September 4, 2026

Office of the Attorney General
Attn: The Honorable Phil Weiser
1300 Broadway
Denver, CO 80203

Re: "Automated Decision-Making Technology Rulemaking"

Dear Attorney General Weiser:

On behalf of the Computer & Communications Industry Association (CCIA), I write to offer comments requested by Attorney General Weiser on the Automated Decision-Making Technology Rulemaking period following the passage of SB26-189.

CCIA is an international, not-for-profit trade association representing a broad cross-section of communications and technology firms.[1] Proposed regulations on the interstate provision of digital services therefore can have a significant impact on CCIA members. While we fully support the legislature's goals of ensuring consumer protection, transparency, and accountability, the broad statutory language creates significant operational uncertainties. Left unaddressed, these ambiguities could chill AI innovation in Colorado, disrupt routine commercial and employment operations, and result in multi-jurisdictional compliance fragmentation.

Definitional Concerns

We suggest that rules incorporate critical enterprise exceptions into the definition of "Accessible to the General Public" in Section 6-1-1701(3.5)(a). While the rule appropriately excludes internal-only workforce deployments, its broad treatment of commercial applications risks capturing routine business software. Rules should establish clear exceptions for conversational AI systems primarily designed and marketed for enterprise use, customer service, sales purposes, or internal business productivity. Furthermore, exceptions should extend to embedded chatbot features within broader non-companion applications, as well as systems purposefully designed to handle narrow, discrete, or topic-specific user requests. Without these targeted exclusions, low-risk operational tools would be swept into a framework intended for consumer-facing companion AI, creating significant compliance drag without advancing consumer protection.

The proposed definition of "Age Category" requires narrowing to align with statutory authority under C.R.S. Section 6-1-1708(2). Because the statute focuses obligations strictly on minor account holders under eighteen years of age, establishing granular categories for users aged eighteen to twenty-four exceeds the scope of the underlying law. Requiring age-band data for young adults carries severe operational implications: under annual reporting obligations such as Rule 13.4, operators would be forced to implement intrusive age-estimation mechanisms

[1] For more than 50 years, CCIA has promoted open markets, open systems, and open networks. CCIA members employ more than 1.6 million workers, invest more than $100 billion in research and development, and contribute trillions of dollars in productivity to the global economy. A list of CCIA members is available at https://www.ccianet.org/members.

for adult users merely to report user distribution across the eighteen-to-twenty-four bracket. Additionally, embedding this category in the regulatory text creates an infrastructure expectation that could support future expansion of obligations to adult accounts without legislative approval. Rules should align the definition directly with statutory duties by limiting tiers to under thirteen, thirteen to seventeen, eighteen and older, and non-relevant organizational accounts, or explicitly confirm that no substantive duties or age-assurance mandates attach to adult categories.

Materially Influence Standard

CCIA strongly advocates for the adoption of Standard 2 regarding the presumption of material influence under Rule 6.2 and C.R.S. Section 6-1-1701(13)(a). Establishing a framework that recognizes de minimis AI contributions moves Colorado significantly closer to parity with comparable state frameworks, most notably California's ADMT[2] threshold requiring technology to replace or substantially replace human decision-making. From the inception of Colorado's AI legislative process, Governor Polis emphasized the critical importance of preventing a fragmented patchwork of state regulations that imposes conflicting compliance obligations on businesses[3]. Standard 2 helps fulfill this directive by preserving the core legal defense that a human decision-maker exercised genuine independent judgment, rendering the AI output a de minimis factor. In contrast, broader scoping approaches risk sweeping routine informational, analytical, and workflow tools into the regulatory framework, capturing business activities that the statute was never intended to govern.

While Standard 2 represents a major step forward, the mechanism for rebutting the presumption of material influence under Section B.2 contains an evidentiary rigidity that threatens to render the defense unworkable. As drafted, Section B.2 outlines a four-part test joined by the conjunctive term "and," suggesting that an operator must satisfy all four conditions simultaneously to prove an output was a de minimis factor. This cumulative hurdle fails to account for diverse enterprise workflows and creates an excessively strict burden of proof. To ensure the presumption remains meaningfully rebuttable, rules should revise the opening phrase of Section B.2 to specify that "relevant evidence may include, but is not limited to" the listed factors, establishing an illustrative rather than exhaustive framework. Additionally, rules should replace the conjunctive "and" after subsection B.2.c with a disjunctive "or." This adjustment ensures that demonstrating key indicators of independent human judgment, such as substantial independent analysis or statutory decision-making authority, can successfully rebut the presumption even if every enumerated factor is not strictly present.

Finally, rules should explicitly harmonize the de minimis assessment in Section A with the administrative tool exclusions set forth in Section D and C.R.S. Section 6-1-1701(2)(b)(II). The statute expressly exempts tools used solely to summarize, organize, translate, draft, route, or present information for human review. Preserving this statutory boundary requires confirming that outputs from these administrative processing tools are categorically excluded from triggering a presumption of material influence in the first instance. By clarifying that Section A's de minimis factors serve as a broad safety valve for substantive AI outputs, while Section D

[2] California Privacy Protection Agency (CPPA) Proposed ADMT Regulations (11 CCR § 7900 et seq.)
[3] Gov. Jared Polis SB 24-205 Signing Statement (May 17, 2024)

provides a baseline statutory exemption for routine administrative utility, rules will deliver the regulatory clarity needed to support responsible AI deployment without overburdening human-centered decision-making processes.

Multiparty Arrangements

To preserve statutory fidelity, the Attorney General's Office should strike the definition of "Midstream Developer" (Rule 2.2) and Rule 4 in their entirety, without substituting alternative regulatory language. C.R.S. Section 6-1-1701 establishes a clear regulatory framework centered strictly on the roles and responsibilities of developers and deployers; it does not contemplate or authorize an intermediary midstream role.

Creating novel regulatory classifications and imposing mandatory supply-chain pass-through requirements via rulemaking exceeds administrative authority to implement the statute. Superimposing these unmandated intermediary categories creates legal liabilities and operational uncertainty for software integrators while disrupting the carefully balanced allocation of duties enacted by the legislature. Preserving the statutory framework requires adhering strictly to the established definitions of developer and deployer without introducing intermediary compliance obligations.

To the extent the Attorney General's Office seeks to address potential information gaps across multi-party AI stacks, commercial contracting provides the proper and intended mechanism. Section 6-1-1702(1)(e) explicitly contemplates contractual arrangements to manage information-sharing between developers and deployers. Market-driven contracts allow parties to flexibly allocate responsibilities, specify documentation requirements, and account for specific technical deployment contexts, resolving supply-chain information needs without imposing rigid, non-statutory regulatory mandates.

Finally, the illustrative example in Rule 4.2(B) introduces significant legal confusion and threatens proprietary trade secrets.[4] Requiring intermediaries to collect downstream documentation for third-party foundational models forces entities to expose sensitive system architectures and confidential technical integrations. If the rules retain illustrative examples, they should focus strictly on downstream deployers operating covered ADMT within statutory boundaries, rather than establishing scenarios that jeopardize intellectual property and trade secret protections across multi-layered software products.

Regarding multiparty arrangements and vendor disclosures, rules should recognize that allocating deployer obligations among commercial entities is fundamentally a contractual issue rather than a matter for rigid administrative designation. In complex operational setups, such as when a staffing agency utilizes ADMT to screen applicants and provides a constrained shortlist to an employer, both entities exercise decision-making authority at distinct points in the hiring pipeline. Attempting to designate a single party as the exclusive deployer across all multi-party scenarios ignores these commercial realities and distorts statutory liability under C.R.S. Section 6-1-1701. Rather than establishing blanket administrative rules specifying which party must fulfill deployer duties, such as delivering consumer disclosures, rules should explicitly acknowledge that commercial parties may contractually allocate their respective

[4] Colorado Uniform Trade Secrets Act (C.R.S. § 7-74-102 et seq.) & C.R.S. § 6-1-1704(5)

compliance responsibilities based on operational control and consumer touchpoints. This approach preserves commercial flexibility, ensures clear accountability, and avoids imposing artificial regulatory designations on complex business relationships.

Developer Obligations

We oppose Rule 5.2(C)(2) and (3) as developer obligations, as these provisions improperly reassign operational duties to software developers who lack the necessary deployment context. Requiring developers to provide instructions for real-time monitoring of specific consequential decisions or individual output adjustments conflates developer capabilities with deployer responsibilities. Rules should reassign these post-deployment monitoring obligations to deployers, who possess direct control over operational implementation. If retained for developers, the rule must maintain precatory "should" phrasing rather than mandatory directives, and its scope must be strictly limited to technical information reasonably possessed by the developer at the time of delivery, consistent with the framework set forth in C.R.S. Section 6-1-1702(1)(d).

Furthermore, we strongly oppose expanding developer disclosure requirements to mandate explicit testing results, negative disclosures regarding non-testing, or hyper-granular training data breakdowns. The inquiry regarding mandatory testing disclosures represents an impermissible attempt to reinstate risk assessment and impact reporting obligations that the legislature deliberately removed from the underlying statute. Imposing rigid mandates regarding testing methodologies and training data details exceeds the Attorney General's statutory authority to implement and clarify the law under Section 6-1-1702(1)(e). Such requirements would create compliance burdens even more restrictive than international frameworks like the EU AI Act[5] and would establish perverse disincentives against voluntary AI safety testing by transforming internal quality assurances into regulatory exposure.

Finally, rules should refrain from prescribing rigid forms, technical formats, or mandatory channels for developer-to-deployer disclosures. Section 6-1-1702(1)(e) already establishes an effective statutory baseline, calibrating developer information to what deployers reasonably require to fulfill their compliance obligations. Commercial entities are best positioned to negotiate the channels and formats that suit their specific software architectures and enterprise relationships. Attempting to dictate the precise mechanics of technical documentation transfers via administrative rulemaking needlessly burdens commercial integration without providing meaningful utility to deployers or end-users.

Post-Adverse Outcome Disclosure Requirements

Rules governing the mechanics of adverse outcome notices under Rule 6.2(A) and Rule 6.4(B) introduce severe operational friction without delivering corresponding benefits to consumers. Mandating dual-channel delivery under Rule 6.2(A)(1) forces enterprises engaging in high-volume, digital-first transactions to maintain redundant, costly physical mailing infrastructure alongside electronic channels. The rules should clarify that a deployer fulfills written notification obligations by communicating exclusively through the primary, ordinary channel through which it routinely interacts with the consumer. Furthermore, the granular

[5] EU Artificial Intelligence Act (Regulation (EU) 2024/1689), Recital 60 & Articles 26–28

requirements in Rule 6.4(B) demanding effective dates, step-by-step descriptions of internal human reviewer roles, and proactive instructions on how consumers can remedy incomplete data exceed statutory intent. C.R.S. Section 6-1-1704 mandates a clear description of the role automated decision-making technology played in an adverse decision, not the conversion of disclosures into exhaustive workflow audits or application remediation guides.

The requirement to disclose specific numeric risk scores, individual inference breakdowns, and exact automatic denial cutoffs under Rules 6.4(B)(5), (6), and (8) violates critical statutory safeguards for trade secrets under C.R.S. Section 6-1-1704(5). Mandatory disclosure of specific numerical scores or proprietary decision thresholds exposes underlying algorithmic models, enabling bad actors to reverse-engineer and game evaluation systems. This creates severe security risks and compromises model integrity across financial, housing, and employment contexts. Rather than requiring individual-level mathematical reconstruction, rules should align disclosures with the categorical framework authorized by Section 6-1-1704(4)(a)-(b). Deployers should be required only to inform consumers that a score, profile, or inference informed the determination and identify the broad categories of personal data utilized, balancing meaningful transparency with the protection of proprietary intellectual property.

We strongly urge the Attorney General's Office to strike Rules 6.4(B)(11) and 6.4(B)(12) in their entirety, as these provisions impermissibly transform a disclosure obligation into a strict liability design constraint. Declaring a deployer non-compliant whenever it cannot explain with total precision how technology materially influenced a specific decision or articulate the principal reasons with absolute accuracy exceeds administrative rulemaking authority under the Act. Modern machine learning models rely on complex multi-variable analysis that does not always translate into atomistic, human-interpretable logic for every singular output. By converting the statutory obligation of a good-faith system explanation into an absolute explainability mandate, these rules effectively penalize advanced analytical tools, conflict with statutory trade secret exemptions, and expose compliant entities to severe regulatory liability for normal technological operations.

The employment example set forth in Rule 6.5(F)(2) establishes an unworkable precedent that severely distorts the scope of an adverse outcome. Classifying non-selection for a discretionary promotion or bonus as an adverse outcome conflates the failure to obtain an unearned opportunity with the affirmative denial or termination of an existing right or entitlement. If applied as drafted, this logic would sweep virtually every competitive selection process, candidate shortlisting tool, or internal talent management program into an overwhelming disclosure regime, forcing employers to issue formal adverse outcome notices to every non-selected applicant in every routine hiring or internal movement process. Furthermore, the example incorrectly frames human reviewers as passive verifiers of an automated score, understating the primary, central role that human managerial judgment plays in personnel evaluations. Rules should explicitly clarify that adverse outcomes are limited to the denial or termination of existing access, rights, or opportunities, while accurately reflecting human reviewers as primary decision-makers.

Finally, Rule 6.6 must be modified to keep consumer information requests operationally scalable and aligned with established standards in neighboring jurisdictions. Requiring deployers to disclose specific, named third-party data sources, such as individual data brokers,

public records databases, or specific courts under Rule 6.6(A), exceeds C.R.S. Section 6-1-1704(3)(b), which requires only the disclosure of types, categories, and general sources of data. Providing centralized, standardized disclosures at the decision-type level fully satisfies consumer informational needs without requiring costly, bespoke outcome reconstructions for every single inquiry. Additionally, the proposed ten-business-day response window in Rule 6.6(B) is drastically short for complex data processing environments. Aligning this timeframe with California's forty-five-day ADMT standard honors Governor Polis's goal of regional regulatory alignment, providing deployers with a realistic timeframe to compile thorough and accurate information without creating conflicting multi-state compliance regimes.

Consumers Rights

The procedural framework for exercising consumer rights under Rule 7.2 requires changing from rigid mandates to flexible performance standards. Compelling deployers to maintain specific channels, fulfill telephonic requests within twenty-four hours, and adopt verbatim heading text introduces unnecessary administrative friction without enhancing consumer comprehension. Instead, rules should permit deployers to receive requests through the primary, established channels utilized in their ordinary course of business. Furthermore, replacing mandated verbiage with a clear disclosure standard would allow for rights to be communicated in plain language tailored to specific digital interfaces.

Rule 7.3 improperly expands the scope of personal data by sweeping automated system outputs into access requests. Under C.R.S. Section 6-1-1705(1)(a)(I), consumers are entitled to the data used in a decision, not proprietary algorithmic outputs, risk scores, or predictive classifications. Conflating input data with system-generated inferences distorts statutory definitions and jeopardizes proprietary analytical models. Additionally, requiring deployers to enumerate every potential input, rather than the specific data actually material to a consequential decision, creates an unworkable technical burden. Rules must clarify that access rights extend strictly to the personal data inputs relied upon in the determination.

The data correction framework in Rule 7.4 requires modification to preserve enterprise functionality. Mandating corrections across all existing systems imposes an expansive infrastructure requirement that ignores the specific nature and purpose of data processing. More critically, the requirement under Rule 7.4(D) to automatically stay adverse outcomes pending correction lacks statutory authorization and creates severe operational risk. Allowing unverified assertions of inaccuracy to freeze commercial actions—such as hiring or credit decisions—is untenable. Rules should strike the mandatory stay provision and ensure correction obligations are evaluated relative to the specific processing context.

Regarding human review protocols under Rule 7.7(A) through (C), rules should focus on accessibility and authority rather than prescribing internal corporate structures. Rule 7.7(B) improperly forces a re-performance of decision-making upon any correction, regardless of whether that data impacted the result. Reconsideration should only be triggered when corrected data is material to the outcome, and deployers should not be required to re-run decisions where reversal is not commercially reasonable. Furthermore, Rule 7.7(C) improperly superimposes a judicial independence standard; C.R.S. Section 6-1-1701(15) requires authority to modify a decision, not structural separation from the original decision-maker.

Rules should clarify that reviewers must exercise independent judgment without being bound by automated outputs.

Rules 7.7(D) and 7.7(E) incorrectly measure the validity of human review by its outcome rather than procedural rigor. Declaring that a decision reversal indicates meaningful review creates an improper outcome-based compliance standard; mechanisms must be evaluated by objectivity and training frameworks. Additionally, the Attorney General's Office should strike the "basic human need" presumption in Rule 7.7(E)(3), which imposes an unauthorized and unworkable commercial reasonableness burden. When evaluating whether human review is commercially reasonable under Rule 7.7(E)(4), the assessment must account for the cumulative, system-wide volume of incoming requests across an enterprise rather than viewing requests in isolation.

Finally, Rules 7.7(F) and 7.7(G) impose disclosure mandates and operational stays that threaten the viability of automated systems. Requiring deployers to externalize internal deliberative reasoning under Rule 7.7(F) exceeds statutory intent; compliance should require only notice of the final outcome and a general description of materials reviewed. Furthermore, the mandatory stay requirement in Rule 7.7(G)(1) must be struck, as freezing operations during a review window has no statutory foundation and renders workflows unviable. Timeline requirements must incorporate risk-tiered flexibility reflecting the cumulative volume of requests to ensure human review capacity remains sustainable for high-impact decisions.

* * * * *

CCIA strongly advocates for the Attorney General's Office to recalibrate the draft regulations to ensure strict adherence to the statutory framework established by SB26-189. By narrowing overly broad definitions, upholding critical protections for proprietary trade secrets, and transitioning from prescriptive mandates to flexible performance standards, Colorado can mitigate significant operational friction and prevent multi-jurisdictional compliance fragmentation. Implementing these necessary revisions will support a landscape for responsible technological innovation while fulfilling the underlying legislative intent regarding consumer transparency and accountability.

Sincerely,

Aodhan Downey
State Policy Manager, West Region
Computer & Communications Industry Association