Seyfarth

LEGAL UPDATE Aug 21, 2026

Legal Update: Colorado Releases Proposed Rules for Its AI and Chatbot Safety Laws: These Create More Operational Work than the Statutes Suggest

---

On August 11, 2026, the Colorado Department of Law released a single set of proposed rules substantially building out two 2026 statutes: the **Automated Decision-Making Technology Act** ("ADMT Act") and the **Conversational Artificial Intelligence Service Operator Requirements** (the "Chatbot Safety Act"). Both laws take effect January 1, 2027, and the proposed rules would become effective the same day. **The Rules are not yet final – for covered organizations and interested parties the weeks until September 4 are the time to submit comments to be considered for a redraft – and there is a specific ask to help shape the definition of covered ADMT.**

The Automated Decision-Making Technology & Conversational Artificial Intelligence Services Rules ("Rules") clarify open terms and add operational obligations, and they signal that Colorado's revised AI framework may demand significantly more compliance infrastructure than the statutes alone suggest.

The **ADMT Act** governs automated decision-making technology used to *materially influence consequential decisions*; specifically, those affecting a consumer's access to, eligibility for, or terms of education, employment, the lease or purchase of residential real estate in Colorado, financial or lending services, insurance, health-care services, and essential government services or public benefits. In the employment context, "consumer" reaches Colorado-resident employees and job applicants.

The **Chatbot Safety Act** is narrower in scope but operationally demanding for the operators it covers. It targets general-purpose, consumer-facing conversational AI services offered to the public, imposing age-estimation, disclosure, minor-protection, crisis-response, and annual-reporting duties. Notably, rulemaking is mandatory for the ADMT Act, but Chatbot Safety Act rulemaking is discretionary. Still, the Attorney General nonetheless chose to address in the same package to clarify operator obligations, including the content of the required annual report.

Key ADMT definitions

- **ADMT:** Technology that processes personal data and uses computation to generate outputs – such as predictions, recommendations, classifications, rankings, or scores – that are used to make, guide, or assist decisions about individuals.
- **Covered ADMT:** ADMT used to *materially influence* a *consequential decision*.
- **Excluded functions:** The proposed rules carve out ordinary infrastructure and certain low-risk functions, including routing, translation, summarization, scheduling, customer-service triage, advertising, marketing, search, and content moderation.

The most important unresolved issue is *when* ADMT "*materially influences*" a decision. The proposed rules do not adopt a single definition of the term. Instead, the Department proposes two possible standards for determining whether an ADMT output is more than de minimis and asks the public which to adopt.

Both versions include a rebuttable presumption of material influence where the output 1) constrains the available options, sets a threshold, or produces a rank, score, classification, or inference that relates to the individual; and 2) is reviewed by the decision-maker or used to screen the data the decision-maker sees, and is consistent with the outcome. Similarly, either formulation could require organizations to preserve evidence that a human decision-maker exercised independent judgment or relied on other documented factors.

Adverse outcome notices (ADMT Act)

The Rules would require deployers to provide detailed disclosures within 30 days after a covered ADMT materially influences an adverse outcome. The notice must identify the decision, explain the purpose of the ADMT, describe the roles of the ADMT and any human reviewers, and state the principal reasons for the outcome with specificity. Generic references to internal policies would not suffice.

The rules also require additional explanation when a decision relies on an inference, profile, risk score, automatic-denial factor, or incomplete information. This could prove difficult in applicant-screening, lending, housing, insurance, and other contexts where systems automatically screen out incomplete or disqualifying applications. Organizations may need decision-level explainability and data lineage that they do not currently receive from vendors.

Consumer rights and request handling (ADMT Act)

The Rules would require deployers to make detailed information about the ADMT available to consumers. This includes the system name and version, the developer, and the types, categories, and sources of personal data used. If data was received through an intermediary, the deployer may need to identify both the intermediary and the original source. The requirement to name the original source behind brokered data may be particularly difficult where deployers receive information through aggregators or data brokers; and standard vendor representations that data was collected in compliance with applicable law do not answer the question. This level of traceability may require new vendor commitments and internal recordkeeping processes.

The timing requirements are equally demanding. Some information must be available immediately through designated channels, while other requests must be handled within short regulatory windows. Consumers may obtain the personal data used in the decision, including ranks, scores, classifications, recommendations, predictions, inferences, or other ADMT inputs relating to them. Information must be presented in a way that lets the consumer decide whether to seek correction or human review, rather than in unexplained internal codes.

Meaningful human review (ADMT Act)

The proposed rules would make human review a staffed, documented process. Where feasible, the reviewer must be independent of the original decision-maker and not that person’s subordinate, appropriately trained, authorized to change the outcome, and able to conduct the review without ADMT assistance. The statute allows review only to the extent commercially reasonable, but the proposed rules identify factors for assessing commercial reasonableness and create a presumption in favor of review where the adverse outcome involves the severe and irreversible denial of a basic human need. The rules instruct deployers to weigh the type of review required, the magnitude and reversibility of the harm, the value of reviewing available primary evidence, the deployer's size and capacity, the cost and technical feasibility of review, and the availability of qualified reviewers, with no single factor controlling.

Deployers would need to acknowledge human-review requests within 10 days and complete the review within 45 days. Where possible, they must stay the adverse outcome while review is pending and while incorrect personal data is being corrected. They must also document the reviewer's qualifications, the evidence considered, the disposition, and a written justification.

Chatbot Safety Act obligations

The Chatbot Safety Act is narrower but operationally significant for covered operators. Internal-only workforce deployments behind authentication appear to fall outside the general-public scope; and narrow, task-specific bots may remain outside the Act where they are limited to bounded functions. Covered consumer-facing conversational AI services would face age-assurance, disclosure, privacy, minor-protection, crisis-response, and reporting requirements.

The proposed rules would require age-assurance methods that go beyond self-declaration, prohibit government-issued identification from serving as the sole method, and require reassessment when new signals indicate a different likely age. Operators would also need to report metrics on age distributions, age-estimation methods, crisis-referral outcomes, resolution times, and changes in age determinations. Meeting these requirements may force product, engineering, privacy, and trust-and-safety teams to build measurement and reporting systems before the effective date.

Practical implications

The proposed rules would require far more than updated notices. Organizations using ADMT in covered domains should expect compliance work extending well beyond legal and privacy teams. Participation from product, HR, risk, compliance, data governance, security, and vendor management will be necessary to have an effective compliance posture.

Organizations using ADMT in covered domains should inventory covered use cases, map vendor systems and data inputs, confirm whether outputs materially influence decisions, build adverse-outcome notice templates, create consumer-request workflows, negotiate vendor support for decision-level explanations, and designate trained human reviewers. Covered chatbot operators may need to assess whether their services are within scope and, if so, build age-assurance, user-safeguard, and reporting capabilities.

Comment period is open…speak now or forever hold your peace

The Department's comment portal is open from August 11 through October 26, 2026, with a public rulemaking hearing scheduled for October 26. The more immediate practical deadline is September 4, because comments submitted by then are expected to be considered for a revised draft; any interim updates to the proposed rules are expected to be posted by September 23. Organizations should consider commenting as the proposed rules require information that vendors may not provide, impose potentially impractical timelines, or create operational assumptions inconsistent with existing business processes.

Bottom line

Colorado's revised AI law was expected to be a narrower disclosure and consumer-rights framework. The proposed rules suggest that compliance may still require significant infrastructure, including decision-level explainability, data-source traceability, staffed human-review processes, and chatbot safety reporting. Companies using ADMT or consumer-facing conversational AI should use the comment period to assess compliance gaps, engage vendors, and identify requirements may be difficult or impossible to implement as drafted.

Authors

Yana
Komsitsky

Senior Counsel

Daniel E.
Riley

Associate

Related Practices

Artificial Intelligence

DATA Law: Digital Asset & Technology Advocacy