July 13, 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.

Colorado Should Focus on Well Tailored Rules

We suggest that rules tighten the threshold for the "Materially Influence" standard in Section 6-1-1701(13). Including any output that is a "non-de minimis factor" affecting a decision is a significantly lower threshold than comparable frameworks, such as California's ADMT regulations, which require the technology to "replace or substantially replace" human decision-making. Rules should establish clear presumptions that general-purpose AI tools where a human retains full decision authority, do not trigger the Act, even if the output is referenced in the process.

Rules should explicitly address whether ADMT outputs that operate purely 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 an operational need for 40 workers on a given shift, without determining which specific individuals fill those slots, arguably constrains downstream assignment decisions. The inclusion of the word "constraining" in the definition of material influence under Section 6-1-1701(13)(II) could be misread to capture such aggregate-level outputs. Rules should clarify 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 from scope, even if their outputs indirectly constrain the universe of individual decisions made by a human downstream.

[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.

Next. ensuring that the technology exclusions remain function-based and not conditioned on methodology is key for future machine learning (ML) adoption and preserves the legislative intent. Of the 14 exclusions in Section 6-1-1701(2)(b)(I), only spreadsheets at subsection L are conditioned on the absence of ML or large language models. The remaining 13 exclusions, including anti-virus, firewalls, spam filtering, and spell-checking, carry no methodology-based conditions. Because modern infrastructure services inherently incorporate ML to improve performance, rules must expressly confirm that these 13 exclusions remain function-based and are not subject to an implicit methodology condition. The deliberate decision in the statute to restrict the ML condition to spreadsheets is confirmed by parallel language in the consequential decision exclusions under Section 6-1-1701(3)(b)(III). Because this boundary was intentionally drawn, future rules should preserve it.

The definition of "adverse outcome" under Section 6-1-1701(1) and its trigger scope require narrowing on multiple fronts to be workable. First, additional guidance is needed on the pricing provision in subsection 1(b). It remains unclear whether AI-assisted functions that influence price, such as medical billing, coding, or insurance rate-setting, constitute adverse outcomes simply because the resulting price is higher than what might otherwise apply, even where access to the service itself is not fundamentally altered. Rules should clarify that a pricing-related adverse outcome exists only where the differentiated terms are reasonably likely to materially limit, delay, effectively deny, or otherwise fundamentally alter access, and not merely because those terms differ from what is offered to other consumers.

In addition, and more broadly, the definition's trigger scope needs a clear materiality threshold. If read broadly enough to encompass routine operational determinations such as shift assignments or schedule changes, the definition would generate a volume of required notifications large enough to undermine the statute's own protective purpose. Because each notification carries an accompanying right to request human review, triggering a 30 day response clock, applying that same obligation to routine operational matters at the same rate as high impact decisions such as hiring rejections or terminations would overwhelm any realistic compliance regime. Compliance resources would inevitably be spread thin across a high volume of low stakes notifications, leaving fewer resources available for the individuals who experience genuinely adverse outcomes and who most need timely, substantive engagement.

Rules should establish a clear materiality threshold, both for pricing determinations and for operational decisions generally, so that the notification obligation is limited to decisions carrying a meaningful and lasting impact on the individual. This would ensure that the right to human review remains a meaningful safeguard rather than a procedural formality triggered indiscriminately across routine business operations or minor pricing variances.

Minimizing the Risk of Producing Unintended Consequences

Several unintended consequences are reasonably foreseeable if the statutory language is implemented without clear guardrails, beginning with a severe chilling effect on AI innovation. The conversational AI exclusion in Section 6-1-1701(2)(b)(III) imposes a strict two-prong conjunctive test. Under prong A, the technology must not be contracted, advertised, marketed, configured, or intended by a person to be used in a consequential decision — a five-part

disjunctive standard in which failing any single element defeats the exclusion entirely. Under prong B, the technology must also be subject to an acceptable use policy that prohibits generated content from being used in a consequential decision. The breadth of prong A, particularly its inclusion of "configured" and "intended," 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 the acceptable use policy requirement in prong B is satisfied by a general prohibition in terms of service or acceptable use documentation, without requiring active monitoring, and that a deployer's independent configuration or intent does not retroactively defeat a developer's exclusion where the developer did not share that intent.

The prospect of multi-jurisdictional compliance fragmentation also presents a formidable operational challenge. Because the "materially influence" threshold is notably lower than the benchmarks established in other states, interstate technology providers face disparate compliance obligations without a commensurate gain in consumer protection. We urge the Attorney General to harmonize these definitions with prevailing national frameworks and to establish clear safe harbors for firms that adhere to substantially similar requirements in other jurisdictions.

The post-adverse outcome disclosure mandate in Section 6-1-1704(3)(b) raises further concerns regarding data-source gaming and cybersecurity. While disclosing the "types" and "categories" of data aligns with established privacy principles, the additional requirement to disclose "sources" is ambiguous and could be read to require identification of specific vendors or proprietary feeds — potentially enabling individuals to manipulate inputs, revealing trade secrets, or exposing system architectures to security threats. Rules should specify that "sources" refers to broad categories of origin, such as employer records or credit reporting data, consistent with the transparency approach of the Colorado Privacy Act. This clarification would preserve essential trade secret and system integrity protections.

Finally, rules must guard against disruption of commercial risk allocation and overbroad application to routine employment and back-office functions. The indemnification limitations in Section 6-1-1707(7) conflict with standard contractual arrangements between sophisticated entities, and the tension between the prohibition in subsection 7(a) and the savings clause in subsection 7(c) creates significant uncertainty. Future rules should delineate the boundary between prohibited indemnification for anti-discrimination liability and permissible allocation of other liabilities, and should confirm the scope of the developer carve-out in subsection 7(b).

Similarly, the pricing mandate within the consequential decision framework of Section 6-1-1701(3)(a)(II) risks unintentionally capturing standard administrative functions, including AI-integrated medical billing, ambient clinical documentation, and clinical coding. These technologies serve as assistive tools for practitioners who retain ultimate professional responsibility and do not themselves generate autonomous clinical determinations. Future guidance should offer illustrative scenarios confirming that such tools do not fundamentally alter consumer eligibility or access, preserving their role as beneficial back-office supports.

An overbroad interpretation of "consequential decisions" in the workforce sector — one that sweeps in everyday operational matters like shift rotations, task assignments, or scheduling — likewise threatens to impose prohibitive compliance costs. Such a burden may compel

organizations to relocate personnel or abandon automated efficiencies in favor of manual systems that invite inconsistency. Rules should narrow this scope to actions that substantially affect an individual's professional status or core opportunities, excluding routine operational logistics that provide efficiency and predictable outcomes for the workforce.

Stakeholder Suggestions for Clarity

To provide meaningful relief to stakeholders, the rulemaking process must resolve several distinct areas of compliance uncertainty that currently make it difficult for regulated entities to operationalize the statute in good faith.

One key area concerns the "materially influence" standard under Section 6-1-1701(13). As currently written, the standard offers no clear, objective indicators that entities can apply to their own systems, leaving many unable to determine with confidence whether their use of AI triggers coverage under the statute at all. This uncertainty is compounded by the fact that AI tools vary widely in how directly their outputs feed into a final decision, meaning entities operating in good faith may reach very different conclusions about their own compliance obligations depending on how they interpret the term. The rules should explicitly confirm that the definition's two prongs, requiring an output to be a non-de minimis factor and requiring that factor to affect the outcome, are strictly conjunctive, such that both elements must be satisfied before a tool is considered to materially influence a decision. Absent this clarification, entities may over-comply out of caution, incurring unnecessary costs, or under-comply due to reasonable misreadings, exposing themselves to enforcement risk they did not anticipate.

A second area involves clarifying the consumer definition and geographic scope under Section 6-1-1701(4)(b). The provision's inclusion of any individual whose access to, eligibility for, or opportunity in Colorado is evaluated creates significant ambiguity around the statute's extraterritorial reach, extending its scope well beyond the geographic limitations established under the Colorado Privacy Act. This raises practical difficulties for national or multi-state employers and technology providers, who may be unable to determine in advance whether a given individual's interaction with their systems falls within Colorado's jurisdiction, particularly in remote work, multi-state hiring, or distributed service contexts where an individual's connection to Colorado may not be immediately apparent. Clearer guidance on how geographic scope is determined would allow entities to build compliance processes with confidence rather than defaulting to the broadest possible interpretation out of caution.

A third area concerns the friction between disclosure requirements and trade secret or security protections. The scope of permissible withholding under Section 6-1-1704(5) and (7) needs clarification so that deployers can confidently identify what information may be protected from disclosure. This is particularly important because the "required by law" limitation in subsection 7 is considerably narrower than the broader functional exclusions found elsewhere in the statute, creating an inconsistency that leaves deployers uncertain about which standard actually governs their disclosure obligations in practice. Without resolution, entities face a difficult choice between over-disclosing sensitive system details, potentially compromising trade secrets or creating security vulnerabilities, and under-disclosing in ways that could expose them to enforcement action.

General Rules with Sector Specific Guidance Will Help Industries Comply in Colorado

The department should divide its regulatory outputs between generally applicable rules that establish systemic baseline standards and sector-specific guidance driven by practical examples.

Generally applicable rules are the appropriate vehicle to address the objective indicators and presumptions for the “materially influence” standard, the consumer definition and geographic nexus requirements, and the conversational AI exclusion's five-part disjunctive test regarding whose intent controls. Universal rules should also confirm that the enumerated technology exclusions remain function-based and unconditional, while setting the foundational timelines for human review and data correction requests.

Conversely, sector-specific guidance and illustrative examples are best suited to handle the highly variable environments of specific industries. This flexible mechanism should be used to outline post-adverse outcome disclosure content and define the boundary between routine billing functions and consequential decisions in healthcare, insurance, and lending. Sector guidance should also clarify how the Act interacts with existing federal regulatory frameworks like ECOA, FCRA, HIPAA, and FERPA within each covered domain. Furthermore, it should define what constitutes a commercially reasonable human review obligation across entities of different sizes and decision volumes, while clarifying how the "similarly situated consumers" comparator in the adverse outcome definition applies across different operational contexts.

To support these specific regulations, industry participants encourage the Department to consider empirical evidence and comparative frameworks. This includes a comparative analysis of other jurisdictions' AI governance frameworks, such as the California ADMT regulations, the EU AI Act implementing measures, and the NIST AI Risk Management Framework, which demonstrate that function-based rather than methodology-based scoping successfully reduces compliance costs while achieving equivalent consumer protection outcomes. The Department should also leverage the Colorado Privacy Act rulemaking record, which successfully addressed similar scope and definition challenges, as well as existing federal frameworks like ECOA Regulation B adverse action notices and FCRA consumer disclosures, which provide proven, workable models for explaining automated decisions to consumers.

Well Crafted Rules Create a Suitable Environment for Well Targeted Regulatory Implementation

To clarify the core definitions of the Act, rules should establish clear, objective indicators and presumptions for when an ADMT output constitutes a non-de minimis factor versus an incidental, trivial, or clerical use. The two prongs of the material influence definition must be confirmed as conjunctive, and illustrative examples should confirm that general-purpose AI tools used for summarization, translation, or drafting do not materially influence a consequential decision even if the output is referenced in the process, especially when a human retains full decision authority.

For the consequential decision pricing provision, rules should provide illustrative examples clarifying the boundary between routine pricing or billing functions and consequential decisions that fundamentally alter access, eligibility, or opportunity, confirming that back-office operational uses of AI like billing, coding, and scheduling fall outside the definition.

Useful objective indicators to determine whether an ADMT materially influences a consequential decision should be explicitly codified. These factors should include whether a human decision-maker retains independent authority and actually exercises discretion over the final outcome, whether the ADMT output is merely one of multiple factors considered where other factors have comparable weight, and whether the ADMT output can be easily disregarded by the decision-maker without professional consequence. Additional indicators should evaluate whether the output is used to narrow the universe of options presented to a human, which is potentially material, versus merely organizing or formatting information, which is not material, and whether the removal of the ADMT output from the workflow would leave the outcome unchanged in the substantial majority of cases. Sector-specific examples would be extremely helpful, contrasting an instance of clear material influence, such as an ADMT that rank-orders job applicants where the top-ranked candidates are exclusively selected for interviews, against an instance of non-material influence, such as an AI tool that drafts an interview summary that a hiring manager reviews de novo alongside their own independent notes.

To determine whether an ADMT component is designed, marketed, intended, documented, advertised, configured, or contracted to be used as part of a covered system, rules should focus on objective commercial facts. Relevant facts should include whether the developer's contract terms affirmatively authorize or contemplate use in consequential decisions, whether the developer provides specific configuration guidance or templates for consequential-decision use cases, and whether the technology is purpose-built for a covered domain versus being a general-purpose tool with many broad applications.

Conversely, rules should clarify that specific facts do not independently establish coverage or constitute actual knowledge under Section 6-1-1701(8)(b)(III). These non-triggering facts include general descriptions of system capabilities, including the mere technical ability to be used in making consequential decisions, a downstream deployer's independent decision to integrate the technology without the developer's knowledge or authorization, the mere receipt of customer inquiries about potential use in consequential-decision contexts, or the publication of general use case examples that happen to include scenarios touching covered domains among other applications.

To ensure the Act interoperates cleanly with existing laws, rules should explicitly recognize and defer to federal adverse action notice frameworks under ECOA Regulation B and the FCRA to avoid duplicative or conflicting obligations, building directly on the framework in Section 6-1-1704(6). Rules should also confirm alignment with the Colorado Privacy Act consumer definitions and data correction mechanisms to minimize compliance fragmentation within Colorado itself.

Safe Harbors Allow Good Actors to Be Compliant While Punishing Those Intending Harm

The Attorney General should consider providing safe harbors for entities that comply with substantially equivalent requirements in other jurisdictions, particularly regarding the material influence threshold, to reduce unnecessary multi-jurisdictional compliance fragmentation. Furthermore, rules should expressly address the interaction between ADMT Act disclosure obligations and HIPAA or FERPA confidentiality requirements to provide clarity for healthcare and education deployers. Finally, rules should recognize that ADMT used within FDA-regulated medical devices or clinical decision support tools is already subject to rigorous federal safety requirements that should be deemed to satisfy corresponding obligations under this Act, thereby avoiding duplicative compliance burdens.

Further Guidance for Post-Adverse Outcome Disclosure Requirements are Needed

In the workplace context, rules must clarify what specifically constitutes the moment of an adverse outcome that triggers the 30-day clock, particularly for multi-step decision processes where ADMT contributes to an intermediate determination that later feeds into a final decision. For example, if an ADMT-generated risk score is calculated on Day 1 but the employment action taken partly in response to that score occurs on Day 30, the rules must clarify whether the 30-day window runs from the initial negative score generation or the final adverse employment action. From a practical perspective, the notification and appeal requirements should only be triggered when a decision has a recognizable impact on the data subject, given that interlocutory appeals are operationally infeasible. Similarly, for ongoing employment relationships, rules should clarify whether each individual payroll cycle, schedule assignment, or performance evaluation that relies on ADMT output is a separate consequential decision with its own 30-day clock, or whether a single annual or periodic disclosure satisfies the requirement, which would align more closely with the standard in the California ADMT regulations.

Additionally, the "similarly situated consumers" comparator in subsection 1(b) is currently completely undefined in the Act. Rules should provide guidance on what factors define the comparator group, who bears the burden of establishing the comparison, and how risk-based pricing, which inherently differentiates between consumers based on legitimate risk factors, interacts with this standard in insurance and lending contexts.

To ensure disclosures are accurate and informative without becoming overly onerous for deployers or confusing for consumers, rules should establish a tiered disclosure framework that permits a concise initial notice meeting the 30-day requirement, with additional technical detail made available upon consumer request. The disclosure of types, categories, and sources of personal data under Section 6-1-1704(3)(b) should be interpreted so that sources means categories of sources, such as credit bureau data or employer-provided records, rather than specific vendor names or database identifiers. This interpretation is consistent with the privacy-protective approach in the Colorado Privacy Act and preserves the Act's trade secret protections at Section 6-1-1704(5). Disclosing specific data sources raises significant concerns

that consumers could manipulate identified inputs to game future decisions and may reveal sensitive system architectures that create security vulnerabilities.

The bill's own limiting principle, which states that disclosure is required only to the extent the deployer receives the necessary information from the developer in compliance with Section 6-1-1702, should be given practical effect in the rules. Rules should firmly confirm that disclosures need not reveal information that would compromise system integrity, enable manipulation, or create security risks, consistent with the existing trade secret and cybersecurity protections in Section 6-1-1704(5) and (7). Providing model disclosure templates or safe-harbor language would significantly reduce compliance uncertainty while ensuring consumers receive meaningful information. Ultimately, rules should emphasize flexibility in compliance with disclosure and notice requirements to accommodate different entity use cases, user experience designs, and operational contexts, rather than prescribing a single format or rigid approach.

Guidance regarding categories, sources, or types of personal data must be specific enough to enable compliance but not so granular as to create gaming or security risks. The bill's use of types, categories, and sources at Section 6-1-1704(3)(b) should be read harmoniously, with sources confirmed to mean categories of sources rather than identification of specific databases, vendors, or data feeds. This approach is consistent with FCRA adverse action notices and CPA disclosure practices, and it preserves the Act's trade secret and system integrity protections.

To properly address how disclosure requirements interact with federal or state laws, rules should confirm that the ECOA and FCRA compliance safe harbor in Section 6-1-1704(6) is broadly construed so that creditors providing compliant federal adverse action notices are not required to engage in separate, duplicative state disclosure processes. Rules must also address how HIPAA-covered entities satisfy disclosure obligations without violating federal health privacy requirements, building on Section 6-1-1708(5). Rules should explicitly confirm that compliance with ADMT Act disclosure requirements does not require or authorize the disclosure of protected health information (PHI) as defined under HIPAA, ensuring that deployers in healthcare contexts can satisfy both regimes simultaneously. Furthermore, guidance should clarify that FERPA-covered deployers may use existing student record access processes per Section 6-1-1704(9). Finally, rules should confirm that Section 6-1-1704(7) protects entities from any disclosure that would compromise the integrity of cybersecurity, fraud prevention, or sanctions compliance programs, while acknowledging the "required by law" limitation and clarifying that it extends to programs required by regulators as a condition of operating licenses, and not just explicit statutory mandates.

Consumers Deserve Clarity on the Technical Feasibility of their Rights Under the Act

While the statutory definition of meaningful human review under Section 6-1-1701(15) provides a useful structure, it requires clear implementation guidance regarding process standards and outcome independence. Rules should specify that meaningful human review requires documented consideration of the case, but does not require that the reviewer replicate the entire complex ADMT analysis from scratch. Additionally, the requirement that the

reviewer does not default to the system output should be clarified to mean the reviewer exercises genuine independent judgment, and not that they must disagree with or override the ADMT output in any particular percentage of cases.

Several relevant factors should be considered in determining what human review requirements are commercially reasonable. These factors must include the nature, severity, and reversibility of the decision's impact on the consumer, as well as the volume of decisions made using the covered ADMT, recognizing that high-volume automated processes may require proportional review mechanisms rather than individualized, manual de novo review of every single decision. Guidance should also evaluate the organizational capacity and physical size of the deployer, the technological feasibility of providing different levels of review, and whether existing appeal or reconsideration processes, such as insurance appeals or loan reconsideration procedures, already satisfy the requirement.

Rules should also weigh the cost of review relative to the value of the decision at stake, and whether the consumer has identified a specific factual error or simply disagrees with the outcome. Rules should explicitly confirm that commercially reasonable does not require deployers to offer unlimited de novo human review in all cases, particularly for high-volume, low-stakes decisions where the cost of individualized human review would be entirely disproportionate.

To provide certainty for both consumers and deployers, a clear timeline for responding to consumer requests for personal data and correction is appropriate. Industry strongly recommends complete alignment with the Colorado Privacy Act response timelines, which allow 45 days with an additional 45-day extension for good cause, rather than creating a new, potentially conflicting timeline. This promotes the Department's stated principle of harmonization while ensuring consumers receive thorough and timely responses.

Consumer Notice Requirements Before a User Interacts with a Covered ADMT Should Be Easily Digestible and Require Further Clarification for Developers

Further guidance and concrete examples are necessary to clarify what constitutes a clear and conspicuous notice. Rules should confirm that the public posting option, which permits a prominent public notice that is reasonably accessible at points of consumer interaction, including through a link or posting that is reasonably proximate to the interaction or transaction, provides a workable compliance mechanism without requiring individualized, real-time notification for every single consumer interaction. This is found in Section 6-1-1704(2). Guidance should include clear examples of acceptable formats, such as website footer links, dedicated transparency pages, privacy policy addenda, or application portal disclosures. It should also confirm that a single notice covering multiple uses of covered ADMTs in the same domain is sufficient, while clarifying that reasonably proximate does not require the notice to interrupt, block, or impede the consumer interaction flow.

Industry experience demonstrates that consumers benefit most from concise, plain-language notices at the point of interaction that clearly identify the type of ADMT being used and the domain in which it operates, paired with a link to more detailed information for consumers who

wish to learn more. This layered disclosure approach—providing a brief notice at the point of interaction with accessible, detailed information behind a link or on request—maximizes utility. Notices that are overly long, technical, or numerous risk creating severe notice fatigue that renders all disclosures less effective, meaning rules should actively promote communication quality over sheer quantity.

To promote accessibility for consumers with disabilities or limited English proficiency, rules should align closely with existing accessibility obligations under the ADA and Section 508 standards, rather than creating novel, potentially inconsistent requirements. Specifically, digital notices should comply directly with WCAG 2.1 AA standards, and deployers should provide notices in the specific languages in which they regularly conduct business with consumers, consistent with existing Colorado consumer protection practices. Rules should explicitly confirm that compliance with these existing federal accessibility mandates fully satisfies the Act's accessibility requirement under Section 6-1-1704(8).

* * * * *

A balanced regulatory approach is vital to maintaining Colorado's position as a premier hub for technological innovation while safeguarding consumer interests. By establishing clear thresholds, prioritizing cross-jurisdictional harmony, and defining explicit exclusions for back-office and routine administrative functions, the Department can craft an effective, workable compliance ecosystem for automated technologies. The department should divide its regulatory outputs between generally applicable rules that establish systemic baseline standards—such as objective indicators for material influence, universal technology exclusions, and core timelines—and flexible, sector-specific guidance driven by practical examples to handle the highly variable operational environments of covered industries.

Sincerely,

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