FUTURE OF PRIVACY FORUM

1350 Eye Street NW, Suite 350, Washington, DC 20005 | 202-768-8950 | fpf.org

October 26, 2026

Via Electronic Submission

Attorney General Phil Weiser
Office of the Attorney General
Colorado Department of Law
1300 Broadway, 10th Floor
Denver, CO 80203

Re: Proposed Rules for the Automated Decision-Making Technology Act

Dear Attorney General Weiser and Staff,

Thank you for your continued work and the opportunity to provide input on the proposed rules concerning Senate Bill 26-189 (the "Colorado Automated Decision-making Act"). FPF is a global non-profit organization dedicated to advancing privacy leadership, scholarship, and principled data practices.[1] We seek to support balanced, informed public policy and equip regulators with the resources and tools needed to craft effective regulation.

Our comments draw upon recent work FPF has done in the AI governance and regulation space, including our *Updated Best Practices for AI and Workplace Assessment Technologies* ("Best Practices"). The Best Practices serve as a framework for understanding the AI value chain, assessing AI system risk, and applying responsible AI principles and practices when AI is used in consequential settings like the workplace. Building on this framework, we offer two recommendations to clarify the Department's interim draft rules and align them with current best practices:

1. **Scope:** The Department should simplify the considerations for assessing covered automated decisionmaking technologies to align with California and Colorado's existing Privacy Act Regulations; and

2. **Developer Disclosures:** The Department should further align developer disclosures with deployer obligations to consumers, consistent with each actor's role in the AI value chain

[1] The views reflected herein do not necessarily reflect the views of the FPF Advisory Board or FPF members.

1. ADMT Scope: The Department Should Simplify the Considerations for Assessing Covered ADMT to Align with California and Colorado's Existing Privacy Act Regulations

FPF recommends that the Department simplify the proposed definition of "De Minimis Factor" in Rule 2.2 by **removing the fifth and sixth enumerated considerations**. Doing so would promote interoperability with California's ADMT regulations and Colorado's existing Privacy Act regulations and facilitate greater operational clarity, while preserving a meaningful inquiry into whether ADMT materially influences a consequential decision.

In its pre-rulemaking comments, FPF urged the Department to align the ADM Act's framework for determining the scope of covered ADMT with the Colorado Privacy Act ("CPA") wherever feasible. As FPF explained, divergent standards can require regulated entities to conduct parallel compliance analyses, potentially producing inconsistent outcomes for functionally identical processing. FPF further identified personal data processing, the role of ADMT in the ultimate decision, the consequentiality of that decision, and the degree of independent judgment and meaningful human involvement as distinct components of ADMT scope.

The Department's proposed non-de-minimis framework should focus on that final component: the degree of independent judgment and meaningful human involvement in how a consequential decision was actually made. The first four factors address this directly–whether the decision-maker reviewed the output, analyzed independent analysis, had relevant expertise, and had authority to decide. These considerations are consistent with California Privacy Protection Agency's (CalPrivacy) California Consumer Privacy Act (CCPA) regulations and build on the Department's own definition of "human involved automated processing" in its Colorado Privacy Act regulations. Organizational or technical circumstances surrounding that decision**,** such as internal policies or the volume or complexity of data a decision-maker must review, in contrast, should serve as evidence of whether human review was meaningful, including when assessing rebuttal evidence under Rule 6.1(B).

A. The First Four Considerations Appropriately Focus the Inquiry on Meaningful Human Involvement and Promote Regulatory Interoperability

Under the October 6 interim draft, the Department proposes six considerations for assessing whether an ADMT output constitutes a de minimis factor in a consequential decision (in order to determine whether AI "materially influenced" a consequential decision, per the regulation's definition of Covered ADMT). The first four considerations explicitly contemplate the nature and extent of human involvement in the decision-making process, including (1) whether the decision-maker reviewed the ADMT output, (2) independently analyzed other relevant information, (3) possessed sufficient subject-matter understanding, and (4) retained authority to make or change the decision.

This approach is consistent with existing regulatory frameworks addressing meaningful human involvement in ADM. California's CCPA regulations define "human involvement" in terms of whether a human reviewer (1) understands how to interpret and use the technology's output, (2) reviews and analyzes the output alongside other relevant information, and (3) possesses authority to make or change the decision. Similarly, Colorado's existing Privacy Act regulations define "Human Involved Automated Processing" in terms of whether a human (1) meaningfully considers available information or processing outputs and (2) has authority to influence the outcome.

Table 1. Comparison of the First Four Proposed De Minimis Considerations with California's ADMT Regulations and Colorado's Privacy Act Regulations

First Four Non De-Minimis Considerations	**CA Analogue:** ***11 CCR § 7001(e)*** *Required Elements of "Human Involvement"*	**CO Analogue:** ***4 CCR 904-3, Rule 2.02*** *Required Elements of "Human-Involved Automated Processing"*
Reviewed the output before making the consequential decision and whether the final decision was consistent with that output	Review and analyze both the technology's output and any other information relevant to making or changing the decision	Engage in meaningful consideration of available data used in the processing or any output of the processing
Reviewed and analyzed information independent of the ADMT output	Review and analyze both the technology's output and any other information relevant to making or changing the decision	Engage in meaningful consideration of available data used in the processing or any output of the processing
Possessed the subject-matter understanding necessary to make the decision	Know how to interpret and use the technology's output to make the decision	n/a (no express analogue)
Had authority to make or change the consequential decision	Have authority to make or change the decision	Has authority to change or influence the outcome of the Processing

Although the three frameworks do not establish identical requirements, the substantial overlap among their human-involvement criteria reflect a common approach to evaluating meaningful human participation: the inquiry centers **the decision-maker's engagement with the technology and ability to exercise independent judgment**, focusing specifically on what the human decision-maker **actually did, understood, and was empowered to do**. Aligning the proposed non-de-minimis considerations with these existing frameworks would promote regulatory interoperability and enable organizations operating across jurisdictions to implement more consistent human-review procedures.

B. The Final Two Considerations Address Circumstances Surrounding Human Involvement Rather Than Human Involvement Itself

The fifth and sixth considerations, by contrast, depart from the approach described above. By introducing circumstances surrounding human involvement rather than focusing directly on the nature and extent of the involvement itself, these considerations complicate an otherwise administrable and operationalizable regulatory approach.

i. Organizational Policies Are Evidence of Human Involvement, Consistent with the Goals of Rule 6.1(B)

First, the fifth consideration concerning organizational policies and representations is more appropriate as supporting evidence in compliance documentation and enforcement investigations than as an independent consideration in the material-influence determination. The proposed language directs consideration of whether a deployer's marketing materials, procurement contracts, internal policies, or training materials describe or direct that ADMT output be used as a primary or determinative basis for a consequential decision.

Although such materials may provide relevant context regarding an organization's intended use of ADMT, they do not necessarily establish how a particular decision-maker actually relied on the technology's output. For example, an organization's procurement materials may describe an ADMT system as central to its decision-making workflow, while individual decision-makers nevertheless independently review relevant evidence and exercise meaningful discretion. Conversely, an organization may maintain policies requiring independent human review without ensuring that such review occurs in practice.

Moreover, the Department's proposed Rule 6.1(B) already provides an evidentiary mechanism through which relevant organizational materials can be considered when evaluating whether an ADMT output was de minimis. Rule 6.1(B) permits a deployer to rebut the presumption of material influence through evidence demonstrating that an ADMT output was a de minimis factor in a consequential decision. Specifically, Rule 6.1(B)(5) permits persuasive evidence, including contemporaneous evidence, demonstrating judgment independent of the ADMT output. Organizational policies and related documentation may serve this purpose by corroborating how decision-makers were instructed and empowered to exercise independent judgment.

Removing this consideration from the enumerated factors, and instead clarifying through Rule 6.1(B)(5) that organizational policies may constitute relevant evidence when assessing whether an ADMT output was a de minimis factor, would retain the focus on actual human decision-making while allowing appropriate consideration of the circumstances surrounding that decision-making. This distinction is consistent with the aforementioned *Updated Best Practices*, which emphasize the importance of establishing organizational governance policies while separately recognizing that meaningful human oversight requires decision-makers to possess the information, authority, capability, and genuine opportunity to exercise independent judgment.

ii. The Complexity of Personal Data Processing Is Duplicative of the Threshold ADMT Inquiry and Is Not a Reliable Measure of Human Involvement

Second, the sixth consideration concerning the volume, type, or complexity of personal data processing risks conflating the technical characteristics of ADMT with its degree of influence on a consequential decision. The proposed consideration asks whether ADMT processes personal data of a volume, type, or complexity that a decision-maker could not reasonably review or replicate independently within the available time.

However, the statutory definition of ADMT already incorporates the processing of personal data. Introducing the volume, type, or complexity of such processing as an additional consideration in the material-influence analysis risks creating an overlapping inquiry that focuses on the underlying technology rather than the role of its output in the ultimate decision.

Importantly, a decision-maker's inability to independently replicate an ADMT system's processing does not necessarily establish that the system's output materially influenced the decision. A human reviewer may be unable to reproduce a complex computational analysis while nevertheless possessing sufficient information, subject-matter expertise, and decision-making authority to reach an independent conclusion based on other relevant evidence.

FPF's *Updated Best Practices* similarly distinguishes the characteristics of an AI system, including the data it processes, from the proximity and influence of its outputs on consequential decisions. The framework identifies the influence of an AI system's output and the availability of meaningful human review as distinct considerations when evaluating how directly a system shapes a decision.

Retaining the sixth consideration could therefore introduce unnecessary ambiguity by suggesting that the complexity of personal data processing itself bears on whether ADMT materially influences a consequential decision, even where the decision-maker exercises meaningful independent judgment.

FPF therefore recommends that the Department **remove the fifth and sixth considerations from the proposed definition of "De Minimis Factor" in Rule 2.2.** The remaining four considerations provide a more direct and administrable framework for evaluating the actual role of ADMT outputs in consequential decisions, while promoting consistency with California's ADMT regulations and Colorado's existing Privacy Act regulations. This targeted revision would preserve the Department's ability to assess meaningful human involvement without introducing potentially duplicative or indirect considerations into the material-influence analysis.

2. **The Department Should Further Align Developer Disclosures to Deployer Obligations, Consistent with Each Actor's Role in the AI Value Chain**

Building on several improvements in the interim draft rules, the Department can further clarify the rules to ensure that what developers must provide aligns with what deployers must disclose and explain to consumers under the law. As detailed in our *Updated Best Practices*, "[e]ffective AI governance depends on a shared foundation: who bears responsibility, what risks are present, which practices apply and under what conditions."[2] Transparency is central to that foundation as it not only ensures that affected individuals receive meaningful and accurate information, but also equips deployers to use AI systems responsibly, without taking on risks they cannot identify, assess, or mitigate.

Our *Updated Best Practices* detail seven specific disclosures developers should make to deployers:

- ➔ **Intended purposes**: the purposes for which the tool is designed, and any uses that are restricted or prohibited;
- ➔ **Known limitations**: intended deployment conditions and any known limitations, including conditions under which performance may degrade, or populations or use cases for which the tool has not been validated;
- ➔ **Training**: a general description of how the tool was trained;
- ➔ **Testing performance:** the system-level safeguards, including testing procedures and performance metrics (e.g., likelihood of data leakage), to support informed deployment.
- ➔ **Risk/Impact assessment**: any measures taken by the Developer to assess and/or mitigate potential risks arising from the use of the AI System, such as reasonably foreseeable risks of unlawful discrimination;
- ➔ **Data practices:** whether the tool uses Deployer or Individual data to generate outputs or for model training or improvement, and under what conditions; and
- ➔ **Configuration and oversight guidance:** instructions for appropriate deployment, including any Developer recommendations regarding human oversight during operation.

FPF appreciates that the interim draft rules reflect much of this approach, however the interim draft rules do not yet fully connect what deployers must provide to consumers and adversely affected individuals under the law, and what developers must provide to deployers. Where developer information falls short, the rules allow deployers to disclose only "as much information and specificity as possible."[3] While this accommodates the gap, it does not close it. Consumers receive incomplete disclosures, and deployers bear compliance risk for information they cannot access.

---

[2] *Id.* at 5.

[3] Rule 7.4(B)(5): If the Deployer is unable to fulfill any of the obligations in section 6-1- 1704(3)(a), C.R.S., and this Rule because it has not received the necessary information from the Developer, its notice should contain as much information and specificity as possible based on the information that it has received from the Developer.

Recognizing that the ADMT Act, as amended by SB 26-189, establishes a narrower set of developer obligations than its predecessor, we do not recommend new categories of developer duties. Instead, we recommend that the Department clarify the terms already present in section 6-1-1702, C.R.S., in the following ways:

A. Tie developer documentation to deployer obligations. We recommend adding a general standard to Rule 5.2 requiring developers to provide documentation sufficient to enable deployers to meet their obligations under section 6-1-1704, C.R.S., and Rule 7, to the extent that information is within the developer's knowledge. At a minimum, this documentation should include the covered ADMT's name and version (not currently detailed in developer disclosures) and the categories of data the covered ADMT *generates* (not only the categories and types it uses and requires), including the types of scores, classifications, inferences, or other outputs it produces.

Such approach aligns with similar controller-processor frameworks in data privacy laws, including the Colorado Privacy Act, which requires processors to provide controllers with all documentation and information needed to prove compliance.[4]

B. Clarify the meaning of "known limitations" and "known risks." Rule 5.2(A) requires developers to detail "known limitations" and "known risks" without defining their content. Yet, deployers are accountable under Rules 3.2 and 7.4 for providing accurate disclosures and responding to consumer requests for explanations. If developers disclose only general limitations, deployers may be unable to explain why a system produced a particular outcome or affected certain groups differently, leaving them accountable for information outside their control. Consistent with the Best Practices, we recommend clarifying that "known limitations" and "known risks" include:

a. Conditions under which performance may degrade;
b. Populations or use cases for which the covered ADMT has not been validated;
c. The intended user population; and
d. Reasonably foreseeable risks the developer has identified through its own testing, including risks of unlawful discrimination, along with any measures taken to mitigate them.

This clarification would apply only to limitations and risks developers already know, and ensures that information reaches the deployers who need it to meet their obligations under the rules. In practice, it may also encourage developers to validate their systems for intended uses, consistent with established best practices. Absent this clarification, deployers bear the risk of gaps in information they cannot access. For example, a developer whose resume-screening tool was trained on data from a single industry should disclose that the tool has not been validated for other populations, so that

[4] See, C.R.S. § 6-1-1305(5)(d)(II)(A)

deployers can configure additional human review and describe the tool accurately to applicants.

C. **Clarify ongoing disclosure obligations.** Section 1702(2) requires developers to notify deployers of material updates and changes in intended use, but the interim draft rules do not elaborate on this obligation. We recommend clarifying that a "material update" includes changes to the covered ADMT's model, version, or intended uses that could reasonably affect its outputs or performance. We also recommend that developers notify deployers within a reasonable time of newly identified material risks or serious incidents discovered through post-deployment monitoring. Because Rule 7.6(A)(3) requires deployers to disclose the covered ADMT's version number, deployers cannot meet that requirement without timely notice of version changes.

D. **Use standardized formats:** As detailed in our Best Practices, we also recommend, where feasible, developers to use standardized documentation formats such as model cards and system cards to communicate technical information about AI Systems to Deployers.

* * *

Thank you for this opportunity to provide comments on these proposed regulations. FPF welcomes any further opportunities to provide resources or information to assist in this important effort. If you have any questions regarding these comments and recommendations, please contact Tatiana Rice at trice@fpf.org.

Sincerely,

Tatiana Rice, Senior Director of U.S. Legislation and Regulation
Michelle Roitgarts, U.S. Legislation Legal Intern