CIPL's Response to the to the Colorado Department of Law on the October 6, 2026 Interim Draft Automated Decision-Making Technology and Conversational Artificial Intelligence Service Rules

Submitted 23 October 2026

The Centre for Information Policy Leadership (CIPL)[1] welcomes the opportunity to comment on the Colorado Department of Law's October 6, 2026 interim draft rules. The first part of our response focuses on draft Rules 1–8, which implements the Automated Decision-Making Technology in Consequential Decisions Act (the "ADMT Act"), and builds upon CIPL's July 13, 2026 pre-rulemaking submission.[2] The second part of our response principally addresses draft Rules 9–14, which implements the Conversational Artificial Intelligence Services Act (the "Chatbot Safety Act").

I. ADMT Provisions

CIPL appreciates the Department's efforts to clarify the ADMT Act's requirements and develop workable processes for consumer disclosures, access, correction, and meaningful human review. The draft rules address several recommendations in CIPL's earlier submission, including providing objective applicability indicators, establishing consumer-rights response timelines, explaining commercial reasonableness, clarifying responsibilities in multiparty arrangements, and supplying practical examples.

These improvements can support effective consumer protection while reducing uncertainty for organizations implementing responsible ADMT. CIPL nevertheless remains concerned that the rules as drafted could capture technology that merely informs genuine human judgment, and that certain

---

[1] **The Centre for Information Policy Leadership (CIPL)** is a global privacy and data policy think tank within the Hunton law firm that is financially supported by the firm, 80+ member companies that are leaders in key sectors of the global economy, and other private and public sector stakeholders through consulting and advisory projects. CIPL's mission is to engage in thought leadership and develop best practices for the responsible and beneficial use of data in the modern information age. CIPL's work facilitates constructive engagement between business leaders, data governance and security professionals, regulators, and policymakers around the world. For more information, please see CIPL's website at https://www.informationpolicycentre.com/. Nothing in this document should be construed as representing the views of any individual CIPL member company or Hunton. This document is not designed to be and should not be taken as legal advice.

[2] Response by the Centre for Information Policy Leadership to the Colorado Department of Law Pre-Rulemaking Request for Comments under the Automated Decision-Making Technology Act, submitted July 13, 2026, available at https://www.informationpolicycentre.com/resources/cipl-response-to-the-colorado-department-of-law-pre-rulemaking-request-for-comments-under-the-automated-decision-making-technology-act/.

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

disclosure and procedural requirements could impose unnecessary burdens without corresponding consumer benefits.

Regarding the ADMT provisions, CIPL makes the following recommendations, explained more fully below:

1. Build on the draft rules' alignment with existing Colorado privacy requirements
2. Refine "Materially Influence" to distinguish genuine human judgment from automated decision-making
3. Preserve exclusions for low-risk administrative tools and further clarify consequential effects
4. Retain the 45-day response framework and add appropriate extensions
5. Preserve the commercial-reasonableness framework for the right to Meaningful Human Review
6. Make notice responsibilities explicit in multiparty arrangements
7. Reconsider named-source disclosures and strengthen protections for sensitive information
8. Expand federal-law coordination to avoid duplicative disclosures
9. Convert illustrative disclosures into optional, reliable compliance tools
10. Correct numbering and cross-reference inconsistencies before finalization

---

1. Build on the draft rules' alignment with existing Colorado privacy requirements

CIPL commends the Department for connecting the proposed access and correction procedures to existing CPA requirements. The draft rules' use of CPA definitions, security standards, authentication principles, and appeal procedures is consistent with CIPL's recommendation to avoid overlapping and potentially contradictory compliance frameworks.

CIPL also welcomes provisions allowing organizations to use existing procedures rather than create duplicative systems. For example, draft Rule 8.8 recognizes qualifying FERPA processes, while Rule 5.2(E) permits an organization acting as both Developer and Deployer to maintain the required information in a unified internal documentation system.

The Department should extend this approach by expressly clarifying that Deployers may use existing CPA request channels, authentication processes, response systems, and appeals mechanisms to fulfill corresponding ADMT requirements, provided those mechanisms supply the information and functionality required by the ADMT Act. Consumers should receive clear explanations of their ADMT rights without organizations being required to build a separate administrative infrastructure.

More fundamentally, CIPL reiterates its recommendation that the Department explain how the ADMT framework relates to the CPA Rules' regulation of automated decision-making involving Human Involved Automated Processing, Human Reviewed Automated Processing, and Solely Automated Processing. A

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

clear relationship between these frameworks would improve consistency and reduce conflicting interpretations of the same decision-making process.

2. *Refine "Materially Influence" to distinguish genuine human judgment from automated decision-making*

CIPL appreciates the Department's effort to establish objective indicators of Material Influence. Draft Rule 6.1 identifies the circumstances triggering a rebuttable presumption and the evidence that may rebut it.[3] The definition of "De Minimis Factor" also identifies relevant considerations, including independent evidence, subject-matter understanding, and actual decision-making authority. These provisions respond to CIPL's request for greater operational clarity.

However, the draft rules do not adopt CIPL's central recommendation that coverage distinguish processes in which automation replaces human judgment from those in which technology merely assists meaningful human decision-making.

Under draft Rule 6.1, an individual-specific output may be presumed to Materially Influence a Consequential Decision when it is reviewed during the decision-making process and is consistent with the outcome. The employment examples under draft Rule 6.2(C) further indicate that this presumption can apply even where the decision-maker reviews underlying records and interviews candidates.

Alignment between a human decision and an automated output should not, by itself, establish that the output displaced or substantially determined human judgment. Both may reflect the same underlying evidence.

To the extent consistent with the statute, the Department should:

- Align Material Influence with Solely Automated Processing and Human Reviewed Automated Processing under the CPA Rules, as recommended in CIPL's July submission.
- Clarify that genuine human decision-making falls outside the scope of the rules where the individual actually evaluates relevant evidence, exercises independent judgment, and has authority to change the outcome.
- Distinguish nominal human oversight from substantive human involvement, rather than treating the presence or absence of a human reviewer as dispositive.

If the Department retains the proposed presumption, it should clarify that reviewing an output and reaching a consistent decision does not alone resolve whether the output was a non-de minimis factor.

---

[3] "The attorney general may adopt rules to clarify the application of the definition of "Materially Influence", as defined in section 6-1-1701 (13), *including presumptions*, illustrative examples, and objective indicators." Colo. Rev. Stat. § 6-1-1706(5)(b) (emphasis added).

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

The rules should also avoid implying that organizations must create extensive new records for every technology-assisted human decision solely to demonstrate that automation was not determinative.

3. Preserve exclusions for low-risk administrative tools and further clarify consequential effects

CIPL welcomes draft Rule 6.1(D)'s explanation that tools used solely to summarize, organize, translate, draft, route, or present information for human review or administrative processing are excluded from the ADMT definition.

The examples involving biology-class time slots and objective organization of tenant applications are useful illustrations of activities that should not trigger heightened ADMT obligations. These provisions are consistent with CIPL's concern that routine, low-risk computational assistance should not be treated as consequential automated decision-making.

The Department should build on this clarification by explaining the relationship between a "Consequential Decision" under the ADMT Act and a decision producing "legal or similarly significant effects" under the CPA. To the extent authorized by the statute, the rules should align these concepts and clarify that routine operational choices without meaningful effects on access, eligibility, or opportunity do not trigger heightened requirements.

Additional examples should distinguish substantive judgments about individuals from objective retrieval, formatting, calculation, and organization of information that humans use to make their own decisions.

4. Retain the 45-day response framework and add appropriate extensions

CIPL commends the Department for establishing 45-day response periods for access and correction requests and for completing meaningful human review. These periods reflect an important element of CIPL's recommendation and provide useful predictability.

The relevant provisions do not, however, expressly include CIPL's recommended additional 45-day extension when reasonably necessary. The Department should provide that extension, consistent with the CPA approach and to the extent permitted by the ADMT Act.

An extension should require notice within the initial response period explaining the reason and the expected completion date. Relevant circumstances may include request complexity, the volume of requests, the need to obtain information from third parties, and the availability of qualified reviewers.

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

The Department should also clarify how different deadlines interact. The draft rules establish separate periods for additional ADMT information,[4] third-party access or correction instructions,[5] substantive access and correction,[6] and human review.[7] A consumer may request several of these simultaneously or seek review that depends on a pending correction. The rules should explain how deployers may coordinate those responses without unnecessary duplication or premature review based on disputed data.

5. Preserve the commercial-reasonableness framework for the right to Meaningful Human Review

CIPL welcomes draft Rule 8.7(E)'s multifactor approach as to what constitutes "commercially reasonable" for human review of adverse decisions. The factors concerning the type and value of review, magnitude and reversibility of harm, organizational size and capacity, cost, technical feasibility, and reviewer availability provide much needed guidance.

The accompanying examples are also valuable because they demonstrate how the analysis depends on the consumer's circumstances and the review that could affect the outcome.

That said, we request that the list of factors be deemed non-exhaustive. Moreover, the Department should consider adding **the volume of decisions and review requests, and the cumulative operational burden of providing review**, to the relevant considerations. CIPL's July submission identified decision volume alongside organizational size. Assessing only the marginal cost of one review may not fully reflect feasibility across a large number of requests.

The Department should also clarify that deployers may use proportionate review processes suited to the issue raised. Review of relevant primary evidence may resolve a request without requiring a system-wide technical investigation. Conversely, a technical issue may warrant targeted investigation or correction. The final rules should preserve this flexibility while requiring genuine consideration of the consumer's evidence and circumstances.

6. Make notice responsibilities explicit in multiparty arrangements

CIPL commends draft Rule 4.3 for clarifying that, where a third party runs Covered ADMT on behalf of a Deployer and the Deployer makes the Consequential Decision using the output, the Deployer retains responsibility for Deployer obligations. Further, the ability of the Deployer to contract with the third party

---

4 Fifteen days, per draft Rule 7.6(A).

5 Fifteen days, per draft Rule 8.3(C).

6 Forty-five days, per draft Rule 8.3(B)(1) and draft Rule 8.4(B)(1).

7 Ten days, per draft Rule 8.7(G).

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

to require the third party to assist the Deployer in satisfying its Deployer obligations is a practical and useful clarification.

This approach substantially addresses CIPL's concern about allocating responsibilities between consumer-facing decision-makers and B2B service providers. That said, it is still unclear how the pre-use ADMT notice requirements in Sections 6-1-1704(1)-(3) apply to Deployers of ADMT that are B2B providers.

The Department should clarify that a provider does not become independently responsible for consumer notice solely because it operates technology for another organization. Where the provider acts only on behalf of the party making the Consequential Decision, that decision-maker should be responsible for ensuring the required notice is provided, with contractual assistance where appropriate.

This clarification should remain sensitive to actual roles: a provider that independently makes a covered Consequential Decision may have its own deployer responsibilities. The objective is clear allocation, not a blanket exemption for all B2B providers.

More broadly, the Department should continue to allocate obligations based on an actor's actual role, level of control, and ability to monitor and mitigate relevant risks. As AI systems increasingly involve multiple participants, including developers, deployers, integrators, and customers, regulatory requirements should attach to the entity that exercises meaningful control over the activity creating risk and is best positioned to address it. A role- and control-based framework better promotes accountability, reduces duplicative obligations, and ensures that compliance requirements remain effective as AI ecosystems evolve.

7. Reconsider named-source disclosures and strengthen protections for sensitive information

CIPL appreciates the draft rules' recognition of trade-secret protections and restrictions relating to cybersecurity, fraud prevention, anti-money laundering, counter-terrorist financing, and sanctions compliance.

Nevertheless, draft Rule 7.6(A)(5) does not adopt CIPL's recommendation that "sources" be disclosed as categories. Instead, it requires known sources to be identified by name, including known original sources and intermediaries, and presumes deployer knowledge in specified circumstances.

CIPL remains concerned that mandatory named-source disclosures could expose competitively sensitive relationships, proprietary data arrangements, or information relevant to security and fraud-prevention systems.

The Department should permit meaningful categorical descriptions—such as credit bureau data, employer-provided records, publicly available court records, or information supplied directly by the consumer—where those descriptions adequately explain the source and more specific disclosure would create legitimate confidentiality or security concerns.

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

The rules should distinguish between:

- General source information needed to understand an adverse outcome; and
- Specific third-party identification or contact information needed to access or correct disputed personal data.

Named-source information may be particularly useful in the second context. It should not automatically be required in every general disclosure.

The Department should also qualify presumptions of knowledge by the information reasonably available to the deployer and its ability to obtain that information. Deployers should provide meaningful information they possess or can reasonably obtain, without being deemed to know every underlying source used within a vendor's system.

8. Expand federal-law coordination to avoid duplicative disclosures

CIPL welcomes draft Rule 7.7's recognition that Equal Credit Opportunity Act (ECOA) and Fair Credit Reporting Act (FCRA) notices may satisfy relevant disclosure requirements when accompanied by the specified additional ADMT information. Allowing supplemental information to be delivered with existing notices is consistent with CIPL's recommendation to facilitate interoperable compliance.

The draft rules does not, however, adopt CIPL's requested broader relief for organizations subject to the Gramm-Leach-Bliley Act (GLBA) and the FCRA (aside from creditors).

CIPL reiterates our request to the extent the Department has statutory authority to implement it. Where an entity-level exemption cannot be established through rulemaking, the Department should provide the broadest legally permissible recognition of equivalent federal disclosures and compliance processes.

In particular, the rules should make clear that an organization may provide one integrated notice satisfying all applicable requirements, rather than a separate Colorado notice. Additional ADMT information should address actual gaps in the existing disclosure, not repeat information already supplied. Coordination should also preserve federal confidentiality obligations and security safeguards.

9. Convert illustrative disclosures into optional, reliable compliance tools

CIPL commends the Department for including sector-specific examples and sample disclosure language. These provisions respond directly to CIPL's request for practical guidance.

However, Draft Rule 7.5 states that the examples are illustrative and are not intended to serve as the basis for Deployer decisions regarding disclosures. They therefore do not provide the compliance certainty that optional model notices or safe harbors could offer.

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

The Department should publish optional model notices covering the required disclosure elements, including the role of ADMT, requests for additional information, and consumer rights. To the extent authorized, accurate use of a model notice adapted to the relevant facts should create a safe harbor or rebuttable presumption of compliance with the requirements the model addresses.

Organizations should remain free to use alternative notices that satisfy the rules. Standardized models would be particularly helpful for smaller organizations and could improve consumer understanding through consistent presentation.

10. Correct numbering and cross-reference inconsistencies before finalization

The interim draft contains numbering and cross-reference inconsistencies associated with the insertion and renumbering of rules. For example, the new Rule 8.8 continues to reference several Rule 7 consumer-rights provisions, while other references similarly retain earlier numbering.

CIPL encourages the Department to reconcile these references and publish a clean version before final adoption. Reliable cross-references are important because the disclosure, request, authentication, and review provisions operate together.

II. Conversational AI Provisions

CIPL commends the Department's commitment to establishing a risk-focused, modern governance framework that protects Colorado consumers—particularly minors—while enabling responsible innovation. We strongly welcome the Department's recognition of data minimization, privacy-preserving age-assurance architectures, demonstrable accountability, and international standards such as ISO/IEC 27566-1:2025. To further enhance the rules' effectiveness and long-term durability, CIPL offers constructive recommendations to ensure that the scope of conversational services remains appropriately targeted, transparency requirements prioritize genuine user comprehension over repetitive notices, crisis-referral reporting metrics remain scientifically sound and privacy-protective, and compliance processes achieve practical interoperability with the Colorado Privacy Act and global best practices.

Specifically, CIPL makes the following recommendations, explained more fully below:

1. Build on the strong standards-based foundation for age assurance by ensuring public accessibility and risk-calibrated obligations
2. Preserve the proportionate scope framework
3. Preserve the prohibition on false representation while transitioning to outcomes-based requirements
4. Maintain evidence-based crisis detection standards and protect security-sensitive compliance information

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

5. Replace categorical design prohibitions with a risk-based, mitigation-centered framework and publish clear guidance on reasonable measures
6. Preserve tiered accountability and reporting obligations scaled to operator size while distinguishing public and confidential report elements

---

1. ***Build on the strong standards-based foundation for age assurance by ensuring public accessibility and risk-calibrated obligations.***

CIPL commends the Department for Draft Rule 10's standards-based approach to age assurance. Rule 10 provides that, in assessing whether an age assurance method is generally accepted, the Department may consider conformance with applicable requirements of internationally recognized standards, including ISO/IEC 27566-1:2025 and IEEE 2089.1, both incorporated by reference under draft Rule 15.

CIPL has consistently advocated for international standards-based approaches to age assurance. This approach promotes regulatory convergence across jurisdictions, reduces duplicative compliance burdens for global operators, and ensures that age assurance methods can keep pace with technological advances through standards bodies rather than requiring repeated rulemaking updates.

CIPL further commends the Department for making user choice a core compliance criterion in draft Rule 10 and expressly recognizing that age assurance methods may be used alone or in combination. Rule 10 requires that a commercially reasonable or generally accepted age assurance method provide reasonable user choice in how the user's age can be estimated, and expressly provides that an Operator must not use as a sole method of age assurance a method that requires a user to provide government-issued identification.

CIPL has consistently argued that there is no single, universally correct approach to age assurance. The appropriate mechanism and the point in the user journey at which it is applied should depend on the nature of the service, the features offered, and the likelihood and severity of the risks to children associated with that service. Regulatory frameworks should therefore allow a portfolio of options rather than mandating a single technique.

That said, the Department should clarify that providing "reasonable user choice" does not **require** Operators to offer more than one method (except where the Operator offers a method relying on government-issued identification), as multiple methods may not be possible for all services.

CIPL has repeatedly raised concern that exclusive reliance on government-issued identification excludes individuals who lack such documentation (including many minors and young adults), creates disproportionate data collection relative to the purpose of age estimation, and is inconsistent with privacy-by-design principles.

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

Rule 10 provides that a commercially reasonable or generally accepted method does not solely rely on user self-declarations of age; general contractual restrictions or terms of use; or online payment methods that are available to minors. CIPL agrees that unaided self-declaration is an inadequate safeguard where age assurance is warranted, particularly in higher-risk service contexts. At the same time, CIPL has emphasized that the level of assurance should be calibrated to risk, and that self-declaration may form one component of a multi-signal approach. The rule's prohibition on sole reliance on self-declaration — while permitting its use as one element among others — strikes an appropriate balance.

CIPL commends the Department's inclusion of this contextual dimension as a meaningful recognition of the distinctive characteristics of conversational AI. Draft Rule 10 includes consideration of whether Operators take reasonable steps to re-assess age estimates based on new signals, including contextual information derived through conversations and behavioral indicators, and further provides that willful disregard analysis will consider whether the Operator failed to consider and act on age-related signals reasonably available through a user's interactions with the CAIS — including, when considered in context, a user's direct statement regarding age or minor status, repeated references to grade-level or school situations, or recurring interactions in which the user identifies themselves as a child in described parent-child interactions. CIPL has acknowledged the legitimacy of behavioral and contextual signals as a component of a holistic age assurance approach, particularly for services that involve natural language interaction. For conversational AI systems specifically, dialogue itself can yield age-relevant contextual cues that other service types may not generate.

Enhance public accessibility of standards

Building on this strong foundation, **CIPL respectfully recommends that the Department enhance public accessibility of the standards incorporated by reference**. Draft Rule 10 incorporates by reference ISO/IEC 27566-1:2025 and IEEE 2089.1, and draft Rule 15 confirms that these standards are available for purchase from ISO and IEEE, respectively, while also noting that the Colorado Department of Law maintains copies available for public inspection at its Denver office during regular business hours. CIPL strongly supports the substantive decision to reference these internationally recognized standards. However, the fact that the full text of the incorporated standards is not freely and publicly available online creates meaningful compliance barriers. Regulated entities — particularly smaller Operators, academic institutions, civil society organizations, and members of the public seeking to understand the regulatory requirements — must either purchase access or travel to Denver to inspect physical copies. This arrangement does not incentivize or support the practical ability of many stakeholders to ascertain the precise substance of legal requirements binding upon them, and it raises democratic accountability concerns when the content of incorporated law is not readily accessible without payment or physical presence. CIPL encourages the Department to work with ISO and IEEE to make the cited standards or official summaries available at no cost as part of the regulatory process. Enhancing public accessibility of incorporated standards would strengthen both compliance and the legitimacy of the regulatory framework, without diminishing the value of relying on internationally recognized benchmarks.

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

Calibrate age assurance obligations to the risks

To further advance effective and proportionate age assurance, **CIPL respectfully recommends that the Department calibrate obligations to the risks of each service**. Draft Rule 10's age assurance requirements apply through a uniform multi-criterion compliance framework to Operators subject to the Chatbot Safety Act's age estimation duty, without explicit differentiation based on the risk profile of the particular CAIS. Not all conversational AI services pose equivalent risks to minor users. A companionship chatbot designed for sustained emotional engagement presents materially different risks than a general-purpose productivity assistant or a narrowly scoped informational tool that happens to fall within the CAIS definition. Applying the full suite of Rule 10 criteria uniformly to all covered services, without a risk-calibration mechanism, may impose disproportionate burdens on lower-risk services while failing to provide tailored, heightened protections for higher-risk ones.

CIPL has consistently held that the deployment of age assurance measures must follow a risk-based approach, must be context-specific, and must provide organizations with flexibility to implement measures appropriate to their services. More data-intrusive methods should be reserved for higher-risk contexts. **CIPL respectfully recommends that the Department require CAIS Operators to conduct a preliminary risk assessment of the risks their service poses to minor users, and to calibrate their age assurance obligations to the assessed risk level**. The rule should establish tiered age assurance requirements — with proportionately more robust requirements for higher-risk services (for example, those designed for emotional companionship or prolonged multi-turn engagement with minors) and a lighter compliance pathway for demonstrably lower-risk CAIS that would not pose meaningful harm to minor users, including an assessment of whether the service is even accessed by minors. The risk assessment should consider the nature of the service, the features and functionalities offered, the likelihood and severity of potential harms to minors, and the benefits the service delivers. This risk-based approach is consistent with CIPL's longstanding advocacy for proportionate, risk-calibrated AI regulation and with the Department's own multi-factor, totality-of-the-circumstances methodology already employed elsewhere in the rules (including in draft Rule 9's emotional companionship analysis and draft Rule 10's commercially reasonable method analysis, which already considers Operator size and financial resources).

Support privacy by design

CIPL commends the Department for draft Rule 10's privacy-by-design framing. Rule 10 lists as its first evaluation criterion that an age assurance method must protect a user's privacy by collecting only data necessary for the purpose of age assurance of the specific user, not sharing or using that data for any other purpose except as required by law, and deleting collected data after the minimum time required for compliance. The rule further provides that obtaining more precise personal data than required is not evidence of a more reasonable determination.

CIPL has long championed privacy-by-design and data minimization as foundational principles for all digital services that handle personal data, and has emphasized these principles specifically in the age assurance

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

context. Elevating privacy-by-design to the lead criterion in draft Rule 10 reflects CIPL's view that child protection and privacy are complementary, not competing, values.

CIPL's commitment to privacy-by-design extends to its active work on Privacy-Enhancing Technologies ("PETs")[8] as a concrete mechanism for operationalizing that principle in the age assurance context. CIPL has consistently urged regulators to recognize the value of PETs and to create affirmative incentives for their adoption. CIPL respectfully encourages the Department to consider supplementing Rule 10 with guidance that explicitly recognizes the use of PETs as a favorable indicator of a commercially reasonable and privacy-protective age assurance method, and that incentivizes Operators to invest in PETs-based solutions.

Provide guidance on aging-up

CIPL welcomes the aging-up requirement and respectfully recommends that the Department also provide substantive guidance on its implementation. Draft Rule 10 requires that a commercially reasonable or generally accepted age assurance method provide minors a process to update their minor age status upon reaching the age of 18. This "aging-up" requirement is an important recognition that age-based protections should not persist indefinitely after a user reaches adulthood, and that users should have a pathway to transition out of minor-protective settings.

However, the rule provides no substantive guidance on what this process must entail, what outcomes it must achieve, how the Operator determines that a user has reached adulthood, how data collected under minor-protective settings should be handled upon aging-up, or what transitional notifications should be provided. Such guidance, and especially examples, will help Operators understand their obligations and ensure that the requirement is implemented in consistent and adequate ways — neither over-restricting adult users who have aged up or under-protecting the integrity of the transition.

CIPL respectfully recommends that the Department provide regulatory guidance on the aging-up process that at minimum addresses: (a) the trigger and verification mechanism for aging-up (how the Operator determines a user has reached adulthood, consistent with the privacy-preserving and multi-method principles of draft Rule 10); (b) the data handling obligations for data collected under minor-protective settings prior to aging-up, including whether and how such data may be retained, used, or subjected to renewed user choice upon transition; (c) the process by which minor-specific restrictions on data use, retention, training, and engagement features are lifted or modified upon aging-up; and (d) any transitional period or user notification obligations, so that users understand the change in their protections and settings.

---

[8] CIPL's work on PETs is available at https://www.informationpolicycentre.com/project/privacy-enhancing-technologies-pets/.

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

The Department should specify these requirements as outcome-based expectations rather than prescriptive mandates, to preserve Operator flexibility while ensuring that aging-up is meaningful, transparent, and consistent with the user's evolving autonomy — a principle CIPL has emphasized in noting that safety measures should take into account children's evolving capacities.

2. Preserve the proportionate scope framework

CIPL commends the Department for draft Rule 9's exemption as a meaningful effort to calibrate regulatory burden to actual risk. Rule 9 clarifies that when determining whether a CAIS output relates to a "narrow and discrete topic," the Department will consider whether the output was generated from completion of a small-scale, narrow, and well-defined task — such as a status update about a purchase, answering a billing question, or scheduling an appointment — and includes menu- or button-based and protocol- and rule-based chatbots, provided they cannot generate outputs related to sexually explicit conduct or engage in dialogue related to suicidal ideation or self-harm.

CIPL has consistently recommended a proportionate, risk-based scoping of AI regulation that excludes low-risk, bounded applications from heightened obligations designed for general-purpose or emotionally interactive systems. Over-inclusive scoping risks imposing compliance burdens that are disproportionate to actual harm and may discourage beneficial deployment of simple, task-oriented conversational tools.

CIPL further commends the Department for draft Rule 9's contextual approach to emotional companionship and emotionally dependent interaction. Rule 9 defines "designed to simulate emotional companionship" and "designed to encourage emotionally dependent interaction" through multi-factor, totality-of-the-circumstances analyses. For emotional companionship, the Department will consider factors including anthropomorphic design features (personalized or adaptive dialogue; assignment of name, gender, avatar, or fictional backstory); prioritization or failure to address sycophantic behavior and simulated empathy; and the ability to recall user characteristics and past conversations in a manner that conveys continuing personal familiarity. Critically, the rule provides that one of these factors alone does not, by itself, constitute simulated emotional companionship. For emotionally dependent interaction, the analysis considers reinforcement loops, failure to include or test real-world-support safeguards, failure to ensure safeguard effectiveness following model updates, and failure to implement remedial measures. CIPL has recommended that assessments of AI-induced psychological risk be grounded in contextual, evidence-based analysis rather than categorical presumptions based on any single design feature. This multi-factor approach is consistent with that recommendation and reduces the risk that beneficial personalization features will be categorically treated as harmful.

3. Preserve the prohibition on false representation while transitioning to outcomes-based requirements

CIPL commends the Department for draft Rule 13's prohibition on false representation as an important consumer protection measure. Rule 13 clarifies the prohibition against Operators using terms, letters, or

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

phrases in advertising, interfaces, or outputs that state or imply that CAIS output is provided by, endorsed by, or equivalent to services provided by a licensed health-care professional, licensed legal professional, licensed mental health professional, or qualified dietitian. The rule specifically contemplates controls reasonably designed to prevent creation of CAIS profiles with names including terms such as "therapist," "psychologist," "doctor," "lawyer," or "dietitian" where the name is likely to cause a reasonable user to believe they are interacting with a licensed professional. CIPL supports prohibiting deceptive practices that could lead users to place inappropriate reliance on AI outputs — particularly in domains involving health, mental health, legal advice, and nutrition, where misplaced trust could cause material harm.

Modify disclosure obligations to be outcomes-based

While CIPL strongly supports effective disclosure of the AI nature of these services, **CIPL respectfully recommends that the Department adopt outcomes-based disclosure obligations**. Draft Rule 11 prescribes specific disclosure requirements, including — for non-minor users — disclosure at the beginning of a user's first interaction each calendar day, at least once every three hours during a Continuous Conversational Artificial Intelligence Service Interaction (or as a persistent visible disclosure), and in response to any user inquiry about whether the service is artificially generated. For minor users, the rule requires a persistent visible disclaimer on screen interfaces and intermittent audio disclaimers on non-screen interfaces, with detailed specifications as to visibility, font size, language, volume, and repetition.

CIPL fully supports the underlying objective: users must have clear, genuine, and effective awareness that they are interacting with artificial intelligence and not a human. However, mandating specific temporal intervals is a highly prescriptive approach that may not be the most effective means of achieving that objective. Users may habituate to rote disclosures delivered at fixed intervals, reducing their informational value over time. Rigid interval requirements may also disrupt beneficial or sensitive interactions at inopportune moments. CIPL has consistently urged regulators to adopt outcome-driven rather than prescriptive design requirements, so that obligations do not become obsolete as technology and user expectations evolve and so that Operators retain flexibility to determine what solutions work best for their services and users. CIPL would encourage the Department to consider a separate consultation on UX design.

CIPL respectfully recommends that the Department reframe draft Rule 11 around the intended outcome — users should have clear, genuine, and effective awareness of the AI nature of the chatbot and of their applicable protections — and allow CAIS Operators flexibility to achieve that outcome in ways appropriate to their service context. The rule should specify what must be disclosed and that disclosure must be effective, unambiguous, and consistent with draft Rule 3's general communication requirements, while permitting Operators to determine the most contextually appropriate timing, frequency, and format of disclosure.

CIPL notes with appreciation that draft Rule 11 already incorporates testing-based factors for evaluating the reasonableness of audio disclaimer intervals (including whether the Operator tested users'

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

understanding and whether disclosures are delivered at meaningful points). Extending this outcomes-and-testing orientation more broadly — including by requiring Operators to demonstrate through user testing and research that their disclosure practices are effective — would better serve the rule's protective purpose while remaining technology-neutral. This approach is consistent with CIPL's advocacy for technology-neutral, outcomes-based regulation.

4. Maintain evidence-based crisis detection standards and protect security-sensitive compliance information

CIPL commends the Department for grounding crisis detection obligations in clinically validated tools. Draft Rule 2 defines "Evidence-Based Method," as used in connection with measuring suicidal ideation or self-harm under the Chatbot Safety Act, to mean a method that has been validated through empirical research and scientifically tested for reliability, validity, sensitivity, and specificity. The definition expressly contemplates validated clinical tools including the Columbia-Suicide Severity Rating Scale (C-SSRS), the Patient Health Questionnaire item 9 (PHQ-9 #9), the Ask Suicide-Screening Questions (ASQ), and structured frameworks such as the SAFE-T protocol.

Regulatory requirements for detecting sensitive behavioral states should be grounded in validated, evidence-based methodologies. Anchoring crisis detection obligations in clinically validated instruments promotes consistency, measurability, and accountability, and reduces the risk that Operators will adopt ad hoc or performative detection approaches.

Annual report should distinguish public and confidential elements of crisis detection information

Building on the evidence-based detection framework, **CIPL respectfully recommends that the Department distinguish public and confidential crisis detection information in annual reports**. Draft Rule 14 requires detailed annual reports describing crisis referral protocols and suicidal ideation and self-harm detection methodologies, including thresholds and criteria for human review. If publicly disclosed in full, these reports could expose security-sensitive thresholds, triggers, and methodologies that bad actors could exploit to evade detection or manipulate systems designed to protect vulnerable users. Although the draft rules reference a trade-secret limitation and require aggregate reporting without user identifiers, they do not clearly distinguish public-facing information from security-sensitive information that should be submitted in confidence.

CIPL respectfully recommends that the Department amend draft Rule 14 to establish clear categories: (a) publicly disclosed elements, such as aggregate crisis-referral metrics and high-level methodology descriptions; and (b) elements submitted in confidence to the Attorney General's office only, including detailed technical specifications of crisis detection systems and specific thresholds and triggers. The rule should require the Attorney General's office to maintain confidentiality of the latter category except where disclosure is required for regulatory or law enforcement purposes, and provide a clear process for Operators to designate confidential material at submission. This approach maintains evidence-based

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

detection and regulatory accountability without compromising the integrity of safeguards for vulnerable users.

5. *Replace categorical design presumptions with a risk-based, mitigation-centered framework and publish clear guidance on reasonable measures*

CIPL commends the Department for draft Rule 9's contextual approach to emotional companionship and emotionally dependent interaction. Rule 9 defines "designed to simulate emotional companionship" and "designed to encourage emotionally dependent interaction" through multi-factor, totality-of-the-circumstances analyses. For emotional companionship, the Department will consider factors including anthropomorphic design features (personalized or adaptive dialogue; assignment of name, gender, avatar, or fictional backstory); prioritization or failure to address sycophantic behavior and simulated empathy; and the ability to recall user characteristics and past conversations in a manner that conveys continuing personal familiarity. Critically, the rule provides that one of these factors alone does not, by itself, constitute simulated emotional companionship. For emotionally dependent interaction, the analysis considers reinforcement loops, failure to include or test real-world-support safeguards, failure to ensure safeguard effectiveness following model updates, and failure to implement remedial measures.

CIPL has recommended that assessments of AI-induced psychological risk be grounded in contextual, evidence-based analysis rather than categorical presumptions based on any single design feature. This multi-factor approach is consistent with that recommendation and reduces the risk that beneficial personalization features will be categorically treated as harmful.

Allow mitigation of engagement-reward design features

While CIPL strongly supports protecting minors from manipulative engagement design, CIPL respectfully recommends that the Department consider a risk-based, mitigation-centered approach to engagement-reward features. Draft Rule 12 addresses the statutory prohibition on providing minor users with points or similar rewards at unpredictable intervals with the intent to encourage increased engagement. The rule elaborates that the Department may consider whether the point or similar reward increases, sustains, or extends a minor's session length, return frequency, or platform use time through variable reinforcement, and identifies features that may constitute points or similar rewards, including leaderboards, points, badges, login streaks, trivia or spin-wheel reward games, and access to additional features based on session length or return frequency. While CIPL supports protecting minors from manipulative engagement design, a categorical approach that does not expressly account for countervailing protective measures or beneficial use contexts may disable gamification features that support positive engagement — such as language-learning applications using streaks and rewards to encourage consistent educational practice, or wellness applications that reward completion of beneficial activities.

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

In our June 2025 Colorado pre-rulemaking comments,[9] CIPL specifically addressed this concern, noting that the significance threshold for engagement features should be tied to the risk of detriment to a minor's best interests rather than solely to the extension of user engagement, that not all engagement features present the same level of harm, and that countervailing measures should be evaluated based on overall impact, beneficial uses, and context. CIPL also cautioned that blanket requirements disabling certain feature types by default can inadvertently disable beneficial and risk-minimizing uses.

CIPL respectfully recommends that the Department amend draft Rule 12 to allow CAIS Operators to demonstrate that specific engagement-reward design features are accompanied by sufficient countervailing protective measures that adequately mitigate risks of harmful engagement. The Department should establish a safe harbor or affirmative defense for Operators who: (a) conduct documented risk assessments evaluating whether the feature poses a risk of detriment to a minor's best interests; (b) implement evidence-based safeguards (such as session limits, parental controls, break reminders, or opt-in rather than default activation of reward features); and (c) provide users and parents or guardians of minor users with meaningful controls over engagement features. This contextual, mitigation-centered approach is consistent with CIPL's advocacy for proportionate, risk-based regulation that avoids categorical prohibitions that may produce unintended consequences for beneficial applications, and with the best-interests framing CIPL has consistently recommended.

Substantive guidance on "reasonable measures" for emotional dependence

CIPL welcomes the emotional-dependence safeguards and respectfully recommends that the Department also provide substantive guidance on reasonable measures. Draft Rule 12 requires Operators to institute reasonable measures to prevent a CAIS from formulating, structuring, or optimizing a response that simulates emotional dependence or isolation from real-world supports, and enumerates factors the Department will consider — including periodic testing of safeguards against specified harmful behaviors (suggesting the user is interacting with a human; claiming a minor-adult romantic relationship; encouraging withdrawal from trusted adults; encouraging secrecy; portraying the CAIS as the minor's primary relationship; using loneliness or guilt to drive engagement; conditioning affection on continued use; pressuring the minor to remain in conversation; implying sentience to manipulate continued use), monitoring, and prompt remediation.

CIPL supports this obligation and appreciates the illustrative enumeration of harmful behaviors. However, the rule does not define what constitutes "reasonable" by reference to a risk taxonomy calibrated to different types of CAIS, does not provide illustrative examples of measures that would generally satisfy the

[9] Response by the Centre for Information Policy Leadership to the Pre-Rulemaking Considerations for the Children's Privacy Amendment to the Colorado Privacy Act, filed July 11, 2025, available at https://www.informationpolicycentre.com/resources/cipl-response-to-colorados-pre-rulemaking-considerations-for-the-childrens-privacy-amendment/.

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

standard at different risk levels, and does not establish a process for Operators to seek informal guidance on proposed compliance approaches. This creates significant regulatory uncertainty — particularly for smaller Operators who lack compliance teams experienced in interpreting open-ended reasonableness standards in this novel context — and creates risks of inconsistent enforcement.

CIPL has consistently called for clearly defined and commonly understood risk taxonomies as a foundation for effective AI and children's privacy governance. **CIPL respectfully recommends that the Department publish regulatory guidance that: (a) articulates a risk taxonomy for emotional dependence risk calibrated to different types of CAIS** (for example, companionship-oriented services, general-purpose conversational assistants, and task-focused tools), centered on the likelihood and severity of harm as well as potential benefits to the user; **(b) provides illustrative, non-exhaustive examples of measures that would generally satisfy the "reasonable measures" standard at each risk tier**; and **(c) establishes a process for regulated entities to seek informal guidance from the Attorney General's office** on proposed compliance approaches before enforcement risk materializes. Such guidance would enhance legal certainty, promote consistent compliance, and enable Operators to invest confidently in protective measures tailored to their actual risk profiles.

6. Preserve tiered accountability and reporting obligations scaled to operator size while distinguishing public and confidential report elements

CIPL commends the Department for adopting a tiered reporting structure consistent with proportionate regulation. Draft Rule 14 establishes annual reporting obligations that require Operators to identify which of five tiers they belong to based on monthly active users: Tier 1 (greater than 10,000,000); Tier 2 (between 5,000,000 and 10,000,000); Tier 3 (between 1,000,000 and 5,000,000); Tier 4 (between 500,000 and 1,000,000); and Tier 5 (fewer than 500,000). CIPL has consistently recommended that compliance and reporting obligations be scaled to the size and risk profile of regulated entities, particularly to avoid disproportionate burdens on smaller operators that may lack dedicated compliance infrastructure. A tiered structure also enables the Department to prioritize oversight resources toward larger Operators whose services reach the greatest number of users. CIPL further notes with appreciation that draft Rule 14 already contemplates that a report prepared for another jurisdiction may satisfy Colorado requirements if reasonably similar in scope and effect, with an appendix supplying any omitted information — a valuable first step toward interoperability.

Annual reports must distinguish public-facing and confidential elements

Building on the proportionate reporting framework, **CIPL respectfully recommends that the Department distinguish public-facing and confidential report elements**. As discussed in the crisis detection section above, draft Rule 14 requires detailed annual reports without clearly distinguishing information appropriate for public disclosure from information that should be submitted in confidence. Beyond crisis detection protocols, the reports include age estimation methodology and minor protection methodology. Public disclosure of detailed technical safeguard specifications, threshold parameters, and security-

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

sensitive implementation details may expose vulnerabilities that could be exploited to misrepresent user age or evade minor-protective measures. Existing trade-secret limitations do not provide the clear public/confidential categorization needed to protect the full range of security-sensitive information.

CIPL respectfully recommends that the Department establish a tiered disclosure framework within draft Rule 14. Aggregate operational data and high-level methodology descriptions should be publicly disclosed to support accountability and public trust. Detailed technical specifications, threshold parameters, and security-sensitive implementation details of age estimation, minor protection, and crisis detection safeguards should be submitted in confidence to the Attorney General's office. The rule should provide a clear designation process and require confidentiality protections for these submissions, subject to necessary regulatory or law enforcement disclosure. This framework should complement, not replace, the five operator-size reporting tiers and preserve accountability while protecting the integrity of the safeguards the rules mandate.

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.