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

August 28, 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: Proposed 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. My practice serves small regulated financial firms: community banks, credit unions, registered investment advisers, and small broker-dealers, together with the consumer-facing fintechs and the software firms that build for them. Rules like these reach any deployer whose covered activity or consumers are in Colorado, wherever the firm sits, and the Financial or Lending Service definition in Rule 2.2 places my clients' core activities squarely in scope. I have held governance, risk, and compliance roles inside regulated financial institutions for more than twenty-five years, and my work includes maintaining a line-by-line register of the duties an AI governance program asks a small firm to perform. This comment is offered from that operating experience, not as legal advice, and coverage or definitional points below are operational readings offered for the Department's clarification, not legal conclusions.

I write in support of the rulemaking's direction, with two commendations, three requests, and one question for the revised draft, together with an answer to the Notice's questions on vendor arrangements.

**First commendation: the communication standards in Rule 3.** The requirements that consumer communications use plain, straightforward language and avoid technical or legal jargon, and that online information follow recognized accessibility standards (WCAG 2.2, incorporated at Rule 14.1), together with Rule 6's requirement that principal reasons be stated with specificity, with the rules' own example that "internal standards or policies" is insufficient, are among the most consumer-protective choices in the draft. A notice, like a policy, protects no one who cannot read it. I urge the Department to keep these standards through revision.

**Second commendation: the combined-notice design in Rule 6.** The worked examples showing that a creditor's adverse-action notice under the federal Equal Credit Opportunity Act, or an employer's notice under the Fair Credit Reporting Act, satisfies the Adverse Outcome disclosure requirement without a separate ADMT disclosure, provided it carries the content Rule 6.4 requires, is exactly right for the firms I serve. The consumer receives one clear, complete letter instead of two overlapping ones. I urge the Department to preserve these examples through revision and, where practical, to extend the same principle to other overlapping notice regimes, always conditioned on the combined notice satisfying

every applicable requirement of each regime, including timing, delivery, and accessibility.

**First request: proportionality, extended beyond the one place it now lives.** Rule 7's commercially-reasonable assessment for Meaningful Human Review already names "the Deployer's size and capacity," "the marginal cost and technical feasibility of the review," and "the availability of qualified reviewers" among its weighed factors, and the Department deserves credit for that. Two refinements would make the proportionality real for small deployers. First, gloss "size and capacity" with illustrative, non-dispositive reference points a deployer's own regulators already use, such as total assets for depository institutions or assets under management for advisers; familiar measures make the factor administrable, while remaining factors among factors, since an asset-rich firm can still be operationally small. Second, the same proportionality logic should reach the rules' fixed clocks, which currently apply uniformly whatever the deployer's size: consider who deploys covered technology in consumer finance, a community bank or credit union whose compliance function is one person holding many responsibilities, with no model risk department and no staff counsel. For that firm, the requirement in Rule 7 that instructions requested by phone or mail be provided within 24 hours of receipt lands hardest when the clock runs across a weekend at a firm with one compliance officer; the incremental consumer benefit of 24 hours over one business day is small, and the difference in performability for the smallest deployers is large. The obligations in this draft, taken together with the AI governance duties these firms already carry under supervisory expectations, are individually reasonable and collectively enormous: the register I maintain currently counts more than 800 discrete duties across a full program, hundreds even after tailoring to a small firm, before this rule adds its own. That count measures what a program must track, not an objection to any single duty, and its methodology is available to the Department on request.

**Second request: recognize documented AI governance programs as the way these duties get performed, not as something beside them.** Rule 7's documentation requirements for Meaningful Human Review, retaining the reviewer's identity, authority, and training, the timestamps, and the evidence available to the reviewer, already embody the right principle: the record of performance is created at the time of performance. I ask the Department to extend that principle in two directions. Wherever a requirement's substance allows, state it in a form a deployer can evidence from records; in operating experience, duties that can be evidenced from records get performed far more reliably than duties that exist only as prose, which are the first casualties of a pressured week. And say plainly that a deployer operating a documented, risk-based AI governance program, with a systems inventory maintained for completeness and named approvers, risk-based review before deployment, impact assessment, human oversight with defined roles, and records built for examination, may satisfy the corresponding duties here through that program's artifacts, provided each artifact meets the rule's own content requirements. To be clear about the limit of the ask: no label, framework, or program membership should confer any presumption; only artifacts that meet the rule's content requirements, in substance, should count. Without that recognition, small deployers will build a second, parallel compliance structure for one statute, and parallel structures are where compliance quietly fails: two inventories reconciled by nobody, two review records that drift apart. With it, this rule strengthens the governance programs the banking agencies' model risk guidance and FINRA's oversight priorities are already pushing these same firms to build.

**Third request: two clarifications on Meaningful Human Review.** First, the independence requirements in Rule 7, a reviewer who did not make the original decision and is not the decision-maker's subordinate, free from steering by upper management and shielded from retaliation, are framed in internal terms, and at a small firm internal arrangements strain against them, while a qualified external reviewer under contract can satisfy them cleanly. I ask the Department to confirm, ideally by example, that an external reviewer may satisfy Rule 7's independence requirements when the appropriate safeguards are met: the deployer remains fully accountable for the review, the engagement's terms insulate the reviewer's decisions on individual reviews from the deployer's influence, and the deployer's vendor oversight of the arrangement addresses administration and performance, never the outcome of any particular review. Second, the provision stating that "ADMT may not assist in the Meaningful Human Review" is right in its aim and would benefit from a functional boundary. As written, it could be read to bar the reviewer from using any automated tooling at all, including a deterministic check that verifies a decision against human-specified, inspectable criteria. I suggest the boundary be drawn by function rather than by technology, with two conditions doing the work: the reviewer must independently evaluate the matter and remain free to depart from anything a tool reports, and the tool's criteria must be human-authored and independent of the Covered ADMT's own logic, so that the original system's reasoning cannot be relabeled as verification. A tool meeting those conditions supports the reviewer; a tool that recommends an outcome, or that re-runs the decision logic under review, substitutes for the reviewer, and that is what the provision should prohibit. In particular, a second AI system cannot be the review: it is more automation, and its output is input to the human reviewer, never a substitute for them. A deterministic check is different in kind: it applies judgments a human wrote down, in a form a human can inspect, and accountability stays with the person who approved the rules it runs.

**The question for the revised draft.** As I read the proposed rules, they speak of a Covered ADMT, its version number, and its Developer in the singular, and nothing addresses how ADMT is identified when a consequential decision is produced by several coordinated automated components, an orchestrating system delegating steps to subordinate ones, none of which alone looks like the decision-maker. Deployments are already ahead of the rules on this point. I suggest the governing principle plainly: coverage should follow the decision, not the architecture. If the integrated process that produces the consequential decision is assessed as a whole, aggregating the coordinated components that materially contribute to it, then a deployer's architecture cannot be arranged into an exemption from what a consumer is owed. The operational mechanics deserve more care than a comment letter's margin, and I can provide further written detail on request.

**On the Department's vendor questions.** The Notice asks whether the rules should define an ADMT Vendor and allocate responsibilities where a Deployer runs a Covered ADMT through one. In the small firms my practice serves, the deployer that operates its own Covered ADMT is the exception; most run models and features a vendor builds and hosts. Two clarifications would help. First, define the arrangement and state that the Deployer remains responsible for its Deployer obligations: performance may be contracted, responsibility may not, and a firm should not be able to move a duty off its own books by moving the system off its own servers. Second, require the vendor to assist: the records and system information a Deployer needs to honor a consumer's rights commonly live in the vendor's

systems, and an assistance duty analogous in structure to processor obligations under the Colorado Privacy Act - scoped to what only the vendor can provide - would give small deployers a statutory baseline instead of a term each firm must negotiate alone. Where a staffing agency or similar intermediary faces the consumer directly, delivery of a disclosure may sensibly sit with the party the consumer actually deals with, provided the rules keep the Deployer accountable that delivery happened; accountability staying put is also what removes any incentive to interpose an intermediary.

On the Chatbot Safety Act rules, one observation from the same seat: the disclosure mechanics in Rule 10, the daily first-interaction disclosure, the three-hour refresh in continuous interactions or a persistent disclaimer, and the answer-when-asked requirement, serve the user's ongoing awareness sensibly, and Rule 9's rule that government-issued identification may not be the sole age-assurance method preserves a path that does not require identity documents at all. The ADMT rules teach compliance through worked, sector-specific examples; I encourage the Department to give the Chatbot Safety Act rules the same treatment before the compliance date, because worked examples materially improve a small operator's ability to translate definitions into operating controls.

I appreciate the Department's early and open process. Supporting cost information on any of the above can be submitted if useful to the Department.

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