

Insights, News & Events

COLORADO RELEASES PROPOSED RULEBOOK FOR AI-ASSISTED HIRING AND EMPLOYMENT DECISIONS: 8 THINGS EMPLOYERS NEED TO KNOW

Insights
Aug 14, 2026

Colorado Releases Proposed Rulebook for AI-Assisted Hiring and Employment Decisions: 8 Things Employers Need to Know

Colorado regulators just released a draft manual to govern how employers use artificial intelligence to make decisions about hiring, promotion, pay, and discipline. The Colorado Attorney General's office published proposed rules on August 11 that cover the new automated decision-making technology (ADMT) law set to take effect on January 1, 2027. But they go well beyond the statute, spelling out what you must tell a rejected applicant, who is qualified to conduct the human review that applicants can demand, and what records you need to keep to prove any of it happened. Regulators are also asking the public to help them choose between two competing tests for when an AI tool is regulated at all, and that choice will determine whether some of the most common recruiting tools land inside or outside the law. The comment window is open now, and October 5 is the practical deadline if you want your input reflected in the draft that regulators will consider at the October 26 hearing. Here are the eight things employers need to know, plus your chance to shape the final version.

⚡ Quick Refresher

Colorado spent two years deciding how it would regulate AI in the workplace. After delays, false starts, court challenges, and multiple amendments, state lawmakers

Related People



Michael R. Greco

Regional Managing Partner

303.218.3655



Usama Kahf, CIPP/US

Partner

949.798.2118

and Governor Polis landed on the ADMT Act. You can read [our summary of that law here](#). In sum, it's a transparency-and-notice framework requiring disclosures, data access, and human review. If you built a compliance program to the 2024 version of this law, note what came out. The impact assessment and risk management program regime is gone. Note also that the Colorado Privacy Act's profiling opt-out does not reach employment-context data. This law does.

1. Biggest Open Question: Is Your AI Tool Covered At All?

The ADMT Act only applies when a covered tool is used to “materially influence” a consequential decision, and certain employment decisions like hiring and firing are squarely on the list of consequential decisions. Everything else in the law turns on that phrase – but the AG has not yet decided what it means. Instead, regulators put two competing standards in the notice and asked the public which it should adopt.

- **Standard 1:** Under this version, an AI recommendation that even just nudges a decision is likely covered even if a human weighed several other things more heavily. It says that “de minimis factors” to be excluded from coverage are only those that are trifling, trivial, or have an incidental impact on the outcome.
- **Standard 2:** This version, on the other hand, gives employers room to argue that a human’s independent analysis drove the decision. As long as other independent factors played a significantly larger role than AI outputs, the AI tool would be “de minimis” and not covered by the law.

The gap between these two options is the single-most consequential open issue in the rulemaking for employers. Standard 1 would pull resume rankers, scoring tools, and interview assessments broadly into the law. Standard 2 would leave space to keep a tool outside the law if a qualified human does real work on top of it.

Given the breadth of Standard 1, you should consider planning for that approach while using the comment process to explain the operational and policy reasons supporting Standard 2.

2. If the AI Tool Ranks or Scores People, Expect it to be Covered



Daniel Pepper, CIPP/US

Partner

303.218.3661

Service Focus

AI, Data, and Analytics

Resource Hubs

AI Governance Hub

Related Offices

Denver

Regardless of which proposed standard is adopted, the draft rules say that an AI output is presumed to have materially influenced a decision if it constrains an option set, sets a threshold, or produces a rank, score, classification, recommendation, prediction, or other inference that (a) pertains to the individual about whom a decision is being made, (b) is reviewed by the decision-maker or used to screen the data the decision-maker sees, and (c) is consistent with the final outcome.

That description fits most applicant tracking system screening tools, resume ranking software, and structured interview scoring products. If a recruiter looks at a ranked list and hires from the top of it, the presumption would apply.

One helpful carve-out survives in both versions: tools used solely to summarize, organize, translate, draft, route, or present information for human review would be excluded from the definition of ADMT altogether.

3. Ignorance Won't Save Employers

Under the proposed rules, an AI deployer would not be in compliance with the adverse outcome disclosure requirement if it cannot explain how the AI materially influenced a decision, how the tool used a person's data, or the principal reasons the tool produced an adverse result.

As a practical consequence, if your vendor cannot produce the principal reasons behind an individual output, the tool will not be usable for Colorado employment decisions once the rules take effect.

4. Your Rejection Notices Would Need to Say Far More Than They Do Today

The proposed rules specify the timing, form, and content of the post-adverse outcome disclosure in great detail.

- **Timing:** You would have 30 days after the decision to provide the disclosure to the affected individual.
- **Form:** The notice would go out in writing through at least two channels when you have two ways to reach the person, in at least 12-point font for anything mailed, and in plain language.
- **Content:** The disclosure would have to describe the decision, the role the AI played, the role of any human

reviewers, and the principal reasons for the outcome with specificity. Reasons based on inferences would have to identify the inference and the underlying data used to generate it. If a score drove the outcome, you would disclose the actual score. If a data point triggered an automatic denial, you would say so. If the person was screened out for incomplete information, you would tell them exactly what was missing.

The proposal includes two employment examples. In the first, an employer uses AI to select 10 interviewees out of 87 applicants and relies on consumer report data, so its FCRA adverse action notices satisfy the rules (as long as they carry the extra information about consumer rights and how to request more detail). In the second, an employer uses AI to make promotion recommendations that two humans review, and still owes a full disclosure naming the performance evaluation scores and technical assessment score that drove the outcome.

Two cautions on the first example. It works only if the vendor is in fact furnishing a consumer report, and many AI assessment vendors dispute consumer reporting agency status, so do not assume the FCRA path is open to you. The more likely failure is an employer assumes its existing FCRA adverse action notice already satisfies these rules and never adds the incremental content the proposal requires.

5. Human Review Would Need to be Truly Independent

The right to meaningful human review is the most significant obligation for employers. But applicants don't have the right to opt out of your use of AI on the front end of the hiring process, and can't force you to have a human review their file instead of an algorithm. The right attaches only after a person has already received an adverse outcome from a decision your AI tool materially influenced.

Under the proposal, the reviewer must be someone who did not make the original decision and is not a subordinate of the original decision-maker whenever feasible. The reviewer needs subject matter understanding proportionate to the harm at stake, documented training on objective decision-making and on the factors the AI considered, authority that is not subject to steering by upper management, and protection from retaliation. And the rules say plainly that AI may not assist in the review.

- You would need to tell people the right exists. Your adverse outcome disclosure must explain these rights and include a clearly labeled link going directly to instructions for using them, along with a mailing address or toll-free number. If someone asks for instructions by phone or mail, you would have 24 hours to provide them.
- Your request mechanism would need to let the person explain why they want review, identify the specific information they are contesting, and submit additional evidence.
- You would confirm receipt of a review request within 10 days.
- You would also need to complete the review and respond within 45 days.
- Where possible, you would need to delay the adverse outcome pending this review, so one review request can hold a requisition open for roughly two months. **This is the most operationally disruptive requirement in the proposal.**
- The response must state the reviewer's decision, the type of review conducted, the factors considered, and reasons specific to the evidence the person submitted rather than a recitation of how the tool generally works.
- You would also need to retain a file for every review showing reviewer identity, authority, and training, timestamps, the evidence available to the reviewer, whether the reviewer approved, modified, or overrode the output, and a written justification.
- If you argue that human review is not commercially reasonable, you carry the burden of proving it with specific evidence.

6. You Would Have to Let People Correct the Data Behind the Decision

Separate from human review, the ADMT Act gives individuals the right to request the personal data an ADMT used about them and to correct what is inaccurate. That is a second intake channel and a second workflow, and the proposed rules do not fold it into the review process. A rejected applicant can ask what data your screening tool relied on, tell you a field is wrong, and require you to fix it. If

the corrected data changes the picture, you are back in review.

7. Your Recruiting Chatbot May Be Regulated Separately

This is a consolidated rulemaking. The same August 11 filing also proposes rules under the Chatbot Safety Act (HB 26-1263), which Governor Polis signed on July 1 and which takes effect the same day as the ADMT Act.

That law requires operators of conversational AI services to estimate user age, disclose that the user is interacting with AI rather than a human, safeguard teen users against sexually explicit content and simulated emotional dependence, maintain suicide and self-harm response protocols, refrain from presenting outputs as equivalent to licensed professional services, and file an annual report with the Attorney General containing whatever metrics the AG decides are necessary to judge whether those safeguards work.

Employers using candidate-facing chatbots to screen applicants, answer questions, and book interviews will need to pay attention to this new law. Those employers taking applications from 16- and 17-year-olds through those same chatbots (especially in retail, hospitality, food service, and grocery) will be especially impacted. This rulemaking will settle whether a hiring chatbot is a covered conversational AI service. If your recruiting funnel opens with a chatbot, read the proposal for both statutes.

8. Staffing Agencies and Vendors Left in Limbo (For Now)

The AG flagged that the statute does not address what happens when a deployer relies on AI operated by a third party, and asked for input on whether to write rules allocating responsibility. One question is directly on point for employers: when a staffing agency operates the AI and the employer relies on the output, which party owes the deployer obligations, and does the answer change when the staffing agency is the one talking to candidates?

If you use staffing firms, recruitment process outsourcing providers, or third-party assessment vendors for Colorado roles, the final answer will determine whether you are the one writing adverse outcome notices and staffing a human review function.

There is a second allocation problem that has nothing to do with vendors: you may be a developer as well as a deployer. An employer that fine-tunes a model, sets its own scoring thresholds, or trains a tool on its own workforce data can land on the developer side of the line, which carries separate and heavier documentation duties. The configuration choices you think of as ordinary implementation are the ones most likely to move you across it.

Upcoming Dates to Track

- **August 11:** Proposed rules filed and comment period opened.
- **Late September:** The AG will post any revised proposed rules and send them to the rulemaking mailing list.
- **October 5:** Submit written comments by this date to be considered for revisions presented at the hearing.
- **October 26:** Public hearing at 10:00 AM, held in person at 1300 Broadway in Denver and by videoconference. Written comments close at 11:59 PM MT the same day.
- **January 1, 2027:** Both the law and the rules take effect.

What Should Employers Do Now?

Here are the steps you should take now as we wait for the final rules to come together.

1. Submit comments and consider doing it as part of a coalition. The material influence question, the vendor and staffing agency allocation question, and the human review requirements are all genuinely up in the air. Employers who explain the operational reality now have a real chance of shaping the final text. Our **FP Gov team** works with clients to draft and submit regulatory comments and to coordinate industry coalitions, and we can help you develop and file comments before the October 5 date. Reach out to your FP attorney or any member of our team to get started. (One caution before you file a comment. Every written comment is posted publicly at comments.coag.gov and becomes part of

the official rulemaking record. A comment that spells out how your screening tool works, how often your reviewers override it, or why meaningful human review is not commercially reasonable in your operation is discoverable and quotable. That is a strong argument for filing through a trade association or coalition, which puts the operational evidence in front of the AG without tying it to a single employer.)

2. Inventory your AI tools against the presumption. Pull together every AI-assisted tool that touches hiring, screening, interviewing, performance evaluation, promotion, compensation, or discipline for Colorado workers. For each one, ask whether it produces a rank, score, classification, recommendation, or inference about an individual that a decision-maker sees, and whether decisions tend to match the output. If the answer is yes, plan on being covered.

3. Decide whether your human review is sufficient, and build the paper trail. Under either proposed standard, the difference between a covered tool and an uncovered one will rest on the quality and documentation of the human judgment you use. Record when decision-makers form judgments, what independent information they review, and what authority they have to override.

4. Draft the adverse outcome notice now. Take one recent AI-assisted rejection and try writing the disclosure the rules would require, including the principal reasons with specificity and the underlying data behind any inference or score. Where you cannot answer, you have found either the need for a vendor conversation or a tool you should not be using for Colorado decisions. Work with your FP attorney if you need assistance or a final review.

5. Ask your vendors the explainability questions in writing. Request the documentation the law requires developers to provide, and ask specifically whether the vendor can identify the principal factors behind an individual output. Add contract terms requiring that support, and press staffing and assessment vendors on who will carry the deployer obligations.

6. Figure out who conducts human review. You need reviewers who are independent of the original decision-maker, trained, qualified in the subject matter, and empowered to reverse decisions, along with a 45-day turnaround and a records retention practice.

Conclusion

We will monitor developments and provide updates as warranted, so make sure you are subscribed to [Fisher Phillips' Insight System](#) to gather the most up-to-date information. If you have questions, please contact your Fisher Phillips attorney, the authors of this Insight, any attorney in our [Denver office](#), or any member of our [AI, Data, and Analytics Practice Group](#) or [Government Relations team](#).