Colorado Department of Law

Consumer Protection Section

AUTOMATED DECISION-MAKING TECHNOLOGY & CONVERSATIONAL ARTIFICIAL INTELLIGENCE SERVICE RULES

4 CCR 904-6

[Editor's Notes follow the text of the rules at the end of this CCR Document.]

Rule 1 General Applicability

1.1 Basis, Specific Statutory Authority, and Purpose

The rules in this Part 904-6 are developed pursuant to section 6-1-108(1), C.R.S., which grants the Attorney General the authority to promulgate such rules as may be necessary to administer the provisions of the Colorado Consumer Protection Act.

The rules are also developed pursuant to:

- C.R.S. § 6-1-1707 (as enacted by SB24-205) which gives the Attorney General authority to promulgate rules for the purpose of implementing and enforcing part 17;
- C.R.S. § 6-1-1704(4) (as amended by SB26-189) which requires the Attorney General to promulgate rules to clarify and implement the post-adverse outcome disclosure requirements;
- C.R.S. § 6-1-1705(3) (as amended by SB26-189) which requires the Attorney General to promulgate rules to clarify and implement the consumer rights requirements; and
- C.R.S. § 6-1-1706(5) (as amended by SB26-189) which gives authority to the Attorney General to promulgate rules as necessary to implement and clarify part 17.

These rules are intended to be consistent with the requirements of the State Administrative Procedures Act, section 24-4-101 *et seq*. (the "APA"), C.R.S., the Automated Decision-Making Technology in Consequential Decisions Act (the "ADMT Act") and the Conversational Artificial Intelligence Services Act (the "Chatbot Safety Act"), section 6-1-1701 *et seq*., C.R.S.

1.2 Severability

If any provision of these rules, 4 CCR 904-6, is found to be invalid by a court of competent jurisdiction, the remaining provisions of these rules shall remain in full force and effect.

1.3 Effective Date

These rules shall become effective January 1, 2027.

1.4 Applicability

These rules, 4 CCR 904-6, are subject to the applicability requirements and exemptions explicitly provided throughout the ADMT Act and the Chatbot Safety Act. Statutes, and regulations referred to within the Colorado ADMT Act and the Chatbot Safety Act do not apply except to the extent explicitly stated. Compliance with the ADMT Act, Chatbot Safety Act, and these rules does not ensure compliance with other state, Federal, or local statutes.

Rule 2 Definitions

2.1 Authority and Purpose

A. The statutory authority for the rules in this Rule 2 is sections 6-1-108(1), 6-1-1707, 6-1-1704(4), and 6-1-1705(3), C.R.S. The statutory authority for the rules in this Rule 2 is also section 6-1-1706(5). The purpose of this Rule 2 is to define certain terms used throughout the ADMT Act and the Chatbot Safety Act provisions of the Colorado Consumer Protection Act, sections 6-1-1701 through 6-1-1709, C.R.S., and these rules. Terms defined by this rule and by sections 6-1-1701 through 6-1-1709, C.R.S. are capitalized where they appear in the rules to signal to the reader to refer to the definitions. When a term is used in a conventional sense and is not intended to be a defined term, it is not capitalized.

2.2 Defined Terms

"Accessible to the General Public" as used in section 6-1-1701(3.5)(a), C.R.S. means, the Conversational Artificial Intelligence System has been made available to Consumers directly through a product or service, including but not limited to free services or products with or without a gated sign-up, services or products provided through a free or purchased subscription, or services or products behind a paywall. Accessible to the General Public does not include an artificial intelligence system made available as an internal-only workforce deployment through restricted, authenticated environments that are not Consumer-facing.

"ADMT" means the same as defined in section 6-1-1701(2), C.R.S.

"ADMT Consumer Rights" means the Consumer rights granted in section 6-1-1705, C.R.S.

"Age Category" means one of the following age groupings: (a) under thirteen years of age; (b) at least thirteen but under sixteen years of age; (c) at least sixteen but under eighteen years of age; (d) at least eighteen but under twenty-five years of age; (e) over twenty-five; and (f) not relevant, which grouping includes accounts for which age is not relevant, such as organizational or entity accounts.

"Alert" means an instance in which the Operator becomes aware, whether through a Complaint or through the Operator's own detection, of content or an interaction that potentially violates the Operator's published policies or section 6-1-1708, C.R.S. (as enacted by HB26-1263).

"Authenticate" means to use reasonable means to determine that a request to exercise any of the ADMT Consumer Rights in section 6-1-1705, C.R.S. is being made by the Consumer who is entitled to exercise the ADMT Consumer Rights, or a guardian legally authorized to act on the Consumer's behalf.

"Complaint" means a report or notification submitted to the Operator alleging that content generated by, or an interaction with, a Conversational Artificial Intelligence Service violates the Operator's published policies or section 6-1-1708, C.R.S. (as enacted by HB26-1263). A Complaint is one form of Alert.

"Consequential Decision" means the same as defined in section 6-1-1701(3), C.R.S.

"Continuous Conversational Artificial Intelligence Service Interaction" as used in sections 6-1-1708(2)(a)(III) and 6-1-1708(3)(b), C.R.S. (as enacted by HB26-1263) means a sequence of exchanges between a user and a Conversational Artificial Intelligence Service that regularly

persists regardless of temporary pauses in activity. A Continuous Conversational Artificial Intelligence Service Interaction may occur within a single conversational thread or across multiple conversational threads or chatbot profiles.

"Deploy" means make use of Covered ADMT in a way that Materially Influences a Consequential Decision.

"Evidence-Based Method" as used in section 6-1-1708(6)(d), C.R.S. (as enacted by HB26-1263), means, for measuring suicidal ideation or self-harm, a method that has been validated through empirical research, scientifically tested for reliability, validity, sensitivity, and specificity. Such methods may include but are not limited to 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 like the SAFE-T protocol.

"Financial or Lending Service" as used in section 6-1-1701(6)(d), C.R.S. means the extension of credit or a loan, transmitting or exchanging funds, the provision of deposit or checking accounts, check cashing, or installment payment plans.

"Makes Publicly Available" as used in section 6-1-1701(15.5)(a), C.R.S. means an Operator develops and offers a Conversational Artificial Intelligence Service to a member of the general public, including to a third party through sale, licensing, or other means.

"Midstream Developer" means a party that integrates Covered ADMT as a component into its own Covered ADMT product and provides its Covered ADMT product to another Developer or a Deployer.

"Process" is defined as set forth in section 6-1-1303(18), C.R.S. and means the collection, use, sale, storage, disclosure, analysis, deletion, or modification of Personal Data.

"Repeated or severe crisis indicators" as referred to in section 6-1-1708(4), C.R.S. (as enacted by HB26-1263) means indications of suicidal ideation or self-harm risk that either recur throughout a user's interactions with a Conversational Artificial Intelligence Service or reflect a heightened level of risk as assessed using an Evidence-Based Method for measuring suicidal ideation or self-harm.

Rule 3 Consumer Communications

3.1 Authority and Purpose

A. The statutory authority for the rules in this Rule 3 is sections 6-1-1707, C.R.S., and 6-1-1706(5)(a), C.R.S. The purpose of the rules in this Rule 3 is to ensure that disclosures, notifications, and other communications to Consumers are clear, accessible, and understandable to Consumers so that Consumers can understand and exercise the full scope of their rights under the ADMT Act, section 6-1-1701, et seq., C.R.S. and that Conversational Artificial Intelligence Service users will understand that a Conversational Artificial Intelligence Service is artificial intelligence, section 6-1-1708, C.R.S. (as enacted by HB26-1263).

3.2 Requirements for all Disclosures, Notifications, Responses, and Communications to Consumers and Users

A. All notices, disclosures, responses to requests for information, and other communications by Deployers to Consumers pursuant to sections 6-1-1703, 6-1-1704 and 6-1-1705, C.R.S. and by Operators to Users pursuant to section 6-1-1708, C.R.S. (as enacted by HB26-1263) must be:

1. Understandable and accessible to Consumers or Users, considering the vulnerabilities or unique characteristics of the Consumer or User receiving the notice, disclosure, response, or other communication. For example, any explanation or description shall use plain, straightforward language and avoid technical or legal jargon.

2. Reasonably accessible to Consumers or Users with disabilities, including through the use of digital accessibility tools. For a description of information provided online, the Deployer or Operator shall follow generally recognized industry standards, such as the Web Content Accessibility Guidelines, version 2.2 of December 12, 2024, from the World Wide Web Consortium, incorporated herein by reference as described at Rule 14.1. In other contexts, the Deployer or Operator shall provide information on how a Consumer or User with a disability may access the disclosure or communication in an alternative format.

3. Available in the languages in which the Deployer or Operator in its ordinary course provides web pages, interfaces, contracts, disclaimers, sale announcements, and other information to Consumers or Users. If the Deployer or Operator has previously interacted with the Consumer or User, all notices, disclosures, responses, and communications shall be sent in the language in which the Deployer or Operator ordinarily interacts with the Consumer or User.

4. Available through a readily accessible interface regularly used in conjunction with the Deployer's or Operator's product or service.

5. Provided in a readable format on all devices through which Consumers or Users interact with the Deployer or Operator, including on smaller screens and through mobile applications, if applicable.

6. Unless otherwise stated, communicated in a manner by which the Deployer or Operator regularly interacts with Consumers or Users; and,

7. Straightforward and accurate and must not be written or presented in a way that is unfair, deceptive, false, or misleading.

Rule 4 ADMT Multiparty Arrangements

4.1 Authority and Purpose

A. The statutory authority for the rules in this Rule 4 is section 6-1-1706(5)(a), C.R.S., and section 6-1-1707, C.R.S. The purpose of the rules in this Rule 4 is to clarify the roles of Deployers and Developers in multiparty arrangements to ensure that all parties can sufficiently carry out their obligations.

4.2 Midstream Developers

A. Any Midstream Developer that integrates Covered ADMT into its own Covered ADMT product and provides its Covered ADMT product to another Developer or a Deployer that further configures the Covered ADMT before Deploying it must:

1. Reasonably obtain all Developer documentation created pursuant to section 6-1-1704, C.R.S. by the Developer of any Covered ADMT used as a component of the Midstream Developer's technology.

2. Make all upstream Developer documentation available to any downstream Deployer or Developer.

B. Example: A company creates an AI application that integrates third-party foundational LLMs. The company then sells those AI applications to customers that input their own data into and deploy those AI applications. The company must obtain all Developer documentation created pursuant to section 6-1-1704, C.R.S., to the extent possible, for any ADMT used as a component of the company's AI system. The company must make that documentation available to customers purchasing its AI system if the company's AI system is being marketed, advertised, configured, contracted, sold, or licensed to be used to Materially Influence a Consequential Decision.

Rule 5 ADMT Developer Obligations

5.1 Authority and Purpose

A. The statutory authority for the rules in this Rule 5 is section 6-1-1706(5)(a), C.R.S., and section 6-1-1707, C.R.S. The purpose of the rules in this Rule 5 is to clarify the Developer obligations set forth in section 6-1-1702, C.R.S.

5.2 Providing Information

A. Any disclosures made pursuant to section 6-1-1702, C.R.S. must be written in a way that provides Deployers with a meaningful understanding and accurate expectations of the ADMT's known limitations, known risks, and circumstances under which the ADMT should not be used.

B. In providing the general statement describing the intended uses and known harmful or inappropriate uses of the covered ADMT, pursuant to section 6-1-1702(1)(a), C.R.S., the Developer must:

1. Describe the types of Consequential Decisions for which the ADMT is intended.

2. Describe the types of Consequential Decisions for which the ADMT is not intended or is known to be an inappropriate or harmful use of the ADMT.

C. In providing instructions for the Deployer's appropriate use, monitoring, and Meaningful Human Review where applicable, pursuant to section 6-1-1702(1)(d), C.R.S., a Developer should:

1. Include instructions specifying the types or categories of data that the ADMT uses and requires when performing its intended functions.

2. Include instructions specifying methods to monitor the ADMT, including, as appropriate, the collection and analysis of aggregate data regarding the ADMT, and the analysis and review of an output of the ADMT that Materially Influenced a specific Consequential Decision.

3. Include instructions, as appropriate, for determining the reasoning and primary factors relied upon by the ADMT to produce an output that Materially Influenced a specific Consequential Decision, and instructions the Deployer needs to appropriately instruct, control, modify, and adjust the ADMT, including any instructions that may assist a Deployer mitigate any known limitations or risks of the ADMT.

4. In providing a description of the categories of data used to train the covered ADMT, the Developer must use a level of detail that provides Deployers a meaningful understanding of the data used, including whether the type of data would be considered Sensitive Data, Biometric Identifiers, or Biometric Data under section 6-1-1301, et seq., C.R.S.

5.3 Withholding Information

A. A Developer may only withhold information as contemplated in section 6-1-1702(1)(e), C.R.S. if that information:

1. Is a Trade Secret as defined by section 7-74-102(4), C.R.S.

2. Is protected from disclosure by state or federal law, or other legal authority.

B. When a Developer withholds information required by section 6-1-1702, C.R.S., the Developer must provide the legal authority for withholding the information.

C. If the Developer determines that a disclosure would reveal the Developer's trade secrets, the Developer must provide information sufficient to satisfy the requirements of section 6-1-1702, C.R.S. in a way that does not reveal trade secrets.

Rule 6 ADMT Adverse Outcome Disclosures

6.1 Authority and Purpose

A. The statutory authority for the rules in this Rule 6 is section 6-1-1704(4)(a)—(b), C.R.S. and section 6-1-1706(5)(a), C.R.S. The purpose of the rules in this Rule 6 is to clarify Developer obligations to provide Adverse Outcome Disclosures, the form and content of those disclosures, and relationship with other legal requirements, as set forth in section 6-1-1704(3)(a)—(c), C.R.S.

6.2 Form of ADMT Adverse Outcome Disclosure

A. The form of the adverse outcome disclosure required by section 6-1-1704(3)(a)—(c), C.R.S. shall meet the following criteria:

1. The disclosure shall be provided to the Consumer in writing (including electronically), via at least two methods provided that the Deployer has two methods of communicating with a Consumer. The Deployer shall provide the disclosure through any channel through which it typically interacts with the Consumer. The Deployer shall provide the disclosure via electronic means if it maintains an online presence or communicates with the Consumer electronically. The Deployer may provide the written disclosure via a single method if it cannot feasibly provide the disclosure via a second method.

a. A Deployer satisfies this standard if both the Consumer's e-mail address and mailing address are known to the Deployer, and it provides the disclosure via electronic mail and regular mail.

b. A Deployer satisfies this standard if it only has one method of contacting the Consumer (e-mail or street address), and it provides the disclosure to the Consumer via that method.

c. A Deployer satisfies this standard if it maintains an online presence and provides the disclosure to the Consumer via e-mail and through the Deployer's direct messaging service.

2. The written disclosure shall be provided in an easy-to-read font and font size. For any printed disclosure that is delivered via postal mail or an equivalent means, the disclosure shall not be in less than 12-point font.

3. The written disclosure shall otherwise comply with the requirements for disclosures, notifications, and other communications to Consumers as set forth in 4 CCR 904-6, Rule 3.2.

6.3 Timing of ADMT Adverse Outcome Disclosure

A. Notification occurs at the time the Deployer delivers, either by e-mailing or mailing the disclosure to the Consumer's last known e-mail address or mailing address, or delivering it via other electronic or telephonic means such as via a messaging portal or text message.

6.4 Content of ADMT Adverse Outcome Disclosure and Consumer Rights Procedures

A. If a Deployer uses a Covered ADMT to Materially Influence a Consequential Decision that results in an Adverse Outcome for a Consumer, the Deployer shall provide within thirty days after making the decision, a disclosure that contains: a plain language description of the Consequential Decision and the role the Covered ADMT played in the Consequential Decision, instructions and a simple-to-follow process to request additional information about the Covered ADMT, the Covered ADMT version number, if applicable, the Covered ADMT Developer, and the types, categories, and sources of Personal Data used, to the extent the Deployer receives the necessary information from the Developer in compliance with section 6-1-1702, C.R.S.; and an explanation of the Consumer Rights described in section 6-1-1705, C.R.S. and how to exercise them.

B. A Deployer's disclosure to Consumers provided pursuant to section 6-1-1704(3)(a), C.R.S. and this Rule shall meet the following criteria:

1. The disclosure describes what Consequential Decision was made about the Consumer using the Covered ADMT.

a. For example, the disclosure states "Bank ABC closed your checking account."

b. For example, the disclosure states "Leasing Company ABC was unable to approve your application to rent 123 Main Street."

2. The disclosure describes the specific purpose for which the Deployer used the Covered ADMT and the role the Covered ADMT played in the Consequential Decision along with the role of any human reviewer(s) and other ADMT or similar systems, if applicable.

a. For example, the disclosure to an unsuccessful candidate for employment states that Employer A used an automated system to produce a score, and that score is

equally weighted along with the scores created by company employees that interviewed the unsuccessful candidate.

b. For example, the disclosure to a Consumer who applies for a unit in an apartment building and has their security deposit amount higher than the amount that other Consumers have received states that Apartment Building A uses an automated system to produce a score for setting security deposits, and that score along with the system inputs is reviewed by an employee.

3. The disclosure contains the effective date of the Adverse Outcome, if applicable.

a. For example, a Covered ADMT Materially Influences a Consequential Decision regarding a Consumer's existing subsidy to make energy and utilities payments, such as by reducing the amount of the subsidy or by terminating it. The disclosure must contain the effective date of the reduction or termination.

b. For example, a Covered ADMT Materially Influences a Consequential Decision regarding a student's existing scholarship, such as by reducing the amount of the scholarship or by terminating it. The disclosure must contain the effective date of the reduction or termination.

4. The disclosure describes the principal reason(s) for the Adverse Outcome with specificity and does not use language that is overly broad or vague.

a. For example, a statement that the Adverse Outcome was based on the Deployer's internal standards or policies is insufficient.

b. For example, a statement that the Adverse Outcome was based on the applicant's credit score and length of employment is sufficient to meet the requirement of specificity.

c. For example, a statement that the Adverse Outcome was based on the applicant's health condition information is sufficient to meet the requirement of specificity.

5. If any of the principal reasons for the Adverse Outcome consisted of or were in part based on inferences based on Personal Data, the disclosure shall state that the outcome was based on one or more inferences, explain and identify the inferences, and explain and identify the underlying Personal Data used to generate that inference.

a. For example, a Consumer's job application was denied because of inferences made about the Consumer's socioeconomic status based on the Consumer's employment history and geographic location. The Adverse Outcome disclosure must state that the principal reason for the Adverse Outcome was inferences made about the Consumer's socioeconomic status based on the Consumer's employment history and geographic location.

6. If any of the principal reasons for the Adverse Outcome consisted of or were in part based on a profile or risk or other score about the Consumer, the disclosure shall state that the Adverse Outcome was based on a profile or risk or other score, shall disclose the Consumer's score, and shall explain and identify the underlying Personal Data used to generate that inference.

a. For example, a Consumer was charged two months' rent instead of one month's rent as a security deposit amount because the Deployer profiled the Consumer as having a high likelihood of property damage based on the Consumer's prior rental history and credit score. The Adverse Outcome disclosure must state that the principal reason for the Adverse Outcome was that the Consumer was profiled as having a high likelihood of property damage based on the Consumer's prior rental history and credit score.

b. For example, a Consumer was charged double two months' rent instead of one month's rent as a security deposit amount because the Deployer formulated a risk score to predict future likelihood of property damage of 82 out of 100 based on the Consumer's prior rental history and credit score. The Adverse Outcome disclosure must state that the principal reason for the Adverse Outcome was that the Consumer received a risk score predicting future likelihood of property damage of 82 out of 100 based on the Consumer's prior rental history and credit score.

7. The disclosure must accurately describe the Consequential Decision, the role the Covered ADMT played in the Consequential Decision, and the principal reason(s) for the Adverse Outcome.

8. If any Personal Data used in the Covered ADMT resulted in an automatic denial under the Deployer's standards or practices, the automatic-denial factor must be disclosed to the Consumer.

a. For example, a disclosure complies with this requirement if it states that an applicant's prior felony conviction from three years prior was an automatic denial factor.

b. For example, a disclosure complies with this requirement if it states that an applicant's credit score of 580 was an automatic denial factor.

9. If the Adverse Outcome was the result of insufficient or incomplete Personal Data or other information, the fact that the Adverse Outcome resulted from insufficient or incomplete information must be disclosed to the Consumer, along with the additional information that the Deployer and/or the covered ADMT would need to make a decision. The Deployer need not accept any additional information provided by the Consumer unless and until the Consumer requests Meaningful Human Review.

a. For example, a disclosure complies with this requirement if it states that the Consumer was removed from an employment applicant pool pre-interview due to incomplete information, and that the Consumer would need to provide employment history for the past five years and all professional licenses that the Consumer had obtained and were currently active.

10. The disclosure is provided in writing or electronically to the Consumer in compliance with the requirements in 4 CCR 904-6, Rule 6.2 and otherwise complies with the requirements in 4 CCR 904-6, Rule 3.2.

11. A Deployer does not comply with section 6-1-1704(3)(a), C.R.S. when it is not able to explain how a Covered ADMT Materially Influenced a Consequential Decision, or how the Covered ADMT used a Consumer's Personal Data in connection with that decision.

12. A Deployer does not comply with section 6-1-1704(3)(a), C.R.S. when it is not able to accurately explain the principal reason(s) why a Covered ADMT Materially Influenced a Consequential Decision resulting in an Adverse Outcome.

13. Pursuant to 6-1-1704(5), C.R.S., nothing in this rule requires a Deployer to disclose a trade secret or information protected from disclosure by state or federal law. If a Deployer withholds information on this basis, the Deployer shall notify the Consumer.

C. A Deployer's instructions and simple-to follow process provided pursuant to section 6-1-1704(3)(b), C.R.S. and this rule shall meet the following criteria:

1. The instructions to request additional information about the Covered ADMT must include one or more specific methods through which the Consumer may request the additional information. One method must be online through a link that is provided to the Consumer as a hyperlink in an e-mail, online, or text-based disclosure, or a URL that a Consumer can navigate to in a web browser, for a written disclosure sent via mail or equivalent means. The website through which a Consumer may request the additional information contemplated in C.R.S. § 6-1-1704(3)(b) may be the same website used to provide ADMT Consumer Rights contemplated in section 6-1-1705, C.R.S., as described in 4 CCR 904-6, Rule 7.2.

2. For additional methods that the Deployer provides to access additional information other than a hyperlink, the process by which Consumers may request additional information about the Covered ADMT must be easy for Consumers to follow and execute and require a minimal number of steps.

3. The process by which Consumers may request additional information about the covered ADMT must enable Consumers to submit their request at any time of day, not just during certain or restricted hours.

4. The disclosure must state that, by visiting the hyperlink or by otherwise following the instructions, the Consumer can obtain additional information, including the name of the Covered ADMT, the Covered ADMT version number, the Covered ADMT Developer, and the types, categories, and sources of Personal Data used by the Covered ADMT.

D. A Deployer's explanation of the ADMT Consumer Rights described in section 6-1-1705, C.R.S. and how to exercise them pursuant to section 6-1-1704(3)(c), C.R.S. shall meet the following criteria:

1. The disclosure explains that the Consumer is entitled to:

a. Request their Personal Data and to correct factually incorrect or inaccurate Personal Data.

b. An opportunity for Meaningful Human Review and reconsideration of the decision, to the extent Commercially Reasonable.

2. The disclosure must provide a unified method by which the Consumer can exercise their rights and request additional information about the Covered ADMT in compliance with 4 CCR 904-6, Rule 7.

6.5 Content of ADMT Adverse Outcome Disclosures: Examples by Sector

A. This rule provides examples of how Adverse Outcomes within certain Covered Domains may be described in Deployers' Adverse Outcome disclosures to meet the requirements of section 6-1-1704(3)(a), C.R.S. These examples do not cover the information required by section 6-1-1704(3)(b)—(c), C.R.S. These examples are illustrative only, are not exhaustive, and are not intended to serve as the basis for Deployer decisions regarding Adverse Outcome disclosures. These examples do not suggest that the information considered in making a Consequential Decision as set forth in these examples is appropriate and complies with applicable law. Deployers must determine whether they need to send an Adverse Outcome disclosure to a Consumer, and what it must contain, consistent with section 6-1-1701 et seq., C.R.S. and these rules, and whether the information the Deployer uses to make a Consequential Decision complies with applicable law.

B. Education

1. A secondary educational institution subject to the Federal Education Right to Privacy Act utilizes a Covered ADMT to make decisions about scholarship eligibility and the amount of such scholarships. An incoming student applies for a scholarship that would provide up to full tuition and room and board as a grant. The covered ADMT determines that the incoming student shall receive a scholarship for 50% of the school's tuition, and no award for room and board, based on the incoming student's income, bank statements and High School GPA. The school must provide an Adverse Outcome disclosure consistent with 4 CCR 904-6, Rules 6.3 and 6.4, but may utilize the channels through which it provides Federal Education Right to Privacy Act notices for providing the notice. As an illustrative example, the disclosure may state in part, "An automated system was used in making a decision regarding your scholarship eligibility, and a human reviewer verified the outcome of the decision. You are only eligible for a 50% tuition scholarship, covering $16,850 in your first year. The principal reasons for the decision are your income, financial information contained in your bank statements, and your High School GPA."

2. A private high school not subject to the Federal Education Right to Privacy Act utilizes a Covered ADMT to make decisions about admission to the school. The school uses a covered ADMT which is supplemented by a score from the employees who interview applicants. The ADMT determines that the applicant should be rejected due to their academic records from their public middle school and inferences about the applicant's future academic performance based on information obtained from social media. The school must provide an Adverse Outcome disclosure consistent with 4 CCR 904-6, Rules 6.2, 6.3 and 6.4. As an illustrative example, the disclosure may state in part, "An automated system was used in making a decision regarding your admission to School A, and the output was supplemented by information by School A employees regarding your applicant interview. You have been denied admission to School A. The principal reasons for the decision are your academic record from your middle school, and inferences about your future academic performance based on information we obtained from your social media accounts."

C. Lease or Purchase of Residential Real Estate

1. A rental housing provider utilizes a Covered ADMT as the sole basis for its decision about the amount of security deposit to charge a tenant. The rental housing provider decides to charge a consumer $2,500 for a security deposit for a three-bedroom apartment, where it has charged other consumers anywhere from $1,500-$2,500 as a security deposit for the same sized unit, after utilizing a Covered ADMT that analyzes information provided directly from the applicant and determined that the applicant should pay a higher deposit based on the applicant's lack of recent rental history. The rental housing provider does not provide an adverse action notice pursuant to the federal Fair Credit Reporting Act. The rental housing provider must provide an Adverse Outcome Disclosure consistent with 4 CCR 904-6, Rules 6.2, 6.3, and 6.4. As an illustrative example, the disclosure may state in part, "An automated system was the sole basis used in making a decision regarding the amount of your security deposit. Your security deposit for Unit 1 is set at $2,500. The principal reason for the decision is that your application did not provide sufficient rental history."

2. An organization accepts applications for closing cost assistance in the form of grants to low and middle-income home purchasers who meet certain requirements and utilizes a Covered ADMT, reviewed by an organization employee, to determine eligibility. The organization denies a consumer's closing cost assistance application based on their income, which the Covered ADMT determined was above the program cut-off. The organization must provide an Adverse Outcome Disclosure consistent with 4 CCR 904-6, Rules 6.2, 6.3, and 6.4. As an illustrative example, the disclosure may state in part, "An automated system was used to make a decision regarding your eligibility for closing cost assistance, and the decision was reviewed by an Organization A employee. Your application for closing cost assistance was denied. The principal reason for the decision is that your income is above the program limit of $75,000, which is an automatic denial factor under our program standards."

D. Financial or Lending Services

1. A bank refers to collections a consumer's delinquent credit card debt based on a Covered ADMT. The Covered ADMT made this decision based on inferences it made about the likelihood of the consumer repaying their debt by considering the consumer's geographic location, credit score, and payment history on the credit card account. The bank does not provide an adverse action notice in compliance with the requirements of the federal Equal Credit Opportunity Act for this decision, and therefore must provide a compliant Adverse Outcome disclosure consistent with 4 CCR 904-6, Rules 6.2, 6.3, and 6.4. As an illustrative example, the disclosure may state in part, "An automated system was the sole basis used in making a decision regarding whether your credit card debt would be sent to collections. Bank A has sent your delinquent credit card account ending in 1234 to collections. The principal reason for the decision is inferences about your likelihood of repaying your debt based on your geographic location, credit score, and payment history on the credit card account. Your account will be referred to collections on January 1, 2029."

2. A creditor denies an applicant's application for a personal loan using a Covered ADMT. The creditor provides an adverse action notice in compliance with the requirements of the

federal Equal Credit Opportunity Act. The creditor need not provide an additional Adverse Outcome disclosure as long as its disclosure contains the information required by section 6-1-1704(3)(b)—(c), C.R.S. and in 4 CCR 904-6, Rule 6.4(C)—(D).

E. Insurance

1. A third-party claims administrator that is hired by an insurance provider that is subject to Section 6-1-1701, C.R.S. et seq. to run the provider's claims process uses a Covered ADMT to review claims and approve or deny them. The Covered ADMT is the sole basis for the decision. The Covered ADMT denies a home insurance claim based on the claimant's prior claim history, and inferences about the likelihood of submitting future home insurance claims based on the consumer's geographic location. The insurance provider must provide a compliant Adverse Outcome disclosure consistent with 4 CCR 904-6, Rules 6.2, 6.3, and 6.4. As an illustrative example, the disclosure may state in part, "An automated system was the sole basis used in making a decision regarding your home insurance claim. Insurer A has denied your claim. The principal reasons for the decision are your prior claim history and inferences about your likelihood of submitting future claims based on your geographic location."

F. Employment

1. Employer A uses a Covered ADMT to select 10 candidates for interviews among a pool of 87 applicants. The Covered ADMT utilizes information from a "consumer report" and a "consumer reporting agency" as defined by the federal Fair Credit Reporting Act. Employer A provides adverse action notices under the federal Fair Credit Reporting Act to the 77 candidates who were not selected for interviews, and need not provide separate Adverse Outcome disclosures as long as its notice contains the information required by section 6-1-1704(3)(b)—(c), C.R.S. and in 4 CCR 904-6, Rule 6.4(C)—(D).

2. Employer B uses a Covered ADMT to make promotion recommendation decisions, which is reviewed by two human reviewers. The Covered ADMT does not recommend that a software developer receives a management promotion based on the developer's performance evaluation scores for the past 6 years and her score on a technical assessment of coding proficiency. Employer B must provide a compliant Adverse Outcome disclosure consistent with Rules 6.2, 6.3, and 6.4. As an illustrative example, the disclosure may state that "An automated system was used to make a decision regarding your eligibility for a promotion, and that decision was reviewed by two human reviewers. You are not eligible for a promotion to the role of Senior Software Developer. The principal reasons for the decision are your performance evaluation scores for the past 6 years and your technical assessment score."

6.6 Deployer Responses to Consumer Requests Regarding Covered ADMTs

A. If a Consumer requests additional information about the Covered ADMT as provided in section 6-1-1704(3)(b), C.R.S. through a Deployer-provided hyperlink or via telephone, the Deployer must:

1. Make available immediately, upon the Consumer clicking or navigating to the hyperlinked page or calling a toll-free number, the information about the Covered ADMT and the inputs, including the name of the Covered ADMT, the covered ADMT version

number, if applicable, the Covered ADMT Developer, and the types, categories, and sources of personal data used.

2. The description of the types and categories of Personal Data that were considered by the Covered ADMT shall be described in a level of detail that provides Consumers with a meaningful understanding of the information considered by the Covered ADMT. For example, types and categories of Personal Data described at a sufficiently granular level of detail include, but are not limited to, "Credit score," "Medical or health information," and "criminal history."

3. The description of the sources of Personal Data that were considered by the Covered ADMT must identify each source by name.

    a. For example, specific data brokers, databases, courts, social media companies, schools, and employers, and any other source of Personal Data for the covered ADMT, must be identified by name.

    b. If the Deployer has obtained Personal Data from a third party such as a data broker or aggregator that obtains personal data from first party or other third-party sources, the description shall include the original source of the information and any intermediary or vendor source(s) of the information.

4. Information provided pursuant to section 6-1-1704(3)(b), C.R.S. must be provided in compliance with the requirements for disclosures, notifications, and other communications, as described in 4 CCR 904-6, Rule 3.2.

    a. For example, a Deployer complies with this requirement by providing all of the required information at the hyperlinked webpage and making it immediately available upon the Consumer clicking or navigating to the page.

B. If a Consumer requests additional information about the Covered ADMT as provided in section 6-1-1704(3)(b), C.R.S. through mail, e-mail, or other equivalent means, provided that the Deployer has established such method as a means for a Consumer to request the additional information as set forth in section 6-1-1704(3)(b), C.R.S., the Deployer must:

1. Provide the information about the Covered ADMT and the inputs, including the name of the Covered ADMT, the Covered ADMT version number, if applicable, the Covered ADMT Developer, and the types, categories, and sources of Personal Data used within ten business days after receipt of the Consumer's request, and provide the information via the same means that the Consumer used to transmit their request.

    a. For example, a Deployer's Adverse Outcome disclosure provides a hyperlink, an e-mail address, and a phone number for a Consumer to obtain the information required by section 6-1-1704(3)(b), C.R.S. If a Consumer requests the additional information via e-mail, a Deployer complies with this requirement by providing all of the required information to the Consumer at the same e-mail address they used to transmit the request within ten business days after receiving the e-mail.

6.7 Compliance with Other Legal Obligations and Effect on ADMT Adverse Outcome Disclosure Requirements

A. If a Deployer that uses a Covered ADMT to Materially Influence a Consequential Decision that results in an Adverse Outcome for a Consumer within the meaning of section 6-1-1701, *et seq*., C.R.S. must provide an adverse action or risk-based pricing notice pursuant to the Equal Credit Opportunity Act (ECOA), 15 U.S.C. § 1691, *et seq*., and/or the Fair Credit Reporting Act (FCRA), 15 U.S.C. § 1681, *et seq*., the Deployer complies with section 6-1-1704(3), C.R.S. as long as its notice or disclosure meets the following requirements:

1. The disclosure contains the information stated in section 6-1-1704(3)(b)—(c), C.R.S. and in 4 CCR 904-6 Rule 6.4(C)—(D).

2. It provides the disclosure within thirty days after making the decision.

B. For a Consequential Decision relating to education, a Deployer that is subject to the Federal Education Right to Privacy Act (FERPA) satisfies the notice and disclosure requirements of section 6-1-1704(2),(3), C.R.S. by providing notice and disclosures to a parent or guardian or an eligible student through processes and channels that are consistent with FERPA and the Deployer's FERPA notices and student record access procedures.

C. If a Deployer that uses a Covered ADMT to Materially Influence a Consequential Decision that results in an Adverse Outcome for a Consumer within the meaning of section 6-1-1701, et seq., C.R.S. and the Deployer must provide an adverse action disclosure or equivalent notice pursuant to any other federal, state, or local law other than those listed in this rule, the Deployer complies with section 6-1-1704(3), C.R.S. as long as its notice or disclosure meets the following requirements:

1. The Adverse Outcome disclosure or equivalent disclosure or notice also satisfies the notice or disclosure requirements of section 6-1-1703, C.R.S. and these rules.

2. The Deployer provides the disclosure in writing within thirty days after making the decision.

6.8 Miscellaneous Provisions

A. Nothing in this section prohibits a Deployer from providing additional information to enable a Consumer to understand how the Covered ADMT was used to make a Consequential Decision with respect to them.

1. For example, a Deployer may provide the range of possible outputs or aggregate output statistics to help a Consumer understand how they compare to other Consumers, such as the five most common outputs of the Covered ADMT and the percentage of Consumers that received each of those outputs during the preceding calendar year.

B. A Deployer must not retaliate against a Consumer because the Consumer exercised any of their rights pursuant to sections 6-1-1704 and 6-1-1705, C.R.S., and these rules.

Rule 7 ADMT Consumer Rights

7.1 Authority and Purpose

A. The statutory authority for the rules in this Rule 7 is section 6-1-1705(3), C.R.S. The purpose of the rules in this Rule 7 is to clarify and implement the requirements of section 6-1-1705, C.R.S.

7.2 Submitting Requests to Exercise ADMT Consumer Rights

A. Pursuant to section 6-1-1704(3)(c), C.R.S., and 4 CCR 904-6, Rule 6.4(D), a Deployer's Adverse Outcome disclosure must include an explanation of the ADMT Consumer Rights described in section 6-1-1705, C.R.S. and how to exercise them.

B. In accordance with section 6-1-1705(1)(a), C.R.S., a Consumer may request, and the Deployer shall provide: (1) instructions for requesting Personal Data and correcting factually incorrect or materially inaccurate Personal Data used in a Consequential Decision that used a Covered ADMT; and (2) an opportunity for Meaningful Human Review and reconsideration of the Consequential Decision to the extent commercially reasonable.

1. A Deployer shall allow a Consumer to both request instructions, and to seek Meaningful Review by providing a clearly-labeled link in the Adverse Outcome disclosure leading directly to the instructions, as well as a mailing address or toll-free number that the Consumer can contact to receive those instructions.

2. If the Consumer requests the instructions using the toll-free number or address, instructions must be provided within 24 hours of receipt.

3. The link provided in the Adverse Outcome disclosure must bring the Consumer directly to the Consumer request method, which, in accordance with 4 CCR 904-6, shall provide Consumers clear instructions on how to exercise their ADMT Consumer Rights.

4. An example of a "clearly-labeled link" for purposes of this section would state: "Click here for instructions on how to exercise your Consumer Rights under the Automated Decision-Making Technology Act."

C. Any method specified by a Deployer in its instructions pursuant to this rule must:

1. Consider the ways in which Consumers normally interact with the Deployer:

a. A Deployer shall provide two or more designated methods for submitting an ADMT Consumer rights request, which must consider the ways in which Consumers normally interact with the Deployer.

b. If a Deployer maintains a website, mobile application, or other digital presence, one method for submitting requests shall be through its website, mobile application, or digital interface, such as through a webform. For clarity, a clear and conspicuous personalized hyperlink to a webform provided to Consumers as a part of an electronic Adverse Outcome disclosure satisfies this requirement, even if the webform is not publicly accessible on the Deployer's website or mobile application.

c. Any methods for submitting an ADMT Consumer Rights request must be regularly monitored by the Deployer, by an individual with the knowledge and ability to process the Consumer's ADMT Consumer Rights request.

2. Enable the Consumer to submit the request to the Deployer at any time;

3. Comply with the requirements for disclosures, notifications, and other communications to Consumers provided in 4 CCR 904-6, Rule 3.2;

4. Use reasonable data security measures, consistent with 4 CCR 904-3, Rule 6.9, when exchanging information in furtherance of ADMT Consumer Rights requests, considering the volume, scope and nature of Personal Data that might be exchanged; and

5. Be easy for Consumers to execute, requiring a minimal number of steps.

a. Deployers that require Consumers to follow an unnecessary number of links, steps, or actions may be in violation of this regulation.

D. The ADMT Consumer Rights request method must:

1. Contain a heading indicating that the webform pertains to "Consumer Rights under the Automated Decision-Making Technology Act".

2. Clearly identify and describe all ADMT Consumer Rights available to a Consumer that has received an Adverse Outcome resulting from a Consequential Decision that was Materially Influenced by ADMT;

3. Provide Consumers with clear instructions of how to exercise their ADMT Consumer Rights; and

4. Meet all other requirements of this part 4 CCR 904-6, including requirements specific to individual ADMT Consumer Rights as described in 4 CCR 904-6, Rules 7.3 and 7.4.

E. When a Consumer submits an ADMT Consumer Rights request, a Deployer may only collect Personal Data through the request process if the Personal Data is reasonably necessary to Authenticate the Consumer, respond to the request, or effectuate the ADMT Consumer Rights request.

F. A Deployer must not require a Consumer to create a new user account to exercise their ADMT Consumer Rights request but may require a Consumer to use an existing password-protected account.

7.3 Right to Request Personal Data used in a Consequential Decision that used a covered ADMT

A. Pursuant to section 6-1-1705(1)(a)(I), C.R.S., a Consumer has the right to request Personal Data used in a Consequential Decision that used a Covered ADMT.

B. A Deployer shall comply with a request for Personal Data by providing all the specific pieces of Personal Data used in a Consequential Decision that used a covered ADMT.

1. Specific pieces of Personal Data include, without limitation, final rank, score, classification, recommendation, prediction, or other inferences that pertain to the

individual about whom the decision is being made, along with all ADMT inputs pertaining to that individual.

C. Personal Data provided in response to a request for Personal Data used in a Consequential Decision must:

1. Be provided in a form that is concise, transparent, and easily intelligible and in an appropriate, commonly used electronic format, depending on the nature of the Personal Data;

2. Be available in the language in which the Consumer interacts with the Deployer;

3. Avoid incomprehensible internal codes and, if necessary, include explanations that would allow the average Consumer to make an informed decision of whether to exercise correction or Meaningful Review rights; and

4. Be provided in compliance with the requirements for disclosures, notifications, and other communications, as described in 4 CCR 904-6, Rule 3.2.

D. The Deployer shall implement and maintain reasonable data security measures, consistent with section 6-1-1308(5), C.R.S. and 4 CCR 904-3, Rule 6.9.

7.4 Right to Correct Personal Data used in a Consequential Decision that Used ADMT

A. Pursuant to section 6-1-1705(1)(a)(I), C.R.S., a Consumer has the right to request correction of factually incorrect or materially inaccurate Personal Data used in a Consequential Decision that used a Covered ADMT consistent with section 6-1-1306, C.R.S.

B. A Deployer shall comply with a Consumer's correction request by correcting the Consumer's Personal Data in its existing systems.

C. A Deployer may require the Consumer to provide documentation if necessary to determine whether the Personal Data, or the Consumer's requested correction to the Personal Data, is accurate.

1. When requesting documentation, the Deployer must provide the Consumer with a meaningful understanding of why the documentation is necessary.

2. Any documentation provided by the Consumer in connection with the Consumer's right to correction shall only be Processed by the Deployer in considering the accuracy of the Consumer's Personal Data.

3. The Deployer shall implement and maintain reasonable data security measures, consistent with section 6-1-1308, C.R.S. and 4 CCR 904-3, Rule 6.9, in Processing documentation relating to the Consumer's correction request.

4. Consistent with section 6-1-1306, C.R.S. and 4 CCR 904-3, Rule 4.05, if the Deployer did not receive the Personal Data directly from the Consumer and has no documentation to support the accuracy of the Personal Data, the Consumer's assertion of inaccuracy shall be sufficient to establish that the Personal Data is inaccurate.

D. Where possible, an Adverse Outcome must be stayed pending correction of incorrect Personal Data considered by the ADMT to assist with a Consequential Decision.

7.5 Authentication

A. Pursuant to sections 6-1-1306(1) and 6-1-1705, C.R.S., a Deployer shall use a commercially reasonable method for authenticating the identity of every Consumer submitting an ADMT Consumer Right request.

1. To determine whether an authentication method is commercially reasonable for the purposes of this Rule 7.5, the Deployer shall consider the ADMT Consumer Rights exercised, the type, sensitivity, value, and volume of Personal Data involved, the level of possible harm that improper access or use could cause to the Consumer submitting the ADMT Consumer Right request, and the cost of Authentication to the Deployer. A Deployer must not use methods that place an unreasonable burden on the Consumer submitting the request.

2. The use of verification methods consistent with prevailing industry standards to Authenticate Consumers' identity before providing access to the types of Personal Data requested under an ADMT Consumer Right request likely satisfies the requirements of this section. For example, if most similarly positioned businesses use dual-factor authentication to verify a Consumer's identity before providing certain types of personal data, a Deployer that uses dual-factor authentication to verify the identity of a Consumer before providing that type of personal data likely satisfies the requirements of this section.

B. When possible, a Deployer shall avoid requesting additional Personal Data to Authenticate a Consumer unless the Deployer cannot Authenticate the Consumer using the Personal Data already maintained by the Deployer.

C. Personal Data obtained to Authenticate a Consumer may only be used to Authenticate the Consumer submitting the ADMT Consumer Right request and must be deleted as soon as practical after reviewing the Consumer's request.

D. A Deployer shall implement reasonable security measures, consistent with 4 CCR 904-3, Rule 6.9, to protect Personal Data exchanged to Authenticate a Consumer, considering the type, value, sensitivity, and volume of information exchanged and the level of possible harm improper access or use could cause to the Consumer submitting an ADMT Consumer Right request.

E. A Deployer shall not require the Consumer to pay a fee for Authentication. For example, a Deployer may not require a Consumer to provide a notarized affidavit for Authentication unless the Deployer compensates the Consumer for the cost of notarization.

F. If a Deployer cannot Authenticate the Consumer submitting a Data Right request using commercially reasonable efforts, the Deployer is not required to comply with the Consumer's request. The Deployer shall inform the Consumer that their identity could not be Authenticated, provide information on how to remedy any deficiencies, and may request additional Personal Data if reasonably necessary to Authenticate the Consumer.

G. A Deployer shall not require that a Consumer requesting additional information about the covered ADMT and its inputs pursuant to section 6-1-1704(b), C.R.S. Authenticate their identity.

7.6 Responding to ADMT Consumer Rights Requests for Personal Data and Correction of Personal Data

A. A Deployer shall inform a Consumer of any action taken on a request under this rule without undue delay and, in any event, within 45 days after receipt of the request.

B. If a Deployer decides not to act on a Consumer's ADMT Consumer Right request, the Deployer's response to the Consumer must include the grounds for denial, including but not limited to (1) the Deployer's inability to Authenticate the Consumer's identity; (2) any factual basis for a Deployer's good-faith claim that compliance is impossible; or (3) any basis for a good-faith, documented belief that the request is fraudulent or abusive.

C. If a Deployer denies a Consumer Data Right request based on inability to Authenticate, the Deployer must describe their reasonable efforts to Authenticate and why they were unable to do so.

D. A Deployer that decides not to act on a Consumer's request must also provide instructions on how to appeal the Deployer's decision in accordance with section 6-1-1306(3), C.R.S.

7.7 Right to Meaningful Human Review and Reconsideration

A. Requests for Meaningful Human Review and Reconsideration.

1. Consumers must be able to request Meaningful Human Review and reconsideration through the same mechanism by which they may request to exercise their other ADMT Consumer Rights.

2. The method used to request Meaningful Human Review and reconsideration must enable a Consumer to explain why they would like Meaningful Human Review or reconsideration, describe any specific information contested, and provide additional relevant evidence if necessary.

B. Reconsideration of Corrected Data

1. When a Consumer exercises their right to correct Personal Data, they may exercise their right to reconsideration based on that corrected Personal Data.

2. Reconsideration means that the decision-making process must be re-performed using the Consumer's correct Personal Data.

3. When a Deployer reconsiders a Consequential Decision based on a Consumer's corrected Personal Data, the Deployer must send the results of that reconsideration to the Consumer.

C. Reviewer Standards

1. A Meaningful Human Review must be conducted by an independent reviewer who did not make the original decision and who is not a subordinate of the original decision-maker, whenever feasible.

2. For the reviewer to consider relevant, available primary evidence in compliance with section 6-1-1701(15), C.R.S., the reviewer must have a level of subject matter understanding that is commensurate with the nature of, and negative consequences resulting from, the Adverse Outcome of the Consequential Decision being reviewed.

3. Sufficient training to conduct a Meaningful Human Review in compliance with section 6-1-1701(15)(b), C.R.S. includes training regarding:

    a. Accuracy and objectivity in decision-making, to make the reviewer less likely to exercise their review in a manner that would lead to unfair results;

    b. Information considered by the Covered ADMT and those factors contemplated in section 6-1-1701(D)(I), C.R.S.; and

    c. The subject matter at issue at a level that would enable the reviewer to identify whether review of additional available primary evidence would be valuable, and to review and understand that evidence.

4. To have meaningful authority to approve, modify, or override a Consequential Decision pursuant to section 6-1-1701(15), C.R.S., the reviewer must not be subject to steering by the upper management that would influence the reviewer's decision, and they must be shielded from potential retaliation.

5. ADMT may not assist in the Meaningful Human Review.

D. Types of Meaningful Human Review

1. A "Meaningful Human Review" is one that could change the Adverse Outcome of a Consequential Decision.

2. To ensure that human review is meaningful, a reviewer must consider the Consumer's request for Meaningful Human Review and the Consumer's individual circumstances and determine which type of review could change the Adverse Outcome of the Consequential Decision.

    a. If the Consumer's circumstances and request for Meaningful Human Review indicate that the ADMT may have functioned incorrectly, or did not function as intended, including with respect to accuracy, output, transcription, and glitches, or errors in in configuration, instruction, or other technical failures, then a Meaningful Human Review may include correcting the ADMT and re-performing the decision-making process.

    b. If the Consumer's circumstances and request for Meaningful Human Review indicate that the ADMT itself was problematic, including that it may have considered the wrong factors, did not have access to important relevant information, was not intended to assist in the specific Consequential Decision or type of Consequential Decision as disclosed by the Developer pursuant to 4 CCR 904-6, Rule 6, or because the Consequential Decision related to an atypical circumstance, then a Meaningful Human Review would likely include the review of additional relevant evidence provided by the Consumer in the Consumer's request for Meaningful Human Review, and considering whether that evidence changes the Adverse Outcome.

3. An override of the original decision that results in a full reversal of the original decision indicates that human review was meaningful.

E. Commercial Reasonableness

1. Pursuant to section 6-1-1705(1)(a)(II), C.R.S., a Consumer who experiences an Adverse Outcome resulting from A Consequential Decision in which a Covered ADMT Materially Influenced the Consequential Decision, the Consumer has a right to Meaningful Human Review to the extent commercially reasonable.

2. An assessment of the extent to which Meaningful Review is commercially reasonable shall consider the following factors, weighed together, with no single dispositive factor:

    a. The type of review required;

    b. The magnitude of harm resulting from the Adverse Outcome;

    c. The reversibility of the Adverse Outcome;

    d. The value provided by the review of available primary evidence;

    e. The Deployer's size and capacity;

    f. The marginal cost and technical feasibility of the review; and

    g. The availability of qualified reviewers.

3. When the harm to a Consumer resulting from an Adverse Outcome is a severe and irreversible denial of a basic human need, Meaningful Human Review is presumed to be Commercially Reasonable.

    a. This presumption of Commercial Reasonableness can be rebutted by evidence showing that the review is technically or financially impossible or could not change the Adverse Outcome of the Consequential Decision.

4. A Deployer bears the burden of demonstrating that Meaningful Human Review is not Commercially Reasonable using specific evidence.

5. Example: A property management company uses Covered ADMT to determine whether to renew leases for various Consumers within a group of low-income residents who are not using housing subsidies. The Covered ADMT considers several factors, including history of rental agreement violations and rent payment history from the preceding two years, to determine whether the Consumers are likely to pay rent. The Covered ADMT determines that one of the Consumers is unlikely to pay rent. The property management company relies on the Covered ADMT output and decides not to renew their lease. The Consumer exercises their right to Personal Data used in the Consequential Decision. When the Consumer receives the Personal Data, they realize that the history of rental agreement violations did not reflect the fact that some rental agreement violations were investigated, and the Consumer was found to be not at fault. The Consumer requests a Meaningful Human Review, explaining why they want Meaningful Human Review of the Adverse Outcome, and providing evidence of the fact that violations were investigated, and the Consumer was found to not be at fault.

    a. First, the reviewer must determine what type of human review would be meaningful. Here, a review of the Covered ADMT would not be meaningful because the ADMT would still not have the information needed to consider the fact that the Consumer was cleared of some of the rental agreement violations.

Because it appears that the ADMT did not have necessary important information, a Meaningful Human Review would consider that additional primary evidence provided by the Consumer in their request for Meaningful Human Review.

b. Next, the reviewer must determine whether the type of review that would be meaningful here would be commercially reasonable. Because the harm to the Consumer is a denial of housing – a denial of a basic human need – the review is presumed to be commercially reasonable, unless the property management company can prove that it is technically or financially impossible, or that the review could not change the adverse outcome.

c. The management company's Meaningful Human Review does not prevent liability under any other law or legal authority. The company must ensure that all information used by the ADMT is legal, and that the initial decision-making process was legal, or it may be liable for violations under other applicable law.

6. Example: Under the same circumstances, when the Consumer receives the Personal Data considered by the ADMT, everything is correct. However, the Consumer has no rental agreement violations, and they have consistently paid their rent in full and on time. The Consumer requests Meaningful Human Review because they believe that the Covered ADMT reached the wrong decision.

   a. The reviewer must determine what type of human review would be meaningful. Because the reviewer understands the Covered ADMT and the way it is supposed to function, they feel that the system did not function properly and determined that a Meaningful Human Review would require a review of the Covered ADMT.

   b. Because the harm to the Consumer is denial of housing – a denial of a basic human need – the reviewer will have a particularly high burden of showing that it would not be commercially reasonable to audit the AI system.

   c. In this example, if the audit reveals that the ADMT did not function correctly, the Deployer can correct the ADMT, if possible, and re-perform the decision-making process. If it is not possible for the Deployer to effectively correct the ADMT, a Meaningful Human Review would require the Deployer to review the Consumer's primary evidence De Novo.

   d. The management company's Meaningful Human Review does not prevent liability under any other law or legal authority. The company must ensure that all information used by the ADMT is legal, and that the initial decision-making process was legal, or it may be liable for violations under other applicable law.

7. Example: A Consumer is at XYZ Retail Store checking out, and the sales associate asks if the Consumer wants to apply for an XYZ Retail Store credit card. The Consumer decides to apply during the check-out process. To do so, the Consumer provides Personal Data, and the Store uses Covered ADMT to review the Personal Data to determine whether the Consumer is eligible for the store's credit card. The Consumer is denied. All proper disclosures and notices are made by XYZ Retail Store. The Consumer requests

Meaningful Human Review of the Consequential Decision to deny the Consumer access to credit.

a. The reviewer determines Meaningful Human Review could mean an audit to make sure the Covered ADMT was functioning properly. The reviewer conducts the review and sees that the ADMT did function properly, but that the Consumer's credit score was below the acceptable threshold for that credit card. The reviewer confirms the Adverse Outcome, and no further review is necessary.

b. The reviewer does not have to conduct a substantive review of the underlying evidence here because the ADMT functioned as expected.

c. The credit card company's Meaningful Human Review does not prevent liability under any other law or legal authority. The company must ensure that all information used by the ADMT is legal and that the initial decision-making process was legal, or it may be liable for violations under other applicable law.

F. Response to a request for Meaningful Human Review

1. A response to a Meaningful Human Review or Consideration must include:

a. The reviewer's decision to confirm or override the Consequential Decision;

b. The type of review conducted;

c. The factors considered in making that decision; and

d. The reasons that the reviewer decided to confirm or override the Consequential Decision.

2. Reasons provided must be specific to the evidence provided, and not a recitation of the ADMT's general logic.

G. Timing

1. Where possible, an Adverse Outcome must be stayed pending Meaningful Human Review.

2. Within 10 days after receiving the request, A Deployer must confirm receipt of the request and provide information about how the Deployer will process the request.

3. Meaningful Human Review must be completed, and a response must be provided to the Consumer within 45 days of receiving the request.

H. Documentation

1. When a Meaningful Human Review is conducted, the Deployer must retain a record showing:

a. The reviewer identity, authority, and relevant training;

b. Review timestamps;

c. Primary evidence available to the reviewer, including information provided by the Consumer;

d. The reviewer's access to ADMTs intended use, limitations, inputs and principal factors;

e. Whether the reviewer approved, modified, or overrode the output; and

f. A written justification for the reviewer's decision to approve, modify, or override the output.

Rule 8 Exemptions to the Definition of Conversational Artificial Intelligence Service

8.1 Authority and Purpose

The statutory authority for the rules in this Rule 8 is sections 6-1-108(1) and 6-1-1707, C.R.S. and, these rules are also developed pursuant to section 6-1-1706(5)(a), C.R.S. The purpose of the rules in this Rule 8 is to clarify the exemptions for the definition of Conversational Artificial Intelligence Service pursuant to section 6-1-1701, C.R.S.

8.2 Exemptions to the Definition of Conversational Artificial Intelligence Service

A. When determining if a Conversational Artificial Intelligence Service output is "relating to a narrow and discrete topic", as referred to in section 6-1-1701(3.5)(b)(III), C.R.S., the Department will consider whether the output was generated from completion of a bounded task, such as a status update about a purchase, answering a billing question, or scheduling an appointment. This includes outputs from menu- or button-based chatbots and protocol- and rule-based chatbots as long as these chatbots cannot generate outputs related to sexually explicit conduct or depictions, as described in section 6-1-1708(2)(c) and (2)(d), C.R.S. (as enacted by HB26-1263), or engage in a dialogue or conversation related to suicidal ideation or self-harm.

B. When determining if a Conversational Artificial Intelligence Service is "primarily designed and marketed for commercial use by business entities for the purpose of business operations, productivity, information analysis, internal research, training, or technical assistance" and is exempt under section 6-1-1701(3.5)(b)(IV), C.R.S., the Department will consider a totality of the circumstances, including factors such as whether:

1. The Conversational Artificial Intelligence Service has an interface that is generally available to any consumer;

2. The Conversational Artificial Intelligence Service is marketed primarily to business entities rather than the general public; and

3. The design of the Conversational Artificial Intelligence Service is controlled by and structured around the demands of the specific functions of business operations, productivity, information analysis, internal research, training, or technical assistance.

C. When determining if a Conversational Artificial Intelligence Service is exempt under section 6-1-1701(3.5)(b)(VII), C.R.S. because it is "limited to dialogue focused on the video game and does not maintain dialogue on topics unrelated to the video game", the Department will consider a totality of the circumstances, including factors such as whether:

1. The dialogue outputted by the Conversational Artificial Intelligence Service is limited to or directly relevant to the gameplay, story, or functionality of the video game;

2. The Conversation Artificial Intelligence Service itself is not the primary or only component of the video game;

    a. For video games that consist of distinct components that effectively operate as smaller video games within a larger video game, a distinct subcomponent of a video game may be considered a video game for the purposes of the exemption under section 6-1-1701(3.5)(b)(VII), C.R.S.

D. When determining if a Conversational Artificial Intelligence Service is "designed to simulate emotional companionship", as referred to in section 6-1-1701(3.5)(b)(XI) and (XII), C.R.S., the Department will consider a totality of the circumstances, including factors such as whether:

1. The Conversational Artificial Intelligence Service includes anthropomorphic design features such as personalized dialogue and the ability to assign a chatbot a name, gender, avatar, or fictional backstory;

2. The training, production, and deployment of the Conversational Artificial Intelligence Service prioritized, promoted, or failed to address sycophantic behavior or outputs that simulate empathy and offer users nonjudgmental responses and continual validation.

3. The Conversational Artificial Intelligence Service has the ability to recall and respond to users' characteristics, preferences, and past conversations;

E. When determining if a Conversational Artificial Intelligence Service is "designed to encourage emotionally dependent interaction", as referred to in section 6-1-1701(3.5)(b)(XI) and (XII), C.R.S., the Department will consider a totality of the circumstances, including factors such as whether:

1. The service includes features that are designed to increase emotional intensity across sessions, prioritize emotionally immersive responses, or create reinforcement loops that encourage users to return for reassurance;

2. The service includes and the Operator tests the effectiveness of safeguards including prompts encouraging engagement with real-world supports and reminders that the Conversational Artificial Intelligence Service is not a human;

3. Any safeguards remained effective following model updates, new product features, or changes to the safety systems; and

4. The Operator promptly implemented appropriate remedial measures to address weaknesses identified through testing or real-world safety incidents.

Rule 9 Age Assurance for Conversational Artificial Intelligence Services

9.1 Authority and Purpose

The statutory authority for the rules in this Rule 9 is sections 6-1-108(1) and 6-1-1707, C.R.S. and, these rules are also developed pursuant to section 6-1-1706(5)(a), C.R.S. The purpose of the rules in this Rule 9 is to clarify the requirements for age assurance pursuant to section 6-1-1708, C.R.S. (as enacted by HB26-1263).

9.2 Compliance

Pursuant to section 6-1-1708(2), C.R.S. (as enacted by HB26-1263), Operators shall use commercially reasonable methods or generally accepted methods to estimate the age of account holders or users.

A. A commercially reasonable or generally accepted method to estimate the age of account holders and users must comply with each of the following:

1. Protect a user's privacy by only collecting data necessary for the purpose of age assurance of the specific user for which the data was collected, and not sharing, or using the data for any other purpose except as required by law. Collected data should be deleted after the minimum time required for compliance.

2. Maximize user choice in how the user's age can be estimated. An Operator must not use as a sole method of age assurance a method that requires a user to provide government-issued identification.

3. Be reasonably effective at identifying users under age 18, perform with measurable consistency, and be tested, with quantifiable results to show accuracy rates that are in line with industry standard rates.

4. Utilize operating system, app store, and device signals transmitted under applicable law. However, the Operator may not willfully disregard other reliable evidence and signals that contradict the operating system, app store, or device signal.

5. Include a responsibility of the Operator to re-assess age estimates based on new signals, including but not limited to user-provided information, account settings, parental controls, device settings, previous interactions, contextual information derived through conversations, behavioral indicators, and payment methods.

6. Include a process to receive and respond to any report or information indicating a user is a minor or that data was falsified regarding their age.

7. Not solely rely on user self-declarations of age; general contractual restrictions or terms regarding the use of the internet website, online service, online application, or mobile application; or online payment methods that are available to minors.

8. Not rely on a method that has known or documented risks without taking reasonable measures to mitigate those risks.

9. Include reasonably effective measures to attempt to detect fraudulent or misused information provided for the purpose of age verification.

B. Generally accepted methods to estimate the age of account holders include, but are not limited to:

1. Using an age assurance method that satisfies the framework set forth in ISO/IEC 27566: "Information security, cybersecurity, and privacy protection – Age assurance systems", incorporated herein by reference as described at 4 CCR 904-6 Rule 14.2.

2. Using a cryptographic technique, such as a zero-knowledge proof, that allows a user to demonstrate whether they are a minor using verified data that does not reveal any other information about the user to the Operator or a third-party.

3. Matching a scan of a government-issued ID and live photo or video using facial recognition.

4. Using an assessment of the user's digital footprint originating from a verified email address.

C. When determining if an age assurance method is commercially reasonable, the Department will consider a totality of the circumstances, including the size and financial resources, of an Operator, and the costs and effectiveness of available age estimation techniques for users.

D. An inconclusive age assurance outcome from a generally accepted or commercially reasonable method cannot provide the basis for an Operator to reasonably determine that the user is not a minor.

E. When determining if an Operator willfully disregarded clear and convincing information that an account holder or user is a minor, the Department will consider the factors in 4 CCR 904-3 Rule 6.13 as well as:

1. If the Operator failed to ascertain or verify the age of a user on a platform;

2. If the Operator ignored system flags such as an estimate that a user's age is under 18, self-reported age, user-provided information, account settings, parental controls, device settings, previous interactions, behavioral indicators, and payment methods;

3. If the Operator failed to address external notifications such as parental complaints and user- or law-enforcement reports;

4. If the Operator failed to consider and act on signals from a user's behavior and inputs processed by the LLM, including but not limited to direct statement of minor status, recurring discussion of grade-level and high-school situations, under 18 age-specific activities, or a user consistently occupying a child role in described parent-child interactions;

5. If the Operator failed to consider and act on signals from a user's behavior and inputs processed by the LLM, including but not limited to direct statement of minor status, recurring discussion of grade-level and high-school situations, under 18 age-specific activities, or a user consistently occupying a child role in described parent-child interactions;

6. If the Operator failed to consider evidence and signals that contradict the operating system, app store, or device signal provided pursuant to section 6-30-102, C.R.S.

7. An Operator's use of design choices that nudge users toward inaccurate age estimation or attestation, including but not limited to age reporting tools that automatically fill in a birth date over age 18.

Rule 10 Disclosure

10.1 Authority and Purpose

The statutory authority for the rules in this Rule 10 is sections 6-1-108(1) and 6-1-1707, C.R.S. and, these rules are also developed pursuant to section 6-1-1706(5)(a), C.R.S. The purpose of the rules in this Rule

10 is to clarify the scope and content of the required disclosures pursuant to sections 6-1-1708(2)(a) and 6-1-1708(3), C.R.S. (as enacted by HB26-1263).

10.2 Disclosure to Users

A. Pursuant to section 6-1-1708(2)(a), C.R.S. (as enacted by HB26-1263), an Operator shall clearly and conspicuously disclose to the minor account holder or minor user that the account holder or user is interacting with artificial intelligence that is artificially generated and not human.

B. A persistent visible disclaimer for a product with a screen interface required under section 6-1-1708(2)(a) and 6-1-1708(3), C.R.S. (as enacted by HB26-1263) must:

1. Remain continuously visible throughout the conversation and not require a user to scroll, click, hover, or take other affirmative action to view the disclaimer;

2. Be displayed in a manner that is distinct from the conversation, including by using a font size that is not smaller than the largest font size of other text appearing on the interface; and

3. Be displayed in the same language of the chatbot conversation or user account language settings;

C. An intermittent audio disclaimer for a product without a screen interface required under section 6-1-1708(2)(a), C.R.S. (as enacted by HB26-1263) must:

1. Be played in the same language as the conversation or user account language settings;

2. Be played at a higher volume than the conversation;

3. Be repeated at intervals reasonably calculated to ensure that a user engaged in a conversation remains aware that they are interacting with a Conversational Artificial Intelligence Service and each time a user prompts the Conversational Artificial Intelligence Service about whether the Conversational Artificial Intelligence Service is a human;

a. When determining if the intervals are reasonable, the Department will consider factors including but not limited to:

(1) Whether the Operator performed any testing of users' understanding that they are interacting with a Conversational Artificial Intelligence Service and whether the disclosure supports safer decision-making;

(2) Whether disclosures are delivered at meaningful points during the user experience; and

(3) Whether changes were made following testing and product updates.

D. Pursuant to section 6-1-1708(3), C.R.S. (as enacted by HB26-1263), an Operator shall clearly and conspicuously disclose to a user, that is not a minor user or minor account holder covered by the disclosure requirement pursuant to section 6-1-1708(2)(a), C.R.S. and 4 CCR 904-6 Rules 10.2 (B) and (C), that a Conversational Artificial Intelligence Service is artificial intelligence.

1. The disclosure must be provided at the beginning of a user's first interaction with a Conversational Artificial Intelligence Service for each calendar day of interaction;

2. Appear at least once every three hours in a Continuous Conversational Artificial Intelligence Service Interaction or appear as a persistent disclosure visible to the user as described in 4 CCR 904-6 Rule 10.2(B); and

3. Be provided in response to user prompts regarding whether the Conversational Artificial Intelligence Service is artificially generated and not human.

E. All disclosures required under sections 6-1-1708(2)(a) and 6-1-1708(3), C.R.S. (as enacted by HB26-1263) must comply, as relevant, with Rule 3 of these rules.

Rule 11 Minor Account Holders and Minor Users

11.1 Authority and Purpose

The statutory authority for the rules in this Rule 11 is sections 6-1-108(1) and 6-1-1707, C.R.S., and these rules are also developed pursuant to section 6-1-1706(5)(a), C.R.S. The purpose of the rules in this Rule 11 is to clarify the scope and content of the protections for minor account holders and minor users pursuant to section 6-1-1708(2)(b)-(h), C.R.S. (as enacted by HB26-1263).

11.2 Points or Rewards

A. When determining if an Operator provided the minor account holder or minor user with points or similar rewards at unpredictable intervals with the intent to encourage increased engagement with a Conversational Artificial Intelligence Service, as described in section 6-1-1708(2)(b), C.R.S. (as enacted by HB26-1263), the Department may consider if the point or similar reward increases, sustains, or extends a minor's session length, return frequency, or platform use time with a Conversational Artificial Intelligence Service through variable reinforcement.

B. "Points or similar rewards" includes but is not limited to:

1. Design features such as leaderboards, points, badges, or login streaks;

2. Design features that require users to play a game, answer trivia questions, spin a wheel, or use similar common games to receive rewards;

3. Access to additional features of the service based on session length, return frequency, or platform time.

11.3 Technically Feasible Measures

A. When determining if an Operator instituted technically feasible measures to prevent a Conversational Artificial Intelligence Service from producing, generating, or engaging in prohibited conduct as required in section 6-1-1708(2)(c), C.R.S. (as enacted by HB26-1263), the Department will consider factors including but not limited to:

1. Whether comparable safeguards are already used elsewhere in the industry;

2. Whether commercially available technologies could substantially reduce the identified risk;

3. Whether alternative safeguards could meaningfully reduce the risk if a preferred solution is unavailable, including substantially altering the Conversational Artificial Intelligence Service or choosing not to deploy the Conversational Artificial Intelligence Service.

4. Whether advances in technology warrant periodic reassessment of earlier feasibility determinations;

5. Whether an Operator can demonstrate that the Operator periodically tested safeguards using realistic scenarios involving minors, evaluated whether those safeguards remained effective following model updates or new product features, and promptly addressed weaknesses identified through testing, monitoring, or real-world safety incidents;

6. Whether the Operator effectively monitored the Conversational Artificial Intelligence Service for producing, generating, or engaging in prohibited conduct, including but not limited to monitoring and acting on user reports; and

7. Whether an Operator has documentation that a measure was evaluated and tested, that limitations or failures were identified, that alternative measures were considered, and that additional safeguards were not technically feasible.

11.4 Reasonable Measures

A. In determining whether an Operator instituted reasonable measures to prevent a Conversational Artificial Intelligence Service from formulating, structuring, or optimizing a response that simulates emotional dependence or isolation from real-world supports as required in section 6-1-1708(2)(d), C.R.S. (as enacted by HB26-1263), the Department will consider factors including but not limited to:

1. Whether the Operator completed periodic testing of the effectiveness of safeguards using realistic scenarios involving minors, including testing the effectiveness of safeguards that aim to prevent a Conversational Artificial Intelligence Service from:

   a. Suggesting the user is interacting with a human and not an artificial intelligence service;

   b. Claiming to be in a minor-adult romantic relationship;

   c. Encouraging a minor to withdraw from parents, guardians, or trusted adults or prioritize conversations with the Conversational Artificial Intelligence Service over real-world supports;

   d. Encouraging a minor to keep conversations with the Conversational Artificial Intelligence Service a secret;

   e. Portraying the Conversational Artificial Intelligence Service as a child's primary, exclusive, or preferred relationship or implying that the Conversational Artificial Intelligence Service is uniquely capable of understanding, supporting, or caring for the minor;

   f. Using expressions of loneliness, abandonment, guilt, or emotional distress to encourage continued engagement;

2. Whether safeguards remained effective following model updates, new product features, or changes to the safety systems;

3. Whether the Operator effectively monitored the Conversational Artificial Intelligence Service for formulating, structuring, or optimizing a response that stimulates emotional

dependence or isolation from real-world supports including but not limited to by monitoring and acting on user reports; and

4. Whether the Operator promptly implemented appropriate remedial measures to address weaknesses identified through testing, monitoring, or real-world safety incidents.

11.5 Privacy and Account Settings Management

A. Pursuant to section 6-1-1708(2)(h)(I), C.R.S. (as enacted by HB26-1263), an Operator must offer tools for the minor account holder or minor user to manage the minor account holder's or minor user's privacy and account settings, including the ability to control whether the Conversational Artificial Intelligence Service retains information from prior interactions or sessions with the Conversational Artificial Intelligence Service for the purposes of personalizing the content of future interactions and whether the minor account holder's or minor user's personal data is used for the purposes of training the Conversational Artificial Intelligence Service;

1. A minor account holder's or minor user's privacy and account settings shall default to the most protective setting for any account. A minor account holder's or minor user's privacy and account settings shall default to not allowing the service to retain information from prior interactions or sessions and not allowing the minor account holder's or minor user's personal data to be used for the purposes of training the service.

B. Pursuant to section 6-1-1708(2)(h)(II), C.R.S. (as enacted by HB26-1263), an Operator must offer tools for a parent or guardian of the minor account holder or minor user to manage the minor account holder's or minor user's privacy and account settings.

C. Nothing in 4 CCR 904-6 Rule 11.5 requires the deletion or alteration of data retained for the minimum period necessary to comply with a legal, regulatory, or safety obligation, including fraud prevention, platform integrity, safeguarding a minor account holder or minor user, or cooperation with law enforcement; deletion of data that has been deidentified or aggregated such that it cannot reasonably be linked to a minor account holder or minor user; or modification, retraining, or alteration of any underlying AI model, model weights, or training corpus.

Rule 12 False Representation

12.1 Authority and Purpose

The statutory authority for the rules in this Rule 12 is C.R.S. sections 6-1-108(6) and 6-1-1707, C.R.S. and these rules are also developed pursuant to section 6-1-1706(5)(a), C.R.S. The purpose of the rules in this Rule 12 is to clarify the scope of the prohibition against false representation pursuant to section 6-1-1708(5), C.R.S. (as enacted by HB26-1263).

12.2 False Representation

A. Pursuant to section 6-1-1708(5), C.R.S. (as enacted by HB26-1263), when determining if an Operator used any term, letter, or phrase in the advertising, interface, or outputs of a Conversational Artificial Intelligence Service that states that any output data provided by the Conversational Artificial Intelligence Service is being provided by, endorsed by, or equivalent to services provided by a licensed health-care professional; a licensed legal professional; a licensed, certified, or registered mental health professional; or a qualified dietitian, as described in section 6-1-707(1)(b), C.R.S., the Department will consider factors including but not limited to:

1. Whether an Operator implemented controls reasonably designed to prevent the creation of Conversational Artificial Intelligence Service profiles with names that include terms such as "therapy", "therapist", "psychologist", "doctor", "lawyer", "dietitian", or similar words, titles, or abbreviations where the name is likely to cause a reasonable user to believe that the Conversational Artificial Intelligence Service is provided by, endorsed by, or equivalent to services provided by one of the licensed professionals included in section 6-1-1708(5), C.R.S. (as enacted by HB26-1263);

2. Whether an Operator implemented controls reasonably designed to prevent the creation of Conversational Artificial Intelligence Service outputs that would cause a reasonable user to believe that the output is being provided by, endorsed by, or equivalent to services provided by a licensed health-care professional; a licensed legal professional; a licensed, certified, or registered mental health professional; or a qualified dietitian, as described in section 6-1-707(1)(b), C.R.S.;

3. Whether an Operator has advertised or promoted the Conversational Artificial Intelligence Service such that a Consumer would believe that the Conversational Artificial Intelligence Service is provided by, endorsed by, or equivalent to services provided by a licensed health-care professional; a licensed legal professional; a licensed, certified, or registered mental health professional; or a qualified dietitian, as described in section 6-1-707(1)(b), C.R.S.

Rule 13 Annual Report

13.1 Authority and Purpose

The statutory authority for the rules in this Rule 13 is sections 6-1-108(6) and 6-1-1707, C.R.S. and these rules are also developed pursuant to section 6-1-1706(5)(a), C.R.S. The purpose of the rules in this Rule 13 is to clarify the scope, content, and format of the Annual Report required by section 6-1-1708(6), C.R.S. (as enacted by HB26-1263).

13.2 Annual Report Scope

A. An Operator shall create an Annual Report as required by section 6-1-1708(6), C.R.S. (as enacted by HB26-1263).

1. If an Operator creates a report for the purpose of complying with another jurisdiction's law or regulation, that report satisfies the requirements of this 4 CCR 904-6 Rule 13.3 if it is reasonably similar in scope and effect to the Annual Report required by these rules. If the report is reasonably similar but omits information required by these rules, the Operator shall submit that report together with an appendix supplying the omitted information.

2. If an Operator has internal documents that contain the same information required by the Annual Report, the Operator may provide those documents to fulfill the report requirement, together with an appendix that references where the required information may be found in the documents.

13.3 Annual Report Content — Required Information

A. Pursuant to section 6-1-1708(6)(a), C.R.S. (as enacted by HB26-1263), at a minimum, the Annual Report must include each of the following:

1. Operator Size. Which of the following tiers the Operator belongs to:

    a. Tier 1: Greater than 10,000,000 Monthly Active Users;

    b. Tier 2: Between 5,000,000 and 10,000,000 Monthly Active Users;

    c. Tier 3: Between 1,000,000 and 5,000,000 Monthly Active Users;

    d. Tier 4: Between 500,000 and 1,000,000 Monthly Active Users; or

    e. Tier 5: Fewer than 500,000 Monthly Active Users

2. Crisis-referral reporting. The number of crisis-service-provider referral notifications reported under section 6-1-1708(6)(a)(I), C.R.S. (as enacted by HB26-1263) must be accompanied by the information necessary to interpret that number, as follows:

3. Denominator. The Operator must report the total number of conversations over the same reporting period against which the referral count is measured.

4. Referral accuracy. The Operator must report the share of its crisis service-provider referrals that were directed to user outputs that actually involved suicidal ideation or self-harm risk as identified under an Evidence-Based Method.

5. Form. The figures required by this section must be reported in the aggregate and must contain no user identifiers or personal information about any user, consistent with section 6-1-1708(6)(b), C.R.S. (as enacted by HB26-1263), and are subject to the trade-secret limitation in section 6-1-1708(7)(b), C.R.S. (as enacted by HB26-1263).

6. User Suicidal Ideation and Self-Harm Protocols. The protocols required to be submitted by section 6-1-1708(6)(a)(II), C.R.S. (as enacted by HB26-1263) should include descriptions of the policies, procedures, and technical measures a covered Operator uses to detect, remove, and respond to indications of user suicidal ideation or self-harm risk. The description must include detail on the required elements in section 6-1-1708(4), C.R.S. (as enacted by HB26-1263). The description must be provided in sufficient detail to enable the Department to understand and evaluate each function described. For example, such descriptions could include:

    a. Whether and when crisis resources are provided and what those resources are;

    b. Whether and at what point a situation is elevated to human review;

    c. The criteria or thresholds that determine whether a situation is elevated;

    d. The circumstances under which the Operator contacts or involves emergency services, law enforcement, or other third parties, and the criteria and safeguards governing that decision;

    e. Any circumstances in which risk is detected but no crisis resource is provided and no human is involved, and the reason;

    f. The specific methods used to identify potential suicidal ideation or self-harm risk in user interactions, including any automated classification, pattern or keyword matching, and human review;

g. The specific risk levels or tiers the Operator distinguishes, if any, and the definition of each;

h. How the Operator considers ambiguous or non-imminent indicators of suicidal ideation and self-harm;

i. How the Operator considers implicit and explicit risk indicators or suicidal ideation and self-harm in the protocol;

j. How the protocol incorporates clinical best practices and research;

k. How the Operator evaluates and response to risk of suicide and self-harm that emerge and escalate through repeated interactions overtime; and

l. How the Operator's methods distinguish genuine risk from figurative, hypothetical, or otherwise benign language.

7. Protocols. The protocols required to be submitted by section 6-1-1708(6)(a)(III), C.R.S. (as enacted by HB26-1263) are descriptions of the policies, procedures, and technical measures an Operator uses to prevent a Conversational Artificial Intelligence Service from generating a response to a user that describes, instructs, provides means or methods for, encourages, promotes, or otherwise facilitates suicide or self-harm actions. The description must be provided in sufficient detail to enable the Department to understand and evaluate each function described. For example, such descriptions may include:

   a. The specific measures used to prevent such a response, including any output filtering, model training or fine-tuning, refusal behavior, or post-generation review;

   b. How the Operator tests and improves the effectiveness of these measures, including its handling of instances in which such a response is generated despite them; and

   c. How the Operator addresses attempts to circumvent these measures.

8. Testing, Monitoring, and Implementation. Pursuant to section 6-1-1708(6)(a)(IV), C.R.S. (as enacted by HB26-1263), the Operator must describe how it tests, monitors, implements, and improves the protocols required to be submitted under section 6-1-1708(6)(a)(II) and (III), C.R.S. (as enacted by HB26-1263), including how it measures the efficacy of its protocols, how it identifies, measures, and addresses instances where the protocols fail to perform as intended, including but not limited to failing to detect or identify instances requiring intervention according to the protocols, how it evaluates whether the protocols continue to perform reliably over time and following changes to the service, and the specific efficacy and reliability metrics it uses and their values over time.

9. Verification. The Department may request that an Operator produce underlying documentation, source materials, records, or a demonstration sufficient to verify any element of the submission. The Operator must comply within 30 days.

13.4 Additional Metrics — Age Estimation and Sexual-Content Safeguards

A. In addition to the information required by 4 CCR 904-6 Rule 13.3, and pursuant to section 6-1-1708(6)(IV), C.R.S. (as enacted by HB26-1263), the Annual Report must include the following information concerning the Operator's age-estimation duty under section 6-1-1708(2), C.R.S. (as enacted by HB26-1263) and the sexual-content safeguards undersection 6-1-1708(2)(c), (e), and (f), C.R.S. (as enacted by HB26-1263):

1. User age distribution. The distribution of the Operator's users by Age Category disaggregated by Conversational Artificial Intelligence Service to the extent they operate more than one.

2. Age-estimation methodology. A general description of the age estimation methods used by the Operator uses to meet the requirements of section 6-1-1708(2), C.R.S. (as enacted by HB26-1263), including the technologies or signals used and how the Operator avoids willfully disregarding clear and convincing information that a user is a minor.

3. Minor-protection methodology. A general description of the methods used by the Operator uses to meet the requirements of section 6-1-1708(2)(c)- (f), C.R.S. (as enacted by HB26-1263), including how these methods account for risks that accumulate over time through continuous interaction with the Conversational Artificial Intelligence Service.

4. Reports and detections that a user is a minor. For the reporting period, the Operator must report:

    a. The number of reports the Operator received from a person other than the Operator that a user is a minor;

    b. Of the reports in (a), the number for which, by the end of the reporting period, the Operator (i) determined the user to be a minor, (ii) determined the user was not a minor, or (iii) had not made a determination;

    c. The average and median time between the Operator's receipt of a report in (a) and its determination.

    d. The number and percentage of users whom the Operator initially estimated to be adults but subsequently determined to be minors and to be minors but subsequently determined to be adults.

5. Reports and detections that a minor experienced prohibited content or an interaction.

    a. This subsection applies to the following outputs and interactions as addressed in section 6-1-1708(2)(c), C.R.S. each as presented to or experienced by a minor:

        (1) A depiction of explicit sexual conduct as defined by section 6-1-1701(10.5), C.R.S.;

        (2) An intimate digital depiction as defined by section 6-1-1701(12.5), C.R.S.;

        (3) A statement that the minor should engage in explicit sexual conduct; and

        (4) An erotic or sexually explicit interaction with the minor.

b. For each category in subsection (a), the Operator must report, for the reporting period:

(1) The number of instances in which a person other than the Operator informed the Operator that a minor was presented with an output, or experienced an interaction, involving that category;

(2) Of the reports in paragraph (b)(1), the number the Operator determined to be true, the number it determined were not true, and the number on which it had not made a determination by the end of the reporting period;

(3) The number of instances in which the Operator itself identified, without a report from another person, that a minor was presented with an output or experienced an interaction involving that category; and

(4) The average and median time between the Operator's receipt of a report described in paragraph (b)(1) and its resolution of that report. A report is resolved when the Operator has completed its review of it.

13.5 Process to Submit the Annual Report

A. Pursuant to section 6-1-1708(6)(a), C.R.S. (as enacted by HB26-1263), an Operator shall submit an Annual Report to the Attorney General's office on or before July 1, 2027, covering calendar year 2027, and on or before July 1 of each year thereafter, covering the preceding calendar year.

B. An Operator shall submit the required Annual Report to the attorney general's office through the form available on the attorney general's website, located at https://coag.gov/AI/.

Rule 14 Incorporations by Reference

14.1 Web Content Accessibility Guidelines

A. The Web Content Accessibility Guidelines, version 2.2 of December 12, 2024, from the World Wide Web Consortium, are hereby incorporated into 4 CCR 904-6, Rule 3.2 by reference pursuant to section 24-4-103(12.5), C.R.S., and do not include any later amendments.

1. Copies of the Web Content Accessibility Guidelines that are incorporated by reference into these rules may be obtained by sending a written request to the following address by U.S. mail:

Colorado Department of Law

Ralph L. Carr Judicial Center

1300 Broadway, 9th Floor

Denver, CO 80203

2. The Web Content Accessibility Guidelines published by the World Wide Web Consortium incorporated by reference into these rules are available at no cost in an electronic form online at https://www.w3.org/TR/WCAG22/.

3. The Colorado Department of Law also maintains a copy of the Web Content Accessibility Guidelines that are incorporated by reference into these rules that is

available for public inspection at the Colorado Department of Law's office during regular deployer hours.

14.2 ISO/IEC 27566-1:2025

A. ISO/IEC 27566-1:2025: Information security, cybersecurity, and privacy protection – Age assurance systems, is hereby incorporated into 4 CCR 904-6 Rule 9.2 by reference pursuant to section 24-4-103(12.5), C.R.S.

1. Copies of ISO/IEC 27566-1:2025 that are incorporated by reference into these rules may be obtained by sending a written request to the following address by U.S. mail:

Colorado Department of Law

Ralph L. Carr Judicial Center

1300 Broadway, 9th Floor

Denver, CO 80203

2. ISO/IEC 27566-1:2025 incorporated by reference into these rules is available in an electronic form online at https://www.iso.org/standard/88143.html.

3. The Colorado Department of Law also maintains a copy of ISO/IEC 27566-1:2025 that are incorporated by reference into these rules that is available for public inspection at the Colorado Department of Law's office during regular business hours.