U.S. Chamber of Commerce

1615 H Street, NW
Washington, DC 20062-2000
uschamber.com

September 4, 2026

The Honorable Phil Weiser
Colorado Department of Law
Office of the Attorney General
Ralph L. Carr Colorado Judicial Center
1300 Broadway, 10th Floor
Denver, Colorado 80203

Re: Comments on Proposed Rules Implementing the Automated Decision-Making Technology Act (SB 26-189) and the Conversational Artificial Intelligence Services Act (HB 26-1263), 4 CCR 904-6

Dear Attorney General Weiser and Staff of the Colorado Department of Law:

The U.S. Chamber of Commerce (the "Chamber") respectfully submits these comments in response to the Colorado Department of Law's (the "Department") proposed rules implementing SB 26-189, the "Automated Decision-Making Technology in Consequential Decisions Act" ("ADMT Act"), and HB 26-1263, the "Conversational Artificial Intelligence Services Act" ("Chatbot Safety Act"), published August 11, 2026 (4 CCR 904-6). The Chamber is the world's largest business federation, representing the interests of more than three million businesses of all sizes, sectors, and regions, as well as state and local chambers of commerce and industry associations.

The Chamber strongly supports the responsible development and deployment of artificial intelligence and automated decision-making technologies. We believe that well-designed regulations can protect consumers while enabling businesses to innovate and compete. However, as detailed below, several provisions of the proposed rules go beyond the statutory text, impose operationally infeasible requirements, and diverge from analogous frameworks in other jurisdictions in ways that will create unnecessary compliance burdens without commensurate consumer benefit. The Chamber urges the Department to revise the proposed rules to align more closely with the statutory standard, promote interoperability with existing federal and state frameworks, and provide practical guidance that enables efficient compliance. We also reserve the right to provide comments on additional issues after the September 4 deadline for submission to be included at the October hearing.

I. Overarching Concerns

A. Tailoring the Rules to Prevent Overbreadth

The proposed rules, as currently drafted, risk sweeping in a broad range of ordinary business activity that the statute was not designed to reach. The definitions of "ADMT," "Materially Influence," and "Consequential Decision" are broader than analogous definitions in the California Consumer Privacy Act ("CCPA") ADMT Regulations, the Connecticut ADMT Law, the EU General Data Protection Regulation ("GDPR"), and the automated profiling provisions of dozens of other state consumer privacy laws. The Chamber recommends that the Department issue regulations that:

- Establish clear, objective indicators and presumptions for when an ADMT output materially influences a decision and is not incidental, trivial, or clerical;
- Confirm that general-purpose AI tools used for summarization, translation, drafting, or organizing information—where a human retains full decision authority—do not trigger the Act even if the output is referenced in the decision process;
- Clarify that routine risk-based pricing, fraud detection, credit underwriting, and compliance screening tools already subject to comprehensive federal regulatory oversight (e.g., FCRA, ECOA, GLBA) are not "consequential decisions" under the Act, or are otherwise excluded from its scope; and
- Provide safe harbors or presumptions of compliance for businesses that comply with analogous requirements under the CCPA ADMT Regulations, GDPR Article 22, or other recognized frameworks such as the NIST AI Risk Management Framework.

To promote consistent implementation, the Department should also consider providing illustrative examples of technologies that generally would not be viewed as materially influencing a consequential decision when operating as designed. Examples could include workflow management tools, information organization tools, deterministic systems implementing predefined human criteria, threshold-based eligibility checks, and similar technologies that automate administrative functions without exercising judgment regarding an individual's eligibility, access, or status. Such examples would improve regulatory certainty while preserving the flexibility necessary to address higher-risk uses.

B. Unintended Consequences

The Chamber is concerned that, without clarifying regulations, the proposed rules could produce several significant unintended consequences:

- Chilling innovation: Overbroad obligations may deter businesses from deploying beneficial AI tools, particularly in healthcare, financial services, and employment, where AI can improve outcomes and reduce human bias.
- Disproportionate impact on small businesses: Large enterprises may be better positioned to absorb compliance costs, while small businesses may be forced to abandon AI tools entirely or exit the Colorado market.
- Multi-state compliance fragmentation: Colorado's unique and broader-than-average definitions create a patchwork compliance environment that increases costs for businesses operating in multiple states.
- Disruption of routine employment operations: Ordinary workforce management tools—scheduling, performance tracking, attendance monitoring—could be swept into the Act's scope, creating significant operational disruption for employers.
- Duplicative burdens on regulated industries: Businesses in financial services and healthcare are already subject to comprehensive federal and state regulatory frameworks governing automated decision-making. Layering additional, potentially inconsistent obligations creates compliance complexity without commensurate consumer benefit.

C. Grace Period and Phased Implementation

The proposed rules take effect January 1, 2027—the same date as the statute—leaving no implementation runway for deployers to build compliance infrastructure after the rules are finalized. The Chamber urges the Department to either provide a reasonable compliance grace period following finalization of the rules or clarify that enforcement discretion will be exercised during an initial implementation window.

II. Definitions and Scope (Rules 2 and 4)

A. "Materially Influence" — Adopt the Substantial Factor Standard

The Chamber recommends that the Department adopt Material Influence Standard 2, which appropriately focuses on whether the ADMT output was a "substantial factor" in the Consequential Decision, rather than Standard 1's lower threshold of merely non-"trifling" impact. Standard 1 risks sweeping in virtually any ADMT output that a decision-maker reviews and happens to reach a result consistent with—even where the human independently arrived at the same conclusion based on separate evidence. Correlation between an ADMT output and a decision does not establish causation.

The Act requires both that an output be a "non-de minimis factor" and that the output "affect[] the outcome," including by "otherwise meaningfully altering" how the decision is made. C.R.S. § 6-1-1701(13). At minimum, the Department should finetune Standard 2's rebuttal mechanics to prevent overbreadth.

The Department should also provide additional guidance regarding how human involvement affects whether an ADMT output "materially influences" a consequential decision. The rules should identify objective indicators that demonstrate when a human decision-maker retained meaningful involvement in the decision-making process. Such indicators may include authority to review relevant information, consider additional factors, and approve, modify, or reject a recommendation before a final decision is made. This clarification would assist businesses in distinguishing between systems that effectively drive decision-making and systems that merely support human judgment.

B. "Consequential Decision" — Exclude Routine Operational Decisions

The Chamber recommends that the Department clarify that "consequential decision" does not include routine risk-based pricing, fraud detection, or compliance screening decisions already subject to comprehensive federal regulatory oversight. Rules should also clarify that decisions affecting large numbers of individuals in a uniform, non-individualized way (e.g., general product pricing, broad eligibility criteria) are not "consequential decisions" unless they involve individualized assessment of a specific person.

The Department should also provide examples to clarify the distinction between real employment opportunities and routine workplace administration and ordinary operational activities (e.g., scheduling, attendance administration, workplace-management functions) to help employers implement the law consistently.

III. Post-Adverse Outcome Disclosures (Rule 6)

A. Eliminate the Two-Channel Delivery Requirement (Rule 6.2)

Proposed Rule 6.2 would generally require adverse-outcome disclosures through at least two methods. For large employers processing thousands of applicants, this creates substantial cost and administrative burden with little corresponding benefit. The Act itself does not require multiple delivery channels. The final rule should permit a single reasonable communication method customarily used to communicate with the consumer.

B. Limit Initial Disclosures to Statutory Elements (Rule 6.4)

Rule 6.4's requirement to disclose "principal reason(s)" with specificity, individual scores/inferences, and underlying data goes beyond what § 6-1-1704(3)(a) expressly requires. The statute distinguishes between the mandatory initial disclosure (a "plain language description") and the separate request process for additional information under § 6-1-1704(3)(b). The rules should preserve that distinction.

Additionally, Rules 6.4(B)(11)–(12) effectively impose an explainability mandate. The provisions stating that a deployer "does not comply" when it cannot explain how the ADMT influenced a decision or cannot "accurately explain the principal reason(s)" functionally prohibit use of any system that cannot be reverse-engineered to produce individualized explanations—even where the statute's trade-secret protections should apply.

C. Named Data Source Disclosure (Rule 6.6) — Limit to Categories, Not Vendor Names

Rule 6.6(A)(3) requires deployers to identify each data source by name, including specific data brokers, databases, courts, social media companies, schools, and employers. This requirement is operationally burdensome and potentially exposes competitive vendor relationships. Many deployers may not even know all upstream sources, especially when using third-party AI systems. It may also pose information security risks to have this information in a publicly available setting.

The Chamber recommends that "sources" of data, for purposes of post-adverse outcome disclosures, be defined to mean categories of data sources (e.g., "consumer reporting agencies," "public records," "application data") rather than specific vendor names or proprietary data providers. This approach is consistent with the Colorado Privacy Act's established framework for consumer-facing data disclosures, which is built around "categories" of personal data—not "types" or "sources." Requiring disclosure of specific vendor names would expose trade secrets, create competitive harm, and provide little additional benefit to consumers seeking to understand the basis for an adverse decision.

D. Remove Rule 6.5(D)(1) Example to Avoid Confusion

The example in Rule 6.5(D)(1) of a bank using ADMT to refer a delinquent credit card to collections does not fit within the definition of "Financial or Lending Service" in Rule 2.2. Removal of this example would help avoid confusion regarding the scope of covered activities.

E. Extend Response Timeline for Consumer Access Requests (Rule 6.6(B))

Rule 6.6(B)(1) requires that a deployer respond to a consumer's access request within ten business days. This is not always operationally feasible and is inconsistent with California's 45-calendar-day requirement. If deployers have 30 days for the initial adverse-outcome disclosure, 30 days should likewise apply to supplemental information requests. Conflicting compliance requirements will create burdens on businesses and result in inconsistent implementations across jurisdictions.

F. Interoperability with ECOA and FCRA

The Chamber strongly recommends that the Department deem compliance for companies that comply with the adverse action notice requirements of the Equal Credit Opportunity Act ("ECOA") and the Fair Credit Reporting Act ("FCRA") for the Act's post-adverse outcome disclosure requirements for covered credit decisions. This would eliminate duplicative and potentially inconsistent disclosure obligations for financial services businesses and ensure that consumers receive consistent, familiar disclosures.

IV. Consumer Rights (Rule 7)

A. Preserve the Statutory Distinction Between Disclosure and Correction (Rule 7.3)

Rule 7.3's requirement that deployers disclose "specific pieces of Personal Data" including "final rank, score, classification, recommendation, prediction, or other inferences" risks conflating disclosure with the right to correct. The final rule must expressly preserve § 6-1-1705(1)(c)'s exclusion: disclosure of an ADMT-generated output should not expand the correction right beyond factually incorrect or materially inaccurate personal data. Applicants should not be able to demand "correction" of the model's evaluative judgment.

B. Adverse Outcome Stay Requirement (Rules 7.4(D) and 7.7(G)(1)) — Limit to Feasible Circumstances

The requirement to stay adverse outcomes pending correction of data or human review could be highly disruptive. For example, requiring a lender to hold a credit denial in abeyance or a landlord to delay a lease non-renewal while a 45-day review clock runs creates significant operational challenges. For high-volume decisioning processes, this requirement could significantly delay hiring timelines and other time-sensitive decisions. The "where possible" qualifier should be given meaningful effect, and the rules should clarify that the stay requirement does not apply where it would be operationally infeasible or where the decision is time-sensitive.

Additionally, Rules 7.4(B)–(D), combined with the stay requirement, could allow a consumer to submit repeated correction requests with new documentation each time, forcing the deployer into an indefinite cycle of re-processing and staying decisions. The rules should include a reasonable limit on sequential correction/reconsideration requests for the same decision.

C. Operational Feasibility: Response Timelines and Administrative Burden (Rule 7.2 and 7.7(G))

The 24-hour response requirement for consumer rights instructions in Rule 7.2 is unrealistic for physical mail. The Department should specify that this applies to electronic mail and the timeline should be extended to at least 7 business days. The Chamber also recommends that the webform heading required by Rule 7.2(D)(1) include the word "Colorado" to distinguish from other jurisdictions' requirements.

The 45-day window for completing Meaningful Human Review (Rule 7.7(G)(3)) provides no extension mechanism for complex cases requiring audit of the underlying ADMT system. The Colorado Privacy Act provides for extensions; the ADMT rules should as well.

D. Meaningful Human Review Standards (Rule 7.7) — Conform to the Statutory Standard

The Colorado ADMT Act gives consumers the right, in certain circumstances, to request human review of a consequential decision that was materially influenced by AI, including in the employment context. For example, this could apply if an employer uses AI as part of a hiring, promotion, or other employment-related decision covered by the statute. The statute requires only that the review be conducted by a trained individual with authority to override the decision.

Proposed Rule 7.7 significantly expands the statutory definition of Meaningful Human Review in ways that are operationally infeasible and untethered from the legislative text. The Chamber urges the Department to revise Rule 7.7 as follows:

- Remove the independent reviewer requirement (Rule 7.7(C)(1)). The requirement that the reviewer be "not a subordinate of the original decision-maker" and be "shielded from potential retaliation" does not appear in the statute and is more prescriptive than California's human review framework. This requirement will be difficult for small businesses with limited staff and may be impractical in highly specialized domains where few qualified reviewers exist. The Chamber urges the Department to remove these additional requirements and align the rule more closely with the statutory standard, which requires only that the reviewer be trained and have authority to override the decision.
- Remove the blanket prohibition on ADMT assistance (Rule 7.7(C)(5)). The rule flatly prohibits ADMT from assisting in the Meaningful Human Review process. For large-scale deployers handling thousands of adverse outcomes, this could make human review extremely costly and slow, particularly in sectors like insurance or financial services where AI-assisted review is standard practice. "Meaningful" should not mean "exclusively human." Reviewers should be able to use tools to assist their independent judgment.
- Correct the problematic implication regarding affirmation of decisions (Rule 7.7(D)(3)). The rule's statement that an override resulting in full reversal "indicates that human review was meaningful" creates a problematic implication that affirming a decision demonstrates the review was not

meaningful. The final rule should expressly state that reaching the same conclusion as the ADMT does not, by itself, indicate the review was not meaningful.

- Define "basic human need" and clarify that employment is not per se a basic human need (Rule 7.7(E)(3)). The "basic human need" presumption of commercial reasonableness is undefined and could sweep in all employment decisions. The term should be defined, and the rules should clarify that employment is not per se a "basic human need" for purposes of this presumption.
- Retain the statute's "to the extent commercially reasonable" limitation as a meaningful constraint, not merely a burden-shifting device. The rule should preserve the statute's flexible, principles-based approach to commercial reasonableness.

E. Record Retention for Meaningful Human Review (Rule 7.7(H)) — Litigation Risk

Rule 7.7(H) requires detailed documentation of every review, including reviewer identity, training, evidence considered, and written justification. While the statute requires three-year retention, the detailed nature of these records could create significant unnecessary exposure in future litigation, particularly given that the Act treats violations as deceptive trade practices. The Chamber urges the Department to limit documentation requirements to what is necessary to demonstrate compliance, consistent with the statute.

V. Chatbot Safety Act (Rules 8–13)

A. Definitions and Scope (Rule 8)

The Chamber recommends that the Department provide additional guidance and examples to clarify the definitions of "conversational artificial intelligence service" and "operator," ensuring these definitions are not interpreted too broadly or too narrowly. In particular, guidance should clarify what constitutes making a conversational AI service "publicly available," as this determination is central to the scope of the Act's obligations.

The Chamber also recommends that the Department clarify the conversational AI exclusion in the ADMT Act. The exclusion imposes a two-part conjunctive test, and the breadth of the first prong—particularly the inclusion of "configured" and "intended by a person"—creates substantial uncertainty for general-purpose conversational AI providers whose technology may be independently deployed by customers in consequential-decision domains. Rules should clarify that: (i) the acceptable use policy requirement is satisfied by a general prohibition in terms of service or acceptable use documentation without requiring active monitoring; and (ii) a deployer's independent configuration or intent does not retroactively defeat the developer's exclusion if the developer did not share that intent.

B. Age Assurance — Commercial Reasonableness (Rule 9)

The Chamber recommends that the Department provide guidance on what factors should be considered in determining whether an age-estimation method is "commercially reasonable," and how regulations should account for changes in technology, industry practices, or the availability of age-estimation tools over time. Rules should also clarify what evidence is relevant in determining whether operators "willfully disregard clear and convincing information" regarding the age of users and account holders, to ensure that this standard is applied consistently and does not impose strict liability on operators acting in good faith.

C. Minor Protections — Objective Indicators (Rule 11)

The Chamber recommends that the Department provide guidance on what factors regulators should consider when evaluating whether measures are "reasonable" or "technically feasible," and what evidence operators should maintain to demonstrate compliance. Rules should also provide objective indicators for evaluating whether a service is encouraging emotional dependence, and should distinguish between permissible categories of interaction (such as roleplay, companionship, and emotional support) and prohibited emotional dependence. Clear categories would help operators design compliant services and avoid over-restriction of beneficial uses.

D. Continuous Interaction Disclosure Every 3 Hours (Rule 10.2(D))

For chatbot operators, the requirement to re-disclose AI status every 3 hours in a continuous interaction (or maintain a persistent on-screen disclaimer) adds UX friction and may be technically complex for voice-based or embedded interfaces. The Chamber urges the Department to provide flexibility in how operators satisfy this requirement, including through persistent disclosure mechanisms that do not require repeated interruptions to the user experience.

E. WCAG Version Flexibility (Rule 3.2 / Rule 14.1)

Incorporating WCAG 2.2 by reference creates a static obligation that may become outdated or misaligned with evolving best practices. The rule should reference the then-current version or allow equivalent standards, consistent with the Chamber's recommendation in its July 2026 pre-rulemaking comments.

F. Annual Report Granularity (Rules 13.3–13.4)

The annual report requirements are extensive—including crisis referral accuracy rates, age distribution by service, detailed protocol descriptions, and specific metrics on minor-protection incidents. The verification provision (Rule 13.3(A)(9)) allows the AG to demand underlying documentation within 30 days, creating ongoing audit exposure. The Chamber urges the Department to calibrate these requirements to the size and capacity of the operator, consistent with the tiered operator-size framework established in Rule 13.3(A)(1), and to provide reasonable timelines for responding to verification requests.

VI. Additional Topics

A. Aggregate-Level ADMT Outputs

Rules should address whether ADMT outputs that operate at an aggregate level—without identifying or evaluating any specific individual—can "materially influence" an individual's consequential decision. For example, a labor-forecasting system that predicts "40 workers are needed for this shift" (without determining which individuals fill those slots) arguably "constrains" downstream assignment decisions about specific individuals. Rules should confirm that aggregate planning and forecasting tools that do not process individual-level personal data and do not identify, rank, or evaluate specific individuals are excluded, even if their outputs indirectly constrain the universe of individual decisions made by a human downstream.

B. Training Data Disclosure Specificity

The Act requires developers to disclose "a description of the categories of data, including personal data, used to train the covered ADMT, to the extent known." However, "data" is not defined, and the Act does not specify the level of detail required for these categories. Additional guidance on the expected specificity of this disclosure would be helpful, including whether developers may disclose data categories at a general level or must provide more granular descriptions. Clarification would help developers and deployers understand the scope of expected disclosures and support consistent implementation.

C. Pre-Use Consumer Notice Requirements

The Chamber recommends that the Department align pre-use notice requirements with the CCPA's flexible approach, which permits notice at or before the point of collection or use. Rules should permit notice at or before the point at which the ADMT is first used in connection with the consumer's specific interaction or transaction. The regulations should also adopt the CCPA's approach to consolidation of multiple disclosures for multiple purposes into a single ADMT notice, facilitating efficient compliance and promoting transparency to consumers.

VII. Conclusion

The Chamber appreciates the Department's commitment to a thoughtful, principle-guided rulemaking process and its stated goals of promoting consumer rights, clarifying ambiguities, facilitating efficient compliance, harmonizing with other frameworks, and allowing for innovation. We believe the recommendations set forth in these comments are consistent with and advance each of those principles.

The Chamber stands ready to work constructively with the Department throughout the rulemaking process and welcomes the opportunity to provide additional information or participate in any stakeholder engagement activities. We urge the Department to use its rulemaking authority to narrow and clarify the ADMT Act's scope, promote interoperability with existing federal and state frameworks, and provide practical guidance that enables businesses to comply efficiently while delivering meaningful protections to Colorado consumers.

Respectfully submitted,

Jordan Crenshaw

Jordan Crenshaw
Senior Vice President
Chamber Technology Engagement Center
U.S. Chamber of Commerce