Bidun Group LLC
Atlanta, Georgia · mike@bidun.com · bidun.com

October 14, 2026

Colorado Department of Law
Office of the Attorney General
1300 Broadway, Denver, CO 80203
Submitted via the rulemaking comment portal (coag.gov/ai)

RE: Supplemental comment on the October 6, 2026 interim draft of the rules implementing the Automated Decision-Making Technology Act (SB 26-189) and the Chatbot Safety Act (HB 26-1263), 4 CCR 904-6

Submitted by: Michael Bidun, Founder, Bidun Group LLC, Atlanta, Georgia. This supplements my comment of August 28, 2026 (Comment 0000000544). My practice serves small regulated financial firms, and I have held governance, risk and compliance roles inside regulated financial institutions for more than twenty-five years. As before, this comment is offered from operating experience, not as legal advice, and any coverage or definitional point is an operational reading offered for the Department's clarification, not a legal conclusion.

I write on five points in the interim draft and on two of the questions in the Department's cover page.

1. Rule 8.7(C)(1)(a): the external reviewer. I welcome the interim draft's treatment of external reviewers. My August comment asked the Department to confirm that a qualified external reviewer under contract may satisfy the independence requirements for Meaningful Human Review, and the interim draft now states the three conditions that make outside review honest: the Deployer remains fully accountable for the review; the contract insulates the reviewer's decisions on individual reviews from the Deployer's influence; and the Deployer's oversight of the arrangement addresses administration and performance, not the outcome of any particular review. For a community bank or credit union whose compliance function is one person, this is the difference between a requirement that can be met and one that cannot. I urge the Department to keep this language through adoption.

2. The Department's question on Material Influence. The cover page asks what would let the parties determine that an ADMT output was a de minimis factor, "outside of the Deployer simply asserting that the output was a de minimis factor." My answer: keep the presumption, and define the rebuttal by the records it must rest on.

The interim draft is already most of the way there. Rule 6.1(B) lists the evidence that can rebut the presumption, including a consistently applied hard cut-off or threshold, an independent review of substantial information other than the ADMT output by a decision maker with authority to decide, and "other persuasive evidence, including contemporaneous evidence, showing judgment independent of the ADMT output." I ask the Department to make contemporaneity the organizing principle of the rebuttal rather than one item on the list: the weight of rebuttal evidence should follow how close to the decision it was made, so that records made at the time of the decision carry the most weight and records assembled afterwards carry the least. Three kinds of record do this work in the ordinary course of an honest process, and none requires a small deployer to build anything new. The decision maker's own judgment, recorded before the output was available to them. A dispositive factor or threshold documented in advance and applied consistently across like cases, which the draft already names. And the record of the decision maker's independent review of substantial other information, made when that

review was done, which Rule 6.1(B)(3) already contemplates; Rule 8.7(H), which now requires the Meaningful Human Review record to be contemporaneous and to show the reviewer's identity, authority, timestamps and the evidence available, is the model for what such a record looks like.

The rules should also say what carries little weight: a judgment first written down after a complaint arrives, or a factor labelled dispositive only after the fact, is an assertion with a date on it, and the Department should be free to treat it as one. This is not a rule against later records. Errors will be found, criteria will be refined, and laws will change; a later record that carries its own date and says what it corrects is the mark of a working process, and a decision should be judged by the standard that applied when it was made. What should carry no weight is a record whose own date and provenance do not show when it was made. The reason this answers the Department's question is simple: an assertion is words, while a record made at the time is an artifact an examiner can inspect and a small firm produces anyway. A standard whose honest application leaves inspectable artifacts is one the Department can examine and a small deployer can operationalize. To be clear about the limit of the ask, as in August: no label, program or framework should confer a presumption; only records that meet the rule's own content requirements, in substance, should count.

3. One decision, several systems. My August comment asked how the rules reach a consequential decision produced by several coordinated automated components, where an orchestrating system delegates steps to subordinate ones and none alone looks like the decision maker, and offered further detail on request. The interim draft adds a Midstream Developer definition for a party that integrates Covered ADMT as a component into its own product. That addresses the supply chain. It does not address a Deployer running several systems, from several developers, whose combined output is the decision, and deployments are already ahead of the rules on this point. The governing principle is that coverage follows the decision, not the architecture. Two mechanics would give it effect. First, in Rule 6, the presumption of Material Influence should attach to the integrated process that produces the Consequential Decision, aggregating the automated components that materially contribute to it, with the Deployer identifying those components; an ADMT output should not be treated as de minimis merely because another automated component, rather than a person, consumed it. The limit is material contribution, the same test the rules already apply to a single system: a component is part of the process when its output, directly or through another component, is an input the decision depends on; a component whose output never becomes such an input is outside it. Second, in Rule 8.7(H), where several systems contribute, the review record should list them, so that the reviewer and an examiner can see the whole chain; the draft's own disclosure provisions already contemplate "other ADMT or similar systems, if applicable." For components a Deployer licenses rather than builds, the Deployer should be able to rely on the Developer's documentation under Rule 5 for what the component does, as evidence rather than as a safe harbor, while remaining responsible for how it configured and used the component; the listing then asks a small deployer for what it can know, not for what only the vendor can.

4. Tools that assist Meaningful Human Review. I welcome the interim draft's clarification in Rule 8.7(C)(5)(a) that tools used solely to summarize, organize, translate, draft, route or present information are not ADMT and may assist the review. I ask the Department to name one more: a deterministic check that tests the file against human-authored, inspectable criteria, such as whether the figure the policy relies on is the figure in the file, or whether a required document is present. Two conditions should govern, as in my August comment: the reviewer evaluates the matter independently and remains free to depart from anything the tool reports; and the tool's criteria are human-authored and independent of the Covered ADMT's own logic, so that the original system's reasoning cannot be relabelled as

verification. The line between the two is what the output is about. A deterministic check reports on the file: whether each stated criterion was met, by how much a figure departs from what the policy requires, and how many criteria out of how many were satisfied, in numbers where numbers are the natural form. A tool whose output is about the consumer or the decision, one that ranks, scores, classifies or predicts the person, recommends an outcome, or re-runs the decision logic under review, is doing the review's work and should remain excluded. So is a tool whose criteria are the eligibility rules themselves, such that passing the check is the decision: a check verifies the facts a reviewer will weigh; it does not apply the test the reviewer is there to apply. A second AI system remains what it was: more automation, never the review. A reviewer at a small firm handling volume needs a checklist to be thorough, and the rule should not leave that reviewer wondering whether a checklist is ADMT.

5. Rule 8.2(B)(2): the instructions clock. The interim draft keeps the requirement that instructions requested by toll-free number or mail be provided within 24 hours of receipt. I ask for one business day. The request concerns only the sending of instructions, not any deadline for acting on the consumer's request itself; the consumer's substantive deadlines under Rule 8 are unchanged, and the cost to the consumer is at most a weekend's wait for a set of instructions. For a deployer with one compliance officer, a 24-hour clock that starts on a Friday evening demands a weekend response for a small incremental benefit; one business day serves the consumer nearly as well and is performable by the smallest deployers.

6. Rule 9: age assurance data. The interim draft keeps Rule 9's privacy design, including the rule that government-issued identification may not be the sole method of age assurance, and I support both. One refinement would extend the design. The rule provides that collected age-assurance data "should be deleted after the minimum time required for compliance." Where a chosen method does retain identity documents or biometric matches, I ask the Department to make that deletion mandatory, with an outer limit that runs from the completion of the assurance, because retained proofs accumulate into stores that attract attackers and normalize upload-your-identification prompts that phishing then imitates. A binding deletion clock does not remove collection or transmission risk, but it keeps the retention risk from growing with time, and it asks of a compliant operator what the rule's minimization principle already asks. The clock should yield to retention that law requires, and to a hold for a specific, documented legal, fraud or security matter for as long as that matter is open, with the hold recorded and reviewed on a stated cadence, never as a standing basis for retention; the Department is better placed than I am to set the outer limit's length, and for methods that assess age continuously the clock should run from each assessment. The rule should distinguish the record that an age assurance was performed, namely the method and its version, the date, the result and any third-party reference, which an operator may keep as evidence, from the identity documents and biometric data used to perform it, which should be deleted on the clock. A record of each method version, its settings, its test results and the dates it was in use, kept once rather than per user, lets the operator show that the check performed met the standard at the time, which is what Rule 9's testing requirement asks and what an examiner will want to see.

7. The Department's invitation on "Financial or Lending Service." The cover page says the Department is open to an alternative definition and invites proposals. I do not propose text. I offer one consideration from the operating side. The current draft definition, "the extension of credit or a loan, transmitting or exchanging funds, the provision of deposit or checking accounts, check cashing, or installment payment plans," reaches the core consumer services of community banks and credit unions, and it should continue to. The question is how it is written. Federal law already defines this ground: "financial product or service" in the Consumer Financial Protection Act, 12 U.S.C. 5481(15)(A), covers

extending credit and servicing loans, deposit-taking, transmitting or exchanging funds, and check cashing, in terms every bank and credit union already maps its activities to, and "credit" in Regulation B, 12 CFR 1002.2(j), is the term those firms apply to every lending decision. A Colorado definition anchored to those terms, by reference to the specific subparagraphs (for the current list, 5481(15)(A)(i), (iv) and (vi)) and to a stated edition so that the incorporation is static, lets a small firm reuse the compliance map it already maintains, and gives the Department a term with examiner practice behind it; a free-standing list, however similar today, will have to be reconciled against those terms by every deployer separately. For contrast, the EU AI Act's Annex III reaches this area by function, naming systems "used to evaluate the creditworthiness of natural persons or establish their credit score," with an exception for fraud detection; I note it only to show that the choice between a function-based and a service-based definition has been made differently elsewhere, and that whichever Colorado chooses, anchoring it to terms deployers already use will reduce the parallel structures my August comment described.

8. Worked examples in the chatbot rules. My August comment asked the Department to give the Chatbot Safety Act rules the same worked, sector-specific examples the ADMT rules use. The interim draft's chatbot rules (Rules 10 to 14) still carry none. I repeat the request, and suggest one of several: a financial institution's customer-service chatbot, which is the deployment most of my clients will run, walked through the disclosure, age-assurance and minor-user rules, including where those rules do not apply. Examples should be expressly illustrative, as the ADMT examples are, and drawn from more than one sector, so that no single example becomes a baseline.

I appreciate the Department's open process and the substance of the interim draft. Supporting detail on any of the above can be provided on request.

Michael Bidun
Founder, Bidun Group LLC, Atlanta, Georgia