demonstrate that the developer or deployer knowingly or repeatedly violated the law). The law does not create a new private right of action or limit existing rights under state or federal law. It applies to consequential decisions made on or after January 1, 2027.

Liability – Fault Allocation

One of the law's most important provisions is its allocation of liability among developers and deployers for state anti-discrimination claims arising from a consequential decision materially influenced by a covered ADMT.

Under Section 6-1-1707, a developer or deployer may be held liable in a state anti-discrimination action arising from the use of a covered ADMT, with fault allocated based on their relative fault for the violation.

???A developer is liable only to the extent its covered ADMT was used by a deployer in a manner intended, documented, marketed, advertised, configured, or contracted for by the developer and the covered ADMT materially influenced a consequential decision giving rise to the violation.

???A developer is not liable for violations arising from a deployer's use of a covered ADMT in a manner not intended, documented, marketed, advertised, configured, or contracted for by the developer.

???The law does not limit a deployer's liability for its own independent acts or omissions in a consequential decision materially influenced by a covered ADMT, including off-contract uses.

???Any contractual indemnity for damages resulting from a party's own acts or omissions related to the use of ADMT in making consequential decisions in violation of Colorado anti-discrimination laws is void. However, this does not apply to a developer whose covered ADMT was used in an unintended manner and who complied with documentation requirements.

???The law clarifies that the use of an ADMT in a consequential decision itself does not provide a defense to any liability under state or federal law, including liability related to discrimination or consumer protection.

???The law does not create joint and several liability except to the extent permitted under existing law.

[1] Employment is one of the covered domains of SB 26-189 and the use of automated tools by employers raises additional workplace law considerations. Employer compliance issues related to the law are beyond the scope of this post. For a deeper dive into those issues, please see our colleagues' discussion on our [Law and the Workplace](#) blog.

[2] This post will not delve deeply into the provisions of the Colorado law, but we have included an appendix that breaks down the revamped Colorado AI law in greater detail.

[3] At the federal level, President Trump has criticized the prospect of a patchwork of state laws, including in an [executive order](#) he signed in December 2025 that advocated for “a minimally burdensome national standard — not 50 discordant State ones.” Given emerging developments related to AI model cybersecurity and other issues, President Trump signed a June 2026 [executive order](#) directing federal agencies to establish a classified benchmarking process to assess advanced cyber capabilities of AI models and determine the threshold at which an AI model should be designated a “covered frontier model.” The order also directs agencies to design a voluntary framework through which AI developers may engage the federal government to determine whether model(s) under development meet the designation of “covered frontier model” and provide the government pre-release access to covered frontier models under appropriate confidentiality. In August 2026, press reports [indicated](#) the administration had completed a draft framework.

[4] In a similar vein, the [California AI Transparency Act](#), which became operative on August 2, 2026, requires covered providers – that is, an entity that creates, codes, or otherwise produces a GenAI system with more than one million monthly visitors or users and is publicly accessible within the state – to make AI-detection tools available to users and to include latent disclosures in AI-generated image, video or audio content created by the covered provider’s system, among other requirements. AB 853 also adds later obligations for large online platforms and GenAI hosting platforms beginning January 1, 2027, and certain capture-device manufacturer obligations beginning January 1, 2028.

[5] A skeptical reader could say that Track 3 is merely a particularly broad version of Track 2, since “consequential decision-making” is itself technically a use case. However, Track 2 is product-/content-/interaction-specific regulation keyed to discrete AI-mediated harms; Track 3 is cross-sector decision-process regulation keyed to consequential decisions and the allocation of obligations among developer/deployer roles.

[6] In September 2025, the California Office of Administrative Law [approved](#) regulations covering businesses that use ADMT to make significant decisions, requiring them to provide consumers with a prominent pre-use notice, including information about how the business plans to use ADMT and a description of consumers’ rights to opt out. The regulations generally went into effect on January 1, 2026, but businesses using ADMT to make significant decisions must comply with the ADMT requirements beginning January 1, 2027.

[7] On May 27, 2026, Connecticut Governor Ned Lamont signed [SB 5](#) (Pub. Act No. 26-15), which most notably requires businesses to notify employees and job applicants when they use “automated employment-related decision technology” that materially influences employment decisions and to disclose certain information about their use of such technology, beginning October 1, 2027. SB 5 also imposes consumer disclosures for certain covered providers, child-related safeguards for AI companion chatbots, requirements that covered providers include provenance or embedded metadata in certain AI-generated content or materially altered audio, image, or video content to the extent commercially and technically reasonable, and certain governance and whistleblower-related obligations on frontier AI developers. It also establishes a pilot program to evaluate the use of independent verification programs to assess the harm-prevention practices of AI models; evidence of good standing would be admissible in a private civil action for injury or property damage caused by an AI model. For a deeper dive into the employment-related issues with SB 5, please see our colleagues’ discussion on our [Law and the Workplace](#) blog.

[8] See, e.g., NJ [S4075](#) (2026), as an example of a pending state bill that centers on downstream uses of AI and automated decision systems. Like Colorado’s SB 26-189 and Connecticut’s SB 5, the bill would regulate automated systems used in consequential contexts (specifically employment decisions and public benefits or services), rather than targeting only developers of frontier AI models. See also South Carolina [S.963](#) (2025-2026).

[9] The statute (§ 6-1-1705(1)(c)) states that this consumer right of correction does not require correction of “opinions, predictions, scores, or protected evaluations.”

[Related Professionals](#)

??Peter J. Cramer

Associate