When Colorado lawmakers rewrote the state’s landmark artificial intelligence law in May 2026, the decision looked like a pragmatic retreat.

The state abandoned much of the country’s most ambitious attempt to regulate algorithmic discrimination: the risk that automated systems reproduce or intensify bias based on race, sex, age or other protected characteristics when making decisions about employment, credit, insurance or healthcare.

In place of that framework, the legislature passed a much narrower statute on May 14. The new Automated Decision-Making Technology Act largely focuses on transparency. Rather than requiring companies to prevent discriminatory outcomes, it requires them to disclose when automated technology has influenced an important decision.

That change may make the law less burdensome for industry. It may also make one constitutional objection disappear. But it does not necessarily place Colorado’s AI policy on safer legal ground.

In fact, by stripping the law down to a collection of notice and explanation requirements, lawmakers may have made another constitutional challenge easier to pursue.

Colorado’s experience captures a broader transformation in state AI policy. Faced with aggressive industry opposition and an increasingly hostile federal government, state lawmakers are moving away from prevention and toward disclosure. Instead of telling companies how to govern high-risk systems, states may increasingly settle for telling them what they must reveal.

That is not merely a technical adjustment. It reflects a fundamentally different view of who should bear responsibility when automated systems cause harm.

From preventing harm to explaining it afterward

The original Colorado Artificial Intelligence Act was built around proactive obligations.

Developers and deployers of high-risk AI systems — including tools used to screen job applicants, assess loan applications, determine insurance eligibility or influence healthcare decisions — were required to establish risk-management programs.

Companies also had to conduct impact assessments and take reasonable steps to prevent algorithmic discrimination. The basic philosophy was clear: Institutions deploying consequential AI systems should examine those systems, identify foreseeable risks and demonstrate that safeguards are in place.

The law did not simply give consumers information. It imposed responsibilities on the organizations with the greatest access to the technology, the data and the decision-making process.

The replacement statute takes a very different approach.

Companies must notify consumers when automated technology plays a substantial role in a consequential decision. If the decision is adverse — such as rejecting a job applicant or denying a loan — the company generally has 30 days to explain how the technology contributed to the result.

Consumers may request corrections to inaccurate personal information. They may also seek meaningful human review.

Those protections are not trivial. Many people currently do not know whether an algorithm influenced a decision affecting their livelihood, finances or health. A right to explanation can

help consumers identify errors, challenge inaccurate information and demand accountability.

But disclosure is not prevention.

The new law eliminates the earlier duty of care, risk-management requirements and impact assessments. It no longer requires companies to examine whether their systems produce systematically different outcomes across demographic groups.

Consider a lender using an automated model to evaluate applications.

Under the original law, the lender would have needed internal procedures intended to detect discriminatory patterns before those patterns harmed consumers. It might have been required to test whether similarly situated applicants received different outcomes depending on race, sex or age.

Under the replacement law, the lender must tell an unsuccessful applicant that automated technology was involved and explain its role. But the lender is not necessarily required to audit the model that produced the decision.

The burden therefore shifts.

Under the old framework, responsibility rested primarily with the institution deploying the system. Under the new one, the affected individual must recognize that something may have gone wrong, request information and pursue a correction or review.

That model assumes consumers have the time, knowledge and resources to challenge automated decisions. Many will not.

One constitutional problem disappears

Colorado’s rewrite followed a federal lawsuit filed in April 2026 by xAI, the artificial intelligence company founded by Elon Musk. The company sought to prevent the original law from taking effect.

Days later, the U.S. Department of Justice intervened in support of xAI. The intervention marked an important escalation: The federal government was no longer merely criticizing state AI regulation but actively seeking to invalidate it.

The Justice Department advanced two major constitutional objections.

The first involved equal protection.

The original law’s antidiscrimination framework required companies to evaluate whether AI systems produced different results for different demographic groups. Where unjustified disparities appeared, companies could be expected to modify their practices.

The Justice Department argued that this structure effectively compelled developers and deployers to make race- or sex-conscious decisions about model performance. Government classifications involving protected characteristics can trigger heightened judicial scrutiny, requiring the state to offer a strong justification and show that its approach is appropriately tailored.

Colorado’s rewrite largely eliminates that argument.

The new statute no longer requires companies to monitor outputs for demographic disparities or adjust systems in response to discriminatory patterns. The provisions at the center of the equal protection challenge are gone.

In that respect, lawmakers clearly reduced the state’s legal exposure.

But they may have intensified the second challenge.

The law is now almost entirely compelled speech

The Justice Department’s other argument rests on the First Amendment doctrine of compelled speech.

The First Amendment protects not only the right to speak but, in many circumstances, the right not to speak. When the government requires a private person or company to deliver a prescribed message, courts examine the nature of the mandate and the government’s justification for imposing it.

Not every disclosure requirement is unconstitutional. Governments routinely require businesses to provide factual information about products, financial terms, health risks and commercial practices.

But the constitutional analysis can become more complicated when disclosure obligations are extensive, controversial or insufficiently connected to preventing consumer deception.

xAI argued that Colorado’s notice, reporting and explanation requirements forced companies to communicate government-mandated messages about their technology.

The original law contained those requirements, but they existed within a broader regulatory framework. The statute also imposed risk-management duties, impact assessments and substantive obligations to prevent harm.

The replacement law retains little of that structure. Its central mechanism is now

communication.

Companies must notify consumers that automated technology was used. They must describe how it contributed to a decision. They must provide information enabling individuals to challenge data or seek review.

As a result, the compelled-speech question is no longer one issue among many. It is close to the whole case.

By narrowing the law, Colorado may have made the First Amendment challenge easier for a court to isolate. The state can no longer argue as convincingly that disclosure is merely an incidental feature of a comprehensive system for preventing algorithmic harm. Disclosure is the system.

The state may still prevail. Courts have upheld many factual commercial disclosure requirements, particularly when they are designed to prevent deception or protect consumers.

But Colorado will need to explain why these specific mandates are constitutionally permissible, sufficiently precise and not unduly burdensome. The answer will matter well beyond Colorado.

A federal campaign against state AI regulation

The Justice Department’s intervention did not emerge in isolation.

It followed Executive Order 14365, signed by President Donald Trump in December 2025. The order called for a “minimally burdensome national policy framework” for AI and directed federal agencies to scrutinize state laws that conflict with that approach.

It instructed the Justice Department to challenge state AI laws considered inconsistent with federal policy. It also directed the Commerce Department to evaluate existing state laws and identify those deemed overly burdensome.

Notably, the order called attention to state-mandated disclosures that might violate the First Amendment.

The Commerce Department’s evaluation was due in March 2026 but had not appeared by the time Colorado rewrote its law. Its eventual release could reveal whether the administration views Colorado as an isolated dispute or the opening case in a broader campaign against state regulation.

That distinction is crucial.

Dozens of states have introduced AI legislation touching on employment, housing, lending, insurance, education and healthcare. Many of those proposals rely on the same tools Colorado has used: notices, impact assessments, risk-management plans, human review and explanations for automated decisions.

A successful constitutional challenge could therefore reshape AI regulation nationwide.

States may respond by weakening their laws before courts ever rule. They may abandon substantive duties and rely almost entirely on consumer disclosure. Or they may hesitate to legislate at all, waiting for a federal framework that may not arrive soon.

Transparency without accountability

For Coloradans, the immediate practical effect remains limited.

A federal judge stayed enforcement of the original law in April 2026, and the stay extends to the replacement statute. None of the new requirements can currently be enforced.

The stay will remain in place until at least 14 days after the court rules on xAI’s expected request for a preliminary injunction. That request is tied to the completion of the state’s rulemaking process, leaving the timeline dependent on the attorney general, the courts and the Justice Department’s willingness to continue its challenge.

But the larger policy question should not be lost in the procedural uncertainty.

Transparency is useful. People should know when an automated system influences a consequential decision. They should be able to correct false data and request human review.

Yet a disclosure delivered after a harmful decision is not equivalent to a safeguard designed to prevent that decision.

Colorado’s original law placed responsibility on institutions to examine the systems they chose to deploy. Its replacement places far more responsibility on individuals to identify and contest possible harm after it occurs.

That may be politically easier and legally safer in some respects. But it is also a diminished vision of accountability.

Colorado began by asking how government could prevent algorithmic discrimination. It is now asking how companies should explain automated decisions once they have already been made.

Whether even that narrower approach survives will help determine how much authority states retain to govern artificial intelligence — and whether the future of AI regulation will be

built around preventing harm or merely disclosing it.