

28 MAY 2024 • 11 MINUTE READ

SHARE

Colorado enacts first-in-the-nation comprehensive AI guardrails

Written by: [Danny Tobey](#), [Tony Samp](#), [Coran Darling](#), [Connor Scott](#), [Todd Mobley](#), [Ted Loud](#)

On May 17, 2024, Colorado became the first US state to enact a regulatory framework governing the development and use of artificial intelligence when Governor Jared Polis (D) signed into enactment [Senate Bill 24-205](#).

The new law, which does not go into effect until February 1, 2026, imposes requirements on developers and deployers of “high-risk” AI systems to, among other goals, avoid “algorithmic discrimination” and require reporting of any instances to the attorney general’s office. Under the law, any AI system that is a substantial factor in making “consequential decisions” qualifies as “high-risk.”

The law also imposes notification requirements, mandating that deployers of high-risk AI systems inform consumers that they are subject to the high-risk AI system making consequential decisions about them.

Controversial law signed “with reservations”

In a message to the state legislature, Governor Polis stated that he was signing the bill “with reservations,” but that he hoped his action “furthers the conversation, especially at the national level.”

The governor noted that the legislation deviates from the standard practice of prohibiting discriminatory conduct and instead regulates “the results of AI system use, regardless of intent,” and he urged the legislature to revisit this issue before the law takes effect.

Governor Polis also said the law “creates a complex compliance regime for all developers and deployers doing business in Colorado, with narrow exceptions for small deployers.” He expressed concern “about the impact this law may have on an industry that is fueling critical technological advancements across our state for consumers and enterprises alike.”

Given the nearly two-year interval until the law goes into effect, Governor Polis urged legislators and stakeholders to work together to “fine tune the provisions and ensure that the final product does not hamper development and expansion of new technologies in Colorado that can improve the lives of individuals across our state.”

The AI legislation has proven to be one of the more controversial measures in the current legislative session, and the governor heard calls to veto the bill from organizations representing the tech industry, including the Chamber of Progress, the Consumer Technology Association, and the US Chamber of Commerce, who said the legislation would stifle innovation. Small business owners also testified in opposition during deliberations on the bill, arguing that it would put small businesses at a disadvantage against large corporations.

Notable provisions of the new law

Much like the recently finalized (though pending official publication in the Official Journal of the European Union) [EU AI Act](#), the law has different requirements for developers of AI systems and for those who deploy those systems. Many terms and provisions, including risk classifications, take inspiration from the EU AI Act – an increasingly common occurrence in developing regulation across the world as legislators learn from the challenges faced by the EU drafters.

the world as legislators learn from the challenges faced by the EU drafters.

For developers, the legislation requires that they “use reasonable care to avoid algorithmic discrimination in the high-risk system[s]” they design, according to a summary provided by the Colorado General Assembly.

In a similar approach to developing product liability rules in the EU through the [AI Liability Directive](#), the law includes the concept of a rebuttable presumption in the context of compliance. The EU has adopted a consumer-favored presumption that states that a party will be held liable, unless proven otherwise, where it is established that they have not complied with their obligations under applicable law. In contrast, Colorado’s approach is more favorable to deployers, as it presumes that a deployer has acted with reasonable care, where they can prove that they have complied with specific provisions in the bill, including:

- Making available to a deployer of the high-risk system a statement disclosing specified information about the high-risk system
- Making available to a deployer of the high-risk system information and documentation necessary to complete an impact assessment of the high-risk system
- Publishing a list summarizing the types of available high-risk systems that they have developed or substantially modified. The list must also state how the developer manages any known or reasonably foreseeable risks of algorithmic discrimination that may arise from each of these high-risk systems
- Disclosing to the attorney general, and known deployers of the high-risk system, any known or reasonably foreseeable risk of algorithmic discrimination within 90 days after the discovery or receipt of a credible report
- Implementing a risk management policy and program for the high-risk system
- Completing an impact assessment of the high-risk system
- Annually reviewing the deployment of each high-risk system to ensure that the high-risk system is not causing algorithmic discrimination
- Providing notice to consumers who are subject to consequential decisions made by the high-risk system
- Providing a consumer with an opportunity to correct any incorrect personal data that a high-risk artificial intelligence system processed in making a consequential decision, and
- Providing a consumer with an opportunity to appeal, via human review if technically feasible, an adverse consequential decision concerning them.

Deployers, on the other hand, have lighter requirements, including:

- Implementing a risk management policy to identify and mitigate the risks of algorithmic discrimination
- Completing an impact assessment for any deployed high-risk AI systems
- Updating the impact assessment annually and following any substantial modifications to the high-risk AI system
- Providing notice to consumers regarding the operation of a high-risk AI system that makes consequential decisions, before those decisions are made
- Publishing a list of the types of high-risk AI systems that are deployed, how the deployer manages the risks of algorithmic discrimination, and the nature and sources of information that the deployer uses for the high-risk AI system, and
- Providing notice to the Attorney General of any detected incidents of algorithmic discrimination within ninety days of discovery of the same.

Notable definitions

To provide clarity to several terms new to the regulation of AI in the US, the law provides several notable definitions to assist compliance with and enforcement of its provisions.

The law, for example, defines “**algorithmic discrimination**” as “any condition in which the use of an artificial intelligence system results in an unlawful differential treatment or impact that disfavors an individual or group of individuals on the basis of their actual or perceived age, color, disability, ethnicity, genetic information, limited proficiency in the English language, national origin, race, religion, reproductive health, sex, veteran status or other classification protected under the laws or this state or federal law.”

The definition does, however, exempt self-testing by developers or deployers from being required to identify, mitigate, or prevent discrimination, or to otherwise ensure compliance with state and federal law. Similar exemptions are present that cover expanding an applicant, customer, or

regulation. Similar exemptions are present that cover expanding an applicant, customer, or participant pool to increase diversity or redress historical discrimination. Private clubs or other establishments not open to the public are also exempt.

“Artificial intelligence system” means any machine-based system that explicitly or implicitly infers from the inputs the system receives how to generate outputs, including content, decisions, predictions, or recommendations that can influence physical or virtual environments. While closely resembling the definition of the OECD, which is often referred to in defining the technology, the definition is absent of several additional characteristics, including levels of autonomy and adaptability.

“Consequential decision” is defined as having material legal or similarly significant effect on the provision or denial to consumers of educational enrollment or opportunity, employment or employment opportunities, financial or lending services, essential government services, healthcare services, housing, insurance, or legal services.

Although the text of the new law says little about employment relationships, the inclusion of “employment or employment opportunities” in this definition could have wide-ranging effects on employers, employees, and job applicants. For example, the new law could be read as encompassing all forms of employment-related decisions, including those related to recruiting and hiring, disciplinary actions and performance management, compensation, terminations, and actions taken with respect to or based on productivity or other employee monitoring efforts.

Employers subject to or familiar with New York City Local Law 144 should take note that scope of this new law is much broader, and that its potential implications are more extensive than the implications of New York City’s law, which applies only to the use of certain AI systems that substantially assist or replace discretionary human decision making with respect to hiring or promotion decisions.

A state-by-state v federal approach

The new law, which is limited to application in Colorado, comes at a time of heightened activity in the US Congress aimed at crafting nationwide AI standards and regulations.

Earlier this month, a bipartisan Senate AI Working Group led by Majority Leader Chuck Schumer (D-NY) released what it called [A Roadmap for AI Policy in the United States](#), providing a framework for legislating at the federal level, including the types of anti-bias and pro-consumer issues redressed in the Colorado statute. See our [client alert](#) for more details and insights.

Various federal agencies are also releasing guidance on the use of AI, including the [Equal Employment Opportunity Commission](#), the [National Labor Relations Board](#) General Counsel, and, most recently, the [Department of Labor and Office of Federal Contract Compliance Programs](#).

In recent congressional hearings on AI-related topics from deepfake election manipulation to copyright protections, federal legislators have repeatedly stressed the importance of adopting uniform federal standards to preempt a “patchwork” of state laws.

In his [signing statement](#), Governor Polis himself stressed the need for federal action, expressing concern that varying state-by-state laws and regulations could lead to confusing compliance burdens and an unlevel playing field.

However, as Colorado has demonstrated, state legislatures can often move more quickly than Congress. Impatience over the slow pace of progress in Washington may continue to be a motivating factor in state capitals across the US as policymakers seek to stay on top of technological advancements that are moving at a rapid pace.

State regulators are also moving forward. For example, on May 17, 2024, the California Civil Rights Council [announced the release of proposed regulations](#) to protect against discrimination in employment resulting from the use of AI, algorithms, and other automated decision-making systems. States may also be encouraged, and take inspiration from, international developments, such as the recent finalization of the EU’s own regulatory framework for AI.

As other states move to adopt their own laws and regulations governing AI, potentially contentious issues *may* arise between state officials who favor stricter, or more permissive, standards than the federal government eventually adopts. A similar federal-state disconnect has previously caused issues for federal privacy protection legislation, with California lawmakers worried their state’s existing strong privacy laws could be undermined or overridden by federal laws.

Given policymakers’ goal of avoiding a duplicative – or contradictory – patchwork of state laws, perhaps the decisive step taken by Colorado to enact its own AI regulatory regime will serve to spur action by Congress. And, since the state law does not go into effect until 2026, there may be time for federal lawmakers to address many of the issues covered in the Colorado statute, either as part of a wider-ranging package or on an individual piecemeal fashion.

DLA Piper

As part of the *Financial Times*’s 2023 North America Innovative Lawyer awards, DLA Piper’s was conferred the Innovative Lawyers in Technology award for its AI and Data Analytics practice.

DLA Piper’s AI policy team in Washington, DC is led by the Founding Director of the Senate Artificial Intelligence Caucus.

For more information on AI and the emerging legal and regulatory standards, visit DLA Piper's [focus page on AI](#).

Gain insights and perspectives that will help shape your AI strategy through our newly released [AI Chatroom series](#).

DLA Piper's AI Practice has over 100 attorneys, data scientists, coders, and policymakers focused on AI worldwide. To learn more about this evolving policy landscape and its implications for your business, please contact any of the authors.

For more information, please join our webinar, "[Artificial Intelligence Roadmap: Unraveling the Senate's latest AI framework](#)."

Related insights

Publication

Senate unveils long-anticipated AI roadmap

15 MAY 2024 • 12 MINUTE READ

Publication

US House of Representatives launches AI Task Force

23 FEBRUARY 2024 • 9 MINUTE READ

Publication

Deciphering OMB's guidance to agencies on AI: A blueprint for responsible AI in government

11 APRIL 2024 • 9 MINUTE READ

Related capabilities

Artificial Intelligence and Data Analytics

Employment

Regulatory and Government Affairs

Data, Privacy and Cybersecurity

Healthcare

Financial Services

