



September 4, 2026

The Honorable Phil Weiser
Office of the Attorney General
Colorado Department of Law
Ralph L. Carr Judicial Building
1300 Broadway, 10th Floor
Denver, CO 80203

RE: The Automated Decision-Making Technology & Conversational Artificial Intelligence Services Rules

Dear Attorney General Weiser,

The Business Software Alliance appreciates the opportunity to share our views on the Department of Law's (the "Department's") Automated Decision-Making Technology & Conversational Artificial Intelligence Services rules (the "draft rules"). Our comments focus on the draft rules implementing the Automated Decision-Making Technology Act (Senate Bill 26-189) ("the Act") and the questions posed in the Notice of Rulemaking Hearing. BSA is the leading advocate for the global software industry, and our members create cutting edge business-to-business products and services that are used by companies of all sizes across every sector of the economy.¹

AI is changing the way we live and work and it has real-world benefits. Realizing the potential of AI requires trusting that the technology is developed and deployed responsibly. The most effective way to address this issue is through a single, national law. However, just as states took the lead in adopting consumer privacy laws, we recognize states are again leading with AI legislation.

We appreciate Colorado's work to create guardrails for AI technologies that are intended to make decisions that significantly impact consumers' lives. BSA supports the overarching goals of the Act. Our comments focus on three key issues that are critical to helping ensure the Act is workable in practice:

- Clarifying roles and responsibilities along the AI value chain;
- Setting clear and workable thresholds for the AI systems in scope of the Act; and

¹ BSA's members include: Adobe, Alteryx, Amadeus, Asana, Atlassian, Autodesk, Avalara, Bentley Systems, Box, Cisco, Cohere, Cohesity, Dassault Systemes, Databricks, Datadog, Docusign, Dropbox, Elastic, EY, Graphisoft, HubSpot, IBM, Kyndryl, MathWorks, Microsoft, Notion, Okta, OpenAI, Oracle, PagerDuty, Palo Alto Networks, PTC, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Trend Micro, TriNet, Veeam, Workday, Zendesk, and Zoom Communications Inc.

- Extending trade secret protections to intellectual property (IP) and other business confidential information.

I. The Draft Rules Should Clarify Roles and Responsibilities Along the AI Value Chain.

The AI supply chain is complex and evolving, and includes a variety of companies, including AI developers and deployers. In many scenarios, a developer may create an AI system, and a deployer will use that system in the real world. BSA strongly supports the Act's recognition of the distinct roles and responsibilities of different actors along the AI value chain. Creating role-based responsibilities for different AI actors is important because it assigns responsibilities to companies that they can implement based on their access to relevant information and position in the AI supply chain, preventing unworkable obligations for things outside of their control; it ensures that consumers and business customers are better protected because the companies that are best positioned to identify and address risks are the entities designated to take relevant steps to protect them; and it creates a policy framework that is workable in practice.

BSA recommends five changes to the draft rules to clarify the roles and responsibilities of different actors along the AI value chain, which is critical to helping ensure the Act works in practice.

[First, the draft rules should ensure that companies are responsible only for information within their knowledge, possession, or control.](#)

Different companies along the AI value chain have access to different information. A Developer may have information about a Covered ADMT's design, intended uses, and known limitations but, particularly in enterprise contexts, may not know how an individual Deployer configures the Covered ADMT, the data the Deployer provides to it, the other systems or information the Deployer uses, or the role the Covered ADMT plays in a particular Consequential Decision. Conversely, a Deployer may understand its own decision-making process and use of a Covered ADMT but may not possess direct insight into the initial training of it. Deployers are also often the companies that have direct relationships with consumers.

The draft rules should recognize these practical limits. A company should not be required to disclose, verify, or retain information it does not know, possess, or control, or to make representations about another company's activities that it cannot independently verify. Similarly, companies should only be responsible for the risks associated with a Covered ADMT that are within their knowledge, possession, or control. Otherwise, companies would

either have to provide speculative or potentially inaccurate information, or conduct burdensome reviews of independent parties that are not workable in practice.

The Act itself appropriately reflects this principle. It requires Developers—which possess information about a covered ADMT’s design, intended uses, training data, and known limitations—to provide relevant documentation to Deployers. It separately requires Deployers—which make Consequential Decisions about consumers—to provide consumers with notices and information about the Covered ADMT’s use. To promote consistent implementation of the Act, the draft rules should preserve this role-based allocation by ensuring that each company is responsible only for information and risks within its knowledge, possession, or control. A company should not be required to obtain information from another actor or be held responsible for information provided by another company.

Recommendation: The draft rules should include a generally applicable provision stating that a Developer, Midstream Developer, or Deployer is only responsible for information and risks within its knowledge, possession, or control. The provision, which could include an express safe harbor, should permit reasonable reliance on information supplied by another party absent actual knowledge that the information is materially inaccurate or incomplete.

[Second, the draft rules should revise Rule 5.2 to include only information and risks within a Developer’s knowledge, possession, or control.](#)

Rule 5.2 seeks to clarify the documentation and disclosure obligations placed on Developers in section 6-1-1702 of the Act. Developers disclosing appropriate information to the next downstream actor about a Covered ADMT is critical to helping ensure that those actors can comply with their obligations and that consumers are protected when Consequential Decisions are made about them. Requiring companies to disclose information or be responsible for risks outside of their control risks compromising companies’ ability to comply with the Act’s obligations, as discussed above.

Several portions of Rule 5.2 concerningly extend beyond information and risks within a Developer’s knowledge, possession, or control:

- Rule 5.2(A) requires that any disclosures made by Developers must be written in a way that provides Deployers with a meaningful understanding and accurate expectations of the [Covered] ADMT’s known limitations, known risks, and circumstances in which the [Covered ADMT] should not be used. The “circumstances in which the [Covered] ADMT should not be used” will heavily depend on Deployer-specific factors affecting the risks associated with a particular Covered ADMT deployment, including, for example, the data used by the Deployer, modifications or configurations made by the Deployer, the domain in which the

Covered ADMT is used, the weight assigned to the Covered ADMT's output, the other information considered, and the level of human involvement in the Consequential Decision. The Developer will likely neither know nor control such Deployer-specific facts.

- Rule 5.2(C)(2) would direct Developers to provide instructions specifying methods to monitor the [Covered] ADMT, including, as appropriate, the collection and analysis of aggregate data and the analysis and review of an output that Materially Influenced a specific Consequential Decision. This requirement is concerning because appropriate monitoring methods often depend on the circumstances of a particular deployment. For example, a Developer will likely not know what data a Deployer provides to the Covered ADMT or how a Deployer configures the Covered ADMT. A Developer also will likely not have access to the aggregate data generated through the Deployer's use of the Covered ADMT.
- Rule 5.2(C)(3) would direct Developers to provide instructions for determining the reasoning and primary factors relied upon by the [Covered] ADMT to produce such an output, as well as instructions enabling the Deployer to instruct, control, modify, and adjust the [Covered] ADMT. This requirement is concerning because it requires a Developer to disclose information about deployment-dependent factors to which it likely does not have access. A Developer may be able to provide general information about how an ADMT produces its outputs, including information about the types of inputs used, the intended outputs, and known limitations. But determining the reasoning and primary factors underlying an output used in a specific Consequential Decision will likely depend on a Deployer's inputs, configuration, modifications, workflow, and use of the Covered ADMT output. Requiring Developers to provide decision-specific instructions could therefore require them to speculate about deployment circumstances they do not know or control.
- Rule 5.2 generally refers to an "ADMT," rather than a "Covered ADMT," as in section 6-1-1702 of the Act.

Recommendation: Rule 5.2 should be revised to:

- Remove "circumstances in which the [Covered] ADMT should not be used" from Rule 5.2(A);
- Require the disclosures in Rule 5.2(C)(2)-(3) to the extent technically feasible and reasonably available, based on information within a Developer's knowledge, possession, or control; and

- Change references to an “ADMT” to a “Covered ADMT” to be consistent with section 6-1-1702 of the Act.

Third, the disclosure requirements imposed on Developers should be limited to the content required in the Act and maintain flexible methods for Developers to provide such disclosures.

The Rulemaking Hearing Notice asks whether the Act’s required Developer disclosures include the information necessary to make those disclosures useful, whether Developers should be required to disclose additional information about known risks—including testing results—and whether the rules should prescribe the channels through which Developers provide information to Deployers.

The disclosures required by section 6-1-1702 already provide the next downstream actor with the information they need to understand and appropriately use a Covered ADMT. Among other requirements in the Act, Developers must provide information about the ADMT’s intended uses, categories of training data, known limitations, and instructions for appropriate use, monitoring, and Meaningful Human Review where applicable. These requirements appropriately focus Developer disclosures on information about the Covered ADMT that is known to and controlled by the Developer. Adding new categories of required information would risk producing longer and more technical disclosures without necessarily providing the next downstream actor with more useful information. The draft rules should instead focus on ensuring that the information required by the Act is relevant, understandable, and appropriately tailored to the Covered ADMT’s intended, documented, marketed, advertised, configured, or contracted uses.

The draft rules also should not require Developers to disclose testing results. Testing may involve numerous versions, configurations, environments, methodologies, and intended uses. Raw or detailed testing results may be difficult to interpret without that context and may not accurately predict how a Covered ADMT will perform when configured or used by a particular Deployer. Testing materials may also contain trade secrets, IP, security-sensitive information, or other confidential information. To the extent testing identifies a known limitation or risk relevant to an ADMT’s intended use, the Developer can describe that limitation or risk through the disclosures already required by the Act. Developers should not be required to disclose testing results, conduct new testing, or provide disclosures that are not relevant to the Covered ADMT’s intended, documented, marketed, advertised, configured, or contracted uses.

Additionally, the draft rules should not prescribe particular channels through which Developers or Midstream Developers must provide the required information downstream. Developers and Midstream Developers make Covered ADMTs available through a wide

variety of products, services, and commercial relationships. Depending on the circumstances, appropriate channels may include technical documentation, model or system cards, customer portals, websites, contractual materials, electronic files, or documentation integrated into a product or service. Prescribing a particular channel or format could interfere with established documentation practices, version control, access restrictions, and security measures. Developers and Midstream Developers should have flexibility to provide the required information through any method reasonably designed to make it accessible to the relevant downstream actor, including through secure portals or role-based access and according to confidentiality terms or reasonable use restrictions.

Recommendation: The draft rules should be revised to:

- Confirm that the disclosures required by section 6-1-1702 contain the categories of information Developers must provide;
- Decline to require Developers to disclose raw testing results, conduct new testing, or provide information beyond the known limitations and risks relevant to a Covered ADMT's intended, documented, marketed, advertised, configured, or contracted uses; and
- Permit Developers and Midstream Developers to provide the required information through any channel reasonably designed to make it accessible to the relevant downstream actor.

Fourth, the draft rules should clarify the role and responsibilities of Midstream Developers.

The draft rules define Midstream Developer as “a party that integrates a Covered ADMT as a component into its own covered ADMT product and provides its Covered ADMT product to another Developer or Deployer.” Rule 4.2 requires a Midstream Developer to reasonably obtain all Developer documentation for any Covered ADMT used as a component of its technology and to make all upstream Developer documentation available to downstream Developers and Deployers. BSA appreciates that the draft rules recognize different types of actors along the AI value chain and recommends further clarifications to the definition of Midstream Developer and the associated obligations to avoid confusion.

It is unclear which Developer obligations Midstream Developers must comply with, in addition to the standalone obligations placed on Midstream Developers in Rule 4.2, for their own Covered ADMT. For example, the draft rules' standalone obligations on Midstream Developers may not account for circumstances in which Midstream Developers substantially modify a Covered ADMT. In those circumstances, documentation from upstream developers alone may not represent the required information about the

Midstream Developer's substantially modified Covered ADMT. The draft rules should, therefore, clarify which Developer-facing obligations apply to a Midstream Developer's Covered ADMT, ensuring that Midstream Developers are only responsible for information within their knowledge and risks within their possession or control, and that Midstream Developers are only required to provide documentation from upstream developers to the next actor downstream to the extent such documentation is applicable to the Midstream Developer's Covered ADMT. Clarification is essential to helping ensure the Act establishes clear roles and responsibilities for different actors along the AI value chain and is workable in practice.

Additionally, Midstream Developers should not be obligated to "obtain" all Developer documentation for any Covered ADMT used as a component of its technology, as required in Rule 4.2(A)(1). Construing this requirement as applying to a Midstream Developer risks holding Midstream Developers responsible for a Developer's failure to provide the required documentation. The Act appropriately requires Developers to provide information about a Covered ADMT to Deployers in section 6-1-1702, and the draft rules should clarify that Developers are required to provide such information to Midstream Developers if a Midstream Developer is the next downstream actor to which the Developer provides its Covered ADMT.

The illustrative example provided in Rule 4.2(B) also presents several concerns, including:

- It is unclear how the third-party foundational LLM is in scope of the Act. The illustrative example seems to suggest that the Developer of the AI application must obtain all Developer documentation from the Developer of the third-party foundational LLM. However, the Act's disclosure requirements apply only when a Developer creates a Covered ADMT that is intended, documented, marketed, advertised, configured, or contracted to be used to make Consequential Decisions or when the Developer becomes aware that the Covered ADMT is being used to make Consequential Decisions in a manner consistent with the intended and contracted uses, per section 6-1-1702(5) of the Act. Therefore, for the foundational LLM developer to be subject to section 6-1-1702 of the Act, two conditions would have to be met: (1) the foundational LLM would have to meet the definition of Covered ADMT; and (2) the developer of the foundational LLM would have had to intended, documented, marketed, advertised, configured, or contracted for the foundational LLM to be used to make Consequential Decisions or become aware that the foundational LLM is being used to make Consequential Decisions in a manner consistent with its intended and contracted uses. Foundational LLMs are generally understood as general-purpose tools that could be used for a wide range of purposes, meaning they typically are not intended, documented, marketed, advertised, configured, or contracted to be used to make Consequential Decisions.

The illustrative example's focus on foundational LLMs risks creating significant confusion regarding the types of AI systems and companies in scope of the Act.

- It is unclear how the AI application is in scope of the Act. The illustrative example fails to detail how the developer of the AI application meets the definition of Midstream Developer, as the draft rules seem to intend, since the illustrative example simply reference that company's "AI system," rather than its Covered ADMT. The illustrative example's lack of explanation regarding how the AI application is a Covered ADMT and therefore in scope of the Act risks creating significant confusion regarding the role of Midstream Developers and the types of AI systems in scope of the Act.

Recommendation: The draft rules should be revised to:

- Clarify which of the Act's obligations placed on Developers apply to Midstream Developers for Covered ADMTs within the Midstream Developer's possession or control, in addition to any obligations directly placed on Midstream Developers in Rule 4.2, ensuring that Midstream Developers are only required to provide documentation from upstream developers to the next actor downstream to the extent such documentation is applicable to the Covered ADMT the Midstream Developer provides;
- Clarify that Developers must provide the disclosures required in section 6-1-1702 of the Act to the downstream actor to which the Developer provides its Covered ADMT, whether that downstream actor be a Midstream Developer or Deployer; and
- Strike the illustrative example in Rule 4.2(B) and replace it with clear examples of how Midstream Developers should implement their obligations with respect to tools that clearly meet the definition of Covered ADMT and are in scope of the Act.

Fifth, the draft rules should avoid unnecessary designations of additional roles, including that of ADMT Vendor.

The Rulemaking Hearing Notice asks whether the rules should recognize a separate "ADMT Vendor" role and impose obligations on vendors that operate ADMT on behalf of Deployers. BSA recommends against creating a new ADMT Vendor role. Defining a new, potentially overlapping role of ADMT Vendor risks creating uncertainty regarding which entity is responsible for which obligations under the Act. The Rulemaking Hearing Notice identifies an edge case that can better be understood through further clarifications and additional illustrative examples, rather than a new regulated role.

To help clarify circumstances in which Deployers rely on third-party vendors in using Covered ADMTs, the draft rules should expressly state that the definition of Deployer includes third parties that deploy a Covered ADMT on a Deployer's behalf and that, in those circumstances, the Deployer remains the party responsible for the obligations the Act places on Deployers. The draft rules should further clarify that a Deployer may contract with a third party to carry out its obligations under the Act.

Recommendation: The draft rules should not create a new defined role of ADMT Vendor. Instead, the draft rules should be revised to:

- Expressly state that the definition of Deployer includes third parties that deploy a Covered ADMT on a Deployer's behalf and that, in those circumstances, the Deployer remains the party responsible for the obligations the Act places on Deployer;
- Expressly state that Deployers may contract with third parties to carry out their obligations under the Act; and
- Include additional illustrative examples of how the above clarifications may work in practice.

II. The Draft Rules Should Establish Clear and Workable Thresholds for the AI Systems in Scope of the Act.

Clear thresholds for the types of AI systems in scope of the Act are critical to ensuring implementation works in practice. Companies must be able to clearly determine if their product or service is in scope to be able to understand and comply with their obligations. An overbroad or unclear scoping would create substantial uncertainty and risk hindering AI innovation in the state.

BSA recommends two changes to help ensure the draft rules create a clear threshold for the AI systems in scope of the Act.

First, the draft rules should generally focus on AI systems that are intended, documented, marketed, advertised, configured, or contracted to make Consequential Decisions, rather than AI systems that are simply used to make or Materially Influence Consequential Decisions.

The Act appropriately focuses on a Developer's intended uses in applying the disclosure requirements in section 6-1-1702 only when a Developer creates a Covered ADMT that is intended, documented, marketed, advertised, configured, or contracted to be used to make

Consequential Decisions or when the Developer becomes aware that the Covered ADMT is being used to make Consequential Decisions in a manner consistent with the intended and contracted uses, per section 6-1-1702(5) of the Act. These indicators give companies a reasonable basis for determining their obligations.

Focusing on downstream uses, conversely, risks creating significant uncertainty. General-purpose AI systems and common software tools can be applied to a wide range of activities at the discretion of downstream users. A downstream user might use an AI system to make a Consequential Decision even if the AI system was not intended, documented, marketed, advertised, configured, or contracted for that purpose. The upstream developer of the AI system may have no knowledge of that use and no ability to control the downstream user's data, configuration, decision-making process, or reliance on the AI system's output. Treating the upstream developer as a Developer of a Covered ADMT under the Act solely because of that downstream use would render the scope of the Act unknowable. The draft rules should avoid retroactively imposing obligations on any upstream developers based on conduct by an independent party downstream.

Furthermore, the draft rules should clarify that providing customer support documentation does not itself establish that a Developer intended a feature to be used to make Consequential Decisions. For example, Developers may provide downstream actors with certain documentation comply with legal requirements to support customer transparency or assist in responding to consumer-rights requests. Providing technical documentation, audit logs, an explanation interface, or data-export functionality should not alone be treated as evidence of a Developer's intended, documented, marketed, advertised, configured, or contracted use. Instead, the draft rules should rely on objective evidence showing a Covered ADMT is intended, documented, marketed, advertised, configured, or contracted to be used to make a Consequential Decision, per section 6-1-1702(5) of the Act.

Recommendation: The draft rules should clarify that the Act's Developer obligations apply when objective evidence shows that a Covered ADMT is intended, documented, marketed, advertised, configured, or contracted to be used to make a Consequential Decision, per section 6-1-1702(5) of the Act. Neither a downstream use outside those parameters nor providing customer support documentation should, by itself, transform an AI system into a Covered ADMT.

Second, the draft rules should incorporate the definition of Materially Influence in Standard 2 and improve upon it.

The Rulemaking Hearing Notice proposes two potential standards for determining whether an ADMT output is a De Minimis Factor in a Consequential Decision. Standard 1 would limit De Minimis Factors to those with a trifling, trivial, or incidental impact and would specify

that a factor is not De Minimis merely because other factors played a more significant role. Standard 2 would define a De Minimis Factor as one that is not a substantial factor in the Consequential Decision and would recognize that an ADMT output may be De Minimis when other relevant and independent information played a significantly larger role.

Standard 2 better reflects the Act's focus on ADMT that "materially" affects a Consequential Decision. The term "materially" must distinguish meaningful reliance on an ADMT output from situations in which the output is merely one minor item among substantial independent information. Standard 1's "trifling, trivial, or incidental" threshold could treat nearly any consideration of an ADMT output as material, even if the output played no meaningful role in the Consequential Decision. That interpretation would substantially diminish the limiting function of the statutory term "Materially Influence."

Standard 2 also more appropriately focuses on the ADMT output's actual role in the decision-making process. Relevant considerations include whether the decision-maker independently reviewed and analyzed other information, the relative weight given to that information and the ADMT output, whether the independent information supported the outcome, and whether the decision-maker had the knowledge and authority to reach or change the decision without relying on the ADMT output. These factors move towards a practical, fact-specific assessment of whether an ADMT output meaningfully affected a Consequential Decision.

Several improvements to Standard 2 are necessary, however, including:

- First, Standard 2 should consider whether the ADMT output was the controlling factor or principal basis in making a Consequential Decision. The fact that an ADMT output was consistent with the eventual outcome does not establish that it affected the outcome. Likewise, the fact that a decision-maker viewed an output does not establish that the decision-maker relied upon it. The relevant question is whether the output was a substantial factor that meaningfully altered the outcome of the Consequential Decision.
- Second, the draft rules should not presume Material Influence merely because an ADMT constrains an option set, sets a threshold, or generates a rank, score, classification, recommendation, prediction, or Other Inference. Many of those outputs are already included in the statutory definition of ADMT. Treating the form of an output as sufficient to establish Material Influence would largely collapse the distinction between an ADMT and an ADMT that Materially Influences a Consequential Decision. Any presumption should therefore depend on evidence that the output was actually used and meaningfully altered the outcome of the Consequential Decision. For example, a presumption may be appropriate when an

output automatically eliminates available outcomes, establishes a binding threshold, or receives a predetermined weight sufficient to determine the decision. An output should not meet the threshold of Materially Influence merely because it was generated, displayed, stored, or made available to a decision-maker.

- Third, Standard 2 should expressly permit a company to rebut any presumption by showing that the ADMT output was not consulted, considered, or incorporated into the Consequential Decision. The absence of reliance necessarily demonstrates that an output did not Materially Influence a decision. Evidence that the decision-maker reached and documented an independent judgment before receiving the ADMT output, never viewed the output, or did not incorporate it into the decision should therefore rebut the presumption. A company should also be able to rebut the presumption by showing that the output was a De Minimis Factor under Standard 2.
- Fourth, the draft rules should retain and clarify the proposed limitations concerning Other Inferences and administrative processing. “Other Inference” should remain limited to an opinion, characterization, or non-objective conclusion generated by an ADMT. It should not include the extraction, transfer, organization, formatting, summarization, or display of existing factual information or the performance of objective calculations. The administrative-processing exception should likewise cover tools that route, organize, format, transmit, store, retrieve, summarize, or check the completeness of information without evaluating an individual’s access to, eligibility for, or compensation related to a covered service or opportunity. A tool should not be considered to Materially Influence a Consequential Decision merely because it affects the timing, order, presentation, or administrative processing of information without meaningfully altering the outcome of the Consequential Decision.

Recommendation: The draft rules should incorporate the definition of Materially Influence in Standard 2 and improve upon it by:

- Requiring consideration of whether the ADMT output was the controlling factor or principal basis in making a Consequential Decision and meaningfully altered the outcome of the Consequential Decision;
- Limiting any presumption of Material Influence to circumstances in which an ADMT output was actually used and meaningfully altered the outcome of a Consequential Decision;
- Expressly permitting companies to rebut the presumption by demonstrating that the output was not consulted, considered, or incorporated into the Consequential Decision or was otherwise a De Minimis Factor; and

- Narrowly defining Other Inference and clarifying that routine administrative processing does not Materially Influence a Consequential Decision unless it meaningfully alters the outcome of the Consequential Decision.

III. The Draft Rules Should Extend the Express Protections for Trade Secrets to IP and Other Business Confidential Information.

Rule 5.3 seeks to clarify when Developers may withhold information from the disclosures Developers are required to provide to Deployers in section 6-1-1702 of the Act. Protecting companies' trade secrets, IP, and other confidential information is critical to ensuring that the Act's disclosure requirements work in practice.

BSA recommends two changes to the draft rules to help ensure Developers can comply with the Act's disclosure requirements without compromising their confidential or business proprietary information.

[First, the draft rules should extend the express protections for trade secrets to IP and other business confidential information.](#)

Rule 5.3(A) provides that Developers may withhold information from the disclosures Developers are required to provide to Deployers in section 6-1-1702 of the Act if that information:

- Is a trade secret, as defined by Colorado law; or
- Is protected from disclosure by state or federal law, or other legal authority.

The draft rules should extend the express protections for trade secrets in Rule 5.3(A)(1) to IP and other confidential information because long-standing federal and state laws protect companies' trade secrets, IP, and other confidential information. Therefore, companies' IP and other confidential information should be protected by Rule 5.3(A)(2). Clarifying expressly that companies may withhold IP and other confidential information from the disclosures required in section 6-1-1702 of the Act will help create a clear threshold for the types of information Developers may withhold and help ensure the Act is workable in practice.

Recommendation: Rule 5.3(A)(1) should be revised as follows: "Is a Trade Secret as defined by section 7-74-102(4), C.R.S., intellectual property, or other confidential information."

Second, the draft rules should clarify how Developers should comply with the Act's disclosure requirements when disclosure would reveal confidential or business proprietary information.

Rule 5.3(C) requires Developers to provide information sufficient to satisfy the disclosure requirements in section 6-1-1702 of the Act "in a way that does not reveal trade secrets" if the Developer determines that a disclosure would reveal the Developer's trade secrets. That requirement is concerning because it is unclear how a Developer can provide the required information without compromising the information that Rule 5.3 expressly permits the Developer to withhold.

Without further clarification, Rule 5.3(C) risks placing Developers in an untenable position: either provide information at a level of detail that reveals the very trade secrets, IP, or other confidential information that Rule 5.3 is intended to protect, or risk noncompliance by providing information that the Department later determines was insufficient. This concern is particularly acute where the required disclosures implicate proprietary details about how a technology operates, the data used to develop it, security measures, or other information that may be difficult to separate from the protected information.

The draft rules should make clear that a Developer may comply by providing an appropriately tailored summary, categorical description, or other generalized information that gives the next downstream actor a meaningful understanding of the relevant information without revealing the underlying confidential or business proprietary information. Whether that alternative information is sufficient should be evaluated under a reasonableness standard that accounts for the nature and sensitivity of the protected information, the information reasonably available to the Developer, and the information the next downstream actor needs to use the Covered ADMT appropriately. The rules should also recognize that, where no reasonable alternative disclosure can be provided without revealing protected information, the Developer may withhold that information consistent with Rule 5.3(A). These clarifications would give practical effect to Rule 5.3's protections and provide Developers with a workable path for complying with the Act.

Recommendation: Rule 5.3(C) should be revised to clarify that a Developer may satisfy the applicable disclosure requirements by providing a summary, categorical description, or other generalized information that is reasonably sufficient to support the next downstream actor's appropriate use of the Covered ADMT without revealing trade secrets, intellectual property, or other confidential information. If no such alternative can reasonably be provided, the Developer should be permitted to withhold the protected information consistent with Rule 5.3(A).

* * *

Thank you for allowing us to provide feedback on behalf of the enterprise software sector and for your consideration of our views. We welcome the opportunity to further engage with you or a member of your staff on these important issues.

Sincerely,

Meghan Pensyl
Director, Policy