

Colorado's Mile High AI Act: 6 Key Takeaways

Client Alert | May 10, 2024

If SB 205 is signed into law, it would go into effect on February 1, 2026, and Colorado would become the first state to enact legislation regulating the development and deployment of high-risk AI systems generally. On May 8, 2024, Colorado's Legislature passed SB24-205, the [Colorado Artificial Intelligence Act](#) ("SB 205"). SB 205 seeks to govern the use of high-risk AI systems in the private sector. If SB 205 is signed into law by Colorado Governor Jared Polis—which he is expected to do—it would go into effect on February 1, 2026, and Colorado would become the first state to enact legislation regulating the development and deployment of high-risk AI systems generally. Although SB 205 would be the most comprehensive AI-specific state law, it is not the only state to move in this area in 2024. This year alone, [Utah](#) and [Tennessee](#) enacted AI legislation (tackling consumer deception by generative AI and AI deepfakes, respectively), while the California Consumer Privacy Protection Agency ("CPPA") has been making progress with its [draft regulations](#) related to automated decision-making technology ("ADMT"). SB 205 is effectively an anti-discrimination law that would regulate the use of high-risk AI systems by imposing a slew of requirements on developers and deployers, including notice, documentation, disclosures, and impact assessments. SB 205's focus on high-risk AI systems is similar to the risk-based approach taken by the European Union's AI Act. Accordingly, companies looking to design a compliance regime to respond to these developments may find opportunities for overlap in these frameworks (e.g., by leveraging ISO's [42001](#)). Structurally, SB 205 would become Part 16 within Colorado's Consumer Protection Act, which already houses the Colorado Privacy Act. SB 205 expressly states that Part 16 does *not* provide the basis for a private right of action and that the Attorney General has "exclusive" enforcement authority. Below are 6 key takeaways. **6 Key Takeaways for the Private Sector**

Related People

[Vivek Mohan](#)

[Cassandra L. Gaedt-Sheckter](#)

[Natalie J. Hausknecht](#)

[Eric D. Vandeveld](#)

- 1. Broad Cross-Sectoral Coverage of High-Risk AI Systems:** High-risk AI systems are defined as any AI system that, when deployed, makes, or is a *substantial factor* in making, a *consequential decision*. A "consequential decision" is one that has a "material legal or similarly significant effect on the provision or denial to any consumer of, or the cost or terms of" education, employment or an employment opportunity, financial/lending services, housing, insurance, healthcare, essential government services, and legal services. Meanwhile, a "substantial factor" must (1) assist in making the consequential decision; (2) be capable of altering the outcome of a consequential decision; and (3) be generated by an AI system. There is ambiguity regarding what this means in practice as it is unclear what would constitute "assisting" in the consequential decision or being "capable" of altering the outcome. This language bears some similarity to that of the CPPA's current draft ADMT regulations, which would cover training ADMT that is merely *capable* of being used for a significant decision concerning a consumer.
 - Further, unlike the Colorado Privacy Act, SB 205 does not provide an exemption for the employment context. Instead, a "consumer" is defined as "an individual who is a Colorado resident," and the law specifically intends to cover consequential decisions including related to employment and employment opportunities.
 - Examples of specifically excluded tools include calculators, databases, data storage, anti-virus software, networking, spreadsheets, spam-filtering,

data storage, cybersecurity, and chatbots subject to an accepted use policy prohibiting the generation of discriminatory or harmful content. The latter exclusion could be fairly significant given the number of “chatbots” being deployed by companies as could the exclusion of tools for cybersecurity—which often are subject to discussion under privacy laws given their unique but sometimes significant use of information.

2. **Exclusive Enforcement by the Attorney General:** SB 205 provides that the Attorney General would have “exclusive” authority to enforce the law and promulgate rules to implement the law regarding documentation, notice, impact assessments, risk management policies and programs, rebuttable presumptions, and affirmative defenses. The text specifies that violations of SB 205 do *not* provide the basis for a private right of action. Notably, SB 205 provides the following two affirmative defenses if the Attorney General commences an action.
 - **Robust AI Governance Programs:** SB 205 would provide an affirmative defense if a deployer has implemented and maintained a risk management policy or program that complies with national or international risk management frameworks such as the National Institute of Standards and Technology’s (“NIST”) AI Risk Management Framework (“AI RMF”) or the International Organization for Standardization’s (“ISO”) 42001.
 - **Cured Violations:** SB 205 would also provide an affirmative defense for a developer or deployer that discovers and cures the violation due to (a) feedback, (b) adversarial testing or red teaming (under NIST’s definition), or (c) an internal review process and is otherwise in compliance with NIST’s AI RMF or ISO’s 42001.
3. **Developers and Deployers Are Subject to an Anti-Algorithmic Discrimination Duty:** SB 205 expressly covers both developers and deployers of high-risk AI systems and would require both to use *reasonable care* to protect consumers from any known or reasonably foreseeable algorithmic discrimination.
 - **Algorithmic discrimination** is defined as any condition in which the use of an AI system results in unlawful differential treatment or impact based on an array of protected classes under Colorado and federal law, including race, disability, age, gender, religion, veteran status, and genetic information. Using an AI system to expand an applicant pool to increase diversity or remedy historical discrimination would *not* constitute algorithmic discrimination under SB 205. The law provides a narrow exemption for certain deployers with fewer than 50 employees that do not use their own data to train or further improve the AI system.
 - **A rebuttable presumption is available in the event of an enforcement action.** The law would establish a rebuttable presumption that *reasonable care* was used to avoid algorithmic discrimination if certain compliance indicators (which differ between developers and deployers) are met:
 - Compliance indicators for developers include: (a) providing sufficient information and documentation to deployers such that an impact assessment can be completed; (b) disclosing to the Attorney General and deployers any known or reasonably foreseeable risk of algorithmic discrimination within 90 days of discovery; (c) publishing a publicly available statement regarding the high-risk systems developed and how any known or reasonably foreseeable risks of algorithmic discrimination are being managed; and (d) the purpose and intended benefits and uses of the AI system.
 - Meanwhile, compliance indicators for deployers include: (a) implementing a risk management policy and program; (b) completing an impact assessment; (c) providing notice to

consumers; (d) disclosing to the Attorney General any algorithmic discrimination within 90 days of discovery; and (e) publishing a publicly available statement summarizing the high-risk AI system being deployed and any known or reasonably foreseeable risks of algorithmic discrimination that may arise.

4. **Impact Assessments Required:** In alignment with trends in other proposed state legislation, SB 205 would require deployers to complete an impact assessment annually, and also within 90-days of any intentional or substantial modification to the high-risk AI system. The impact assessment must include the purpose, intended use cases, benefits, known limitations, and deployment context of the high-risk AI system, any transparency measures taken, post-deployment monitoring and safeguards implemented, and the categories of data used as inputs and the outputs produced. Notably, deployers would be permitted to use a comparable impact assessment that was completed for purposes of complying with another applicable law or regulation. As noted above, completing an impact assessment is one of the indicators that would support a deployer in establishing a rebuttable presumption that reasonable care was used to avoid algorithmic discrimination.
5. **Notice to Consumers is Key:** Similar to other, more narrow AI state and local laws already in effect (e.g., Utah's AI Policy Act and New York City's [Local Law 144](#)), deployers must notify consumers of the use of a high-risk AI system, the purpose of the system, the nature of the consequential decision, a description of how the system works, and, *if applicable*, the consumer's right to opt out of the processing of personal data for purposes of profiling under Section 6-1-1306 of the [Colorado Privacy Act](#). Notably, consumers subject to an adverse consequential decision must be provided with an opportunity to appeal the decision. In alignment with the European Union's AI Act, if it is "obvious" that a consumer is interacting with an AI system, SB 205 would not mandate such a disclosure.
6. **A Violation is Also a "Deceptive Trade Practice" Under Colorado Law:** On its final page, SB 205 provides that a violation of Part 16 would constitute a "deceptive trade practice" under Colorado Revised Statutes, Section 6-1-105, which resides in Part 1 of Colorado's Consumer Protection Act. Note that under Part 1, consumers injured by a "deceptive trade practice" are provided with the ability to bring a civil action. At this stage, it remains unclear whether this was intended to indirectly create a private right of action under Part 1, or if the legislature inadvertently failed to make an express disclaimer (e., "Notwithstanding any provision in Part 1, Part 16 does not authorize a private right of action.").

The following Gibson Dunn lawyers assisted in preparing this update: Vivek Mohan, Cassandra Gaedt-Sheckter, Natalie Hausknecht, Eric Vandeveld, and Emily Maxim Lamm.

Gibson, Dunn & Crutcher's lawyers are available to assist in addressing any questions you may have regarding these issues. Please contact the Gibson Dunn lawyer with whom you usually work, any leader or member of the firm's Artificial Intelligence practice group, or the authors: Cassandra L. Gaedt-Sheckter – Palo Alto (+1 650.849.5203, cgaedt-sheckter@gibsondunn.com) Natalie J. Hausknecht – Denver (+1 303.298.5783, nhausknecht@gibsondunn.com) Vivek Mohan – Palo Alto (+1 650.849.5345, vmohan@gibsondunn.com) Robert Spano – Paris/London (+33 1 56 43 14 07, rspano@gibsondunn.com) Eric D. Vandeveld – Los Angeles (+1 213.229.7186, evandeveld@gibsondunn.com) Emily Maxim Lamm – Washington, D.C. (+1 202.955.8255, elamm@gibsondunn.com) © 2024 Gibson, Dunn & Crutcher LLP. All rights reserved. For contact and other information, please visit us at www.gibsondunn.com. Attorney Advertising: These materials were prepared for general informational purposes only based on information available at the time of publication and are not intended as, do not constitute, and should not be relied upon as, legal advice or a legal opinion on any

GIBSON DUNN

specific facts or circumstances. Gibson Dunn (and its affiliates, attorneys, and employees) shall not have any liability in connection with any use of these materials. The sharing of these materials does not establish an attorney-client relationship with the recipient and should not be relied upon as an alternative for advice from qualified counsel. Please note that facts and circumstances may vary, and prior results do not guarantee a similar outcome.

Related Capabilities

[Artificial Intelligence](#)

[State Attorneys General](#)