

Automated Decision-Making and Municipal Chatbots: Preparing for Colorado's New AI Laws

Municipal operations depend heavily on administrative work, much of it repetitive, time-sensitive, and difficult to scale as workloads increase. Against that backdrop, automation offers a compelling promise of helping staff process more information, respond faster, and devote more attention to work that requires human judgment. Why spend hours sorting applications, routing requests, checking forms, or answering the same questions from scratch when technology can do it faster?

Indeed, municipal employees are increasingly using generative AI to draft correspondence, summarize documents, prepare public information, and assist with administrative tasks. Municipalities are also using automation tools that can screen employment applications, identify potentially eligible benefit recipients, prioritize service requests, evaluate permit materials, and communicate with residents through public-facing chatbots.

While it can be tempting to automate as many tedious, repetitive tasks as possible, overreliance on automated tools may increase risks when efficiency begins to replace judgment. This can be particularly problematic when automated tools are tasked with making or influencing decisions that have a real impact on stakeholders, or when AI-powered communication tools give the user the impression they're speaking to a human.

The Colorado General Assembly responded to those perceived risks by adopting two significant AI laws during the 2026 legislative session: Senate Bill 26-189 ("SB 26-189")¹ and House Bill 26-1263 ("HB 26-1263"). Beginning January 1, 2027,² the Acts will impose new transparency, notice, review, and accountability requirements on "developers" and "deployers" of automated decision-making technologies, and on "operators" of public-facing conversational AI services, extending to both public and private entities meeting those definitions.³

SB 26-189 regulates automated decision-making technology used to materially influence decisions involving employment, housing, financial services, insurance, health care, education, and essential government services and public benefits. The Act does not prohibit municipalities from using AI in consequential decision-making; instead, it emphasizes transparency, accuracy in data, meaningful human review, and accountability for technology-assisted decisions. HB 26-1263 is separately intended to address distinct consumer-safety risks created by publicly available conversational AI systems that simulate human interaction, with particular emphasis on protecting minor users.

When these Acts take effect, the central challenge for municipalities will be to capture the benefits of automation without losing the human oversight and accountability that public decision-making requires. To prepare, municipalities must understand where these Acts' requirements may apply.

SB 26-189: AUTOMATED DECISION-MAKING TECHNOLOGY IN CONSEQUENTIAL DECISIONS

SB 26-189 regulates "automated decision-making technology," or "ADMT," that is used to "materially influence" a "consequential decision." Understanding those terms is critical. An ADMT is a technology that processes personal data⁴ and uses computation to generate an output, including predictions, recommendations, classifications, rankings, scores, or other information that is used to make, guide, or assist a decision, judgment, or determinations about an individual.⁵

The Act excludes many routine technologies and uses. Examples include basic calculators and databases, qualifying spreadsheets, administrative routing, scheduling, translation, summarization, and tools that merely organize or present

information for human review. The Act also excludes certain natural-language tools that provide information, answer questions, make referrals, or generate content when they are not intended, marketed, or configured for making consequential decisions and are subject to an acceptable use policy prohibiting that use.⁶

A “consequential decision” is a decision, determination, or action concerning an individual related to the provision of, or the individual’s access to, eligibility for, selection for, or compensation relating to a “covered domain.” “Covered domains” include educational opportunities, employment opportunities, the lease or purchase of residential housing, financial and lending services, insurance, health-care services, and essential government services and public benefits, including eligibility and renewal determinations.

A decision concerning differentiated pricing, cost sharing, compensation, or other material terms may also qualify as a consequential decision when the differentiation is reasonably likely to materially limit, delay, effectively deny, or otherwise fundamentally alter the individual’s access or opportunity to a covered domain.

For example, suppose a municipal housing authority operates a housing program that owns dwelling units that it offers to qualifying applicants at below-market rental rates. The housing authority uses a technology that assigns applicants a “needs and risk assessment score” based on financial information included in the application. The technology then uses the score to recommend approval or denial of the application. If the authority approves the application for placement in a unit but the authority’s housing stock is limited, the technology recommends the applicant’s “place in line” on a waitlist based on the applicant’s needs and risk assessment score.

While a decision to place an already-approved applicant on a waitlist does not formally deny assistance, the decision may result in a material delay or effective denial of access to residential housing—particularly if the score assigned to the applicant places them far down on the waitlist—and thus may rise to the level of a consequential decision.

Suppose, instead, that the housing authority offers rent payment assistance to qualifying applicants, and uses the technology to recommend how much assistance to provide to individual applicants based on their needs and risk assessment score. That recommendation, too, could rise to the level of a consequential decision because it involves a decision concerning differentiated compensation that could materially limit the individual’s access to the program funding.

On the other hand, the definition of “consequential decision” expressly does not include low-stakes or routine decisions, actions, and business processes that do not materially influence eligibility for, selection for, denial of, compensation for, pricing of, or access to an opportunity or service for a covered domain, including routine scheduling, administrative routing, or workflow management.

A technology “materially influences” a consequential decision when its output is a “non-de minimis factor” affecting the decision-making process or its outcome. This includes functions such as filtering, ranking, scoring, recommending, or classifying individuals as part of the decision-making process for a consequential decision. This means the presence of a human decision-maker does not automatically remove the system from the Act’s application, even if the human will make the final decision related to the covered domain.

Employment-Related Decisions. Importantly, while employment and employment opportunities are a covered domain, municipalities will want to review closely with their own counsel whether SB 26-189’s notice, disclosure, correction, and human-review requirements apply to municipal use of ADMT in decisions impacting this domain. The Act incorporates the Colorado Wage Act’s definition of “employer,” which excludes municipal corporations and certain other public entities. Accordingly, a municipality’s use of automated technology in hiring or other employment decisions could potentially fall outside SB 26-189’s employment-related coverage. However, Attorney General rulemaking or future judicial interpretation may clarify the issue.

Nevertheless, use of automated employment screening tools poses risks. Suppose, for example, that a municipality uses an applicant-tracking platform that parses résumés, processes applicants' personal information, assigns applicants a score, ranks them against the position's qualifications and other candidates, and recommends which candidates should receive interviews. Even if SB 26-189's requirements do not apply to a municipality's use of this type of employment screening tool, SB 26-189 expressly preserves existing rights and remedies available under state or federal law, including the Colorado Anti-Discrimination Act (CADA). A municipality could, therefore, face liability if its selection, configuration, or reliance on an automated screening tool causes unlawful discrimination in employment, including where the system screens out qualified applicants based on criteria that disproportionately affect a protected group.

Obligations for Deployers of ADMT. SB 26-189 regulates "developers" that create or commercially provide covered ADMT and "deployers" that use it. The Act defines a "deployer" as a person doing business in Colorado that deploys covered ADMT. Depending on the context, a municipality could meet either definition.⁷ Though, in most use cases, a municipality is more likely to be considered a deployer of ADMT.

The principal obligations of a deployer of ADMT concern notice, adverse outcomes, data access and correction, human review, and record retention. Before using covered ADMT to materially influence a consequential decision, the deployer must provide clear and conspicuous notice that covered technology will be used in a consequential decision affecting the individual and provide instructions about how the affected individual may obtain additional information.⁸ The notice requirement is met if a prominent notice is maintained at accessible points of the individual's interaction. For example, if ADMT will be used in the hiring process, it may suffice to include the notice within the employment portal or on the application form.

If the technology materially influences a consequential decision resulting in an "adverse outcome," the deployer must make certain "post-adverse outcome disclosures" to the affected individual within 30 days of making the decision. An "adverse outcome" includes denying, terminating, revoking, or materially restricting an individual's access, eligibility, selection, compensation, opportunity, or service. It also includes decisions resulting in materially less favorable pricing, compensation, cost-sharing, or other terms that are reasonably likely to limit, delay, deny, or fundamentally alter an individual's access to an opportunity or service compared with similarly situated individuals.

The post-adverse outcome disclosure must include a plain-language explanation of the decision and the role the covered ADMT played. It must also explain how the individual may request available information regarding the system, its developer, its version, and the types and sources of personal data used, to the extent the deployer receives the necessary information from the developer as required of the developer under the Act.⁹

The disclosure must also include notice about the individual's applicable rights and instructions about how to exercise them, including how the individual may access the personal data used and request correction of factually incorrect or materially inaccurate personal data, and of the individual's right to request meaningful human review of the decision.¹⁰ While a deployer is not required to "correct" opinions, predictions, scores, or protected evaluations made by the technology, the deployer may need to correct an inaccurate employment date, income amount, address, credential, or other factual input that affected the decision, when requested by the affected individual.

Upon the affected individual's request, the deployer must provide the affected person with an opportunity for meaningful human review and reconsideration following an adverse outcome, to the extent commercially reasonable.¹¹ Meaningful review requires more than confirming that the technology generated a particular score. The reviewer must have appropriate training, understand the system's intended use and material limitations, and have the authority to consider relevant primary evidence and approve, modify, or override the decision.

Deployers must retain records reasonably necessary to demonstrate compliance for three years after the consequential decision upon which ADMT had a material influence, or for a longer period if required by applicable state or federal law.

Enforcement of SB 26-189. The Act’s developer and deployer requirements are enforced exclusively by the Colorado Attorney General. A violation constitutes a deceptive trade practice under the Colorado Consumer Protection Act, allowing the Attorney General to pursue the remedies available under that law, including injunctive relief and civil penalties. Until January 1, 2030, the Attorney General generally must provide a developer or deployer with written notice and 60 days to cure an alleged violation before filing an enforcement action, if the Attorney General determines that a cure is possible. But that cure opportunity is not required for knowing or repeated violations. The Act does not create a new private right of action, meaning an individual cannot directly sue for civil liability for a violation of the Act.

However, the absence of a new private action does not insulate a municipality from liability under other laws. SB 26-189 expressly preserves existing rights and remedies available under state or federal law, including claims alleging unlawful discrimination arising from a consequential decision materially influenced by covered ADMT. For municipalities acting as deployers, potential liability may arise from their own selection, configuration, oversight, or use of the technology under laws such as the Colorado Anti-Discrimination Act, Title VII of the Civil Rights Act of 1964, the Americans with Disabilities Act, the Age Discrimination in Employment Act, and other federal, state, or local civil-rights, employment, housing, disability-access, and program-specific laws.

HB 26-1263: CONVERSATIONAL ARTIFICIAL INTELLIGENCE SERVICE OPERATOR REQUIREMENTS

HB 26-1263 applies to those who develop, make publicly available, or offer to a consumer (i.e., an “operator”) a “conversational artificial intelligence service.” That term is defined as an AI system that is accessible to the general public and that primarily simulates human conversation through adaptive text, visual, or audio communications.

The Act excludes several categories of tools from the definition of a “conversational artificial intelligence service,” which clarify the Act’s primary focus on open-ended generative systems designed to sustain adaptive, human-like conversations. Most relevant to municipalities, the Act excludes tools used solely for internal business purposes; systems designed primarily for business operations, productivity, information analysis, internal research, training, or technical assistance; narrow-topic tools that cannot generate sexually explicit content or sustain dialogue concerning suicide or self-harm; and commerce-related or transactional assistants used for functions such as payments, ordering, delivery, returns, customer support, or customer service. For example, a scripted tool that retrieves fixed information, routes a resident to a webpage, or performs a narrow transaction may not fall within the Act’s coverage.

Accordingly, not every municipal chatbot will qualify, and it may be that most will not—at least not those that are commonly deployed.¹² It seems unlikely that a municipality would operate the type of open-ended “AI companion” that appears to be the Act’s principal focus. However, municipalities should assess the chatbot’s actual functions rather than relying on the vendor’s label or the municipality’s informational intent, as it’s possible municipal chatbots may fall within the Act’s coverage, particularly as tools become more adaptive and conversational.

For covered systems, operators generally must clearly disclose that the user is interacting with AI at the beginning of the user’s first interaction each day, at least once every three hours during a continuing interaction unless the disclosure remains persistently visible, and whenever the user asks whether the service is human or artificially generated. HB 26-1263 imposes additional requirements intended to protect the safety of minors and further mental health and suicide prevention objectives, including protocols for suicidal ideation and self-harm, restrictions on misleading professional-service claims, and age-estimation and other safeguards for known minor users.

When a Chatbot May Trigger Both Acts. While it may be that most municipal chatbots will not fall within HB 26-1263's coverage, it's possible a chatbot could implicate SB 26-189's coverage. On the one hand, a chatbot used only to retrieve general information should not ordinarily be covered ADMT under SB 26-189. This is because natural-language technologies used to provide information, answer questions, or make referrals are expressly exempt from SB 26-189's coverage when the chatbot is not intended or configured for consequential decisions and is subject to an acceptable-use policy prohibiting such use.

The result may change if the chatbot collects personal data and begins determining probable eligibility, calculating benefit amounts, ranking programs, recommending approval or denial, prioritizing service requests, or generating individualized recommendations on which municipal employees rely. At that point, it's possible the chatbot could implicate both Acts: HB 26-1263 because it may simulate public-facing conversation, and SB 26-189 because it may materially influence consequential decisions.

Municipalities can reduce this risk by maintaining clear informational boundaries. Public-facing chatbots should not access personal accounts or protected databases, make eligibility determinations, or provide individualized legal or professional advice unless the municipality has intentionally designed the process to satisfy applicable legal requirements. Chatbots should prominently identify themselves as AI, direct users to authoritative sources, provide access to human assistance, support accessibility and language access, and operate under clear retention and acceptable-use rules.

Finally, a disclaimer is helpful but not decisive. A municipality cannot characterize an AI system as "informational only" if, in practice, employees or residents rely on the AI system to determine access to municipal services or benefits.

CONCLUSION

Colorado's new AI laws do not prevent, nor should they discourage, municipalities from pursuing tools that can improve efficiency, expand access, and reduce the burden of repetitive administrative work. But the Acts reinforce that when technology begins to influence decisions with real consequences for real people, efficiency cannot come at the expense of transparency, accuracy, human judgment, or accountability.

For municipalities, the key is to understand what their automated and artificial intelligence tools actually do. In all cases, the applicability of Colorado's new AI laws depends less on how a product is labeled and more on how it functions, how employees use its output, and what effect it has on individuals.

Importantly, the most significant municipal risk may not arise from deliberately purchasing a product advertised as an "AI decision-maker" or "adaptive AI companion." It may arise, instead, from a feature already embedded in familiar, existing software program. Indeed, an existing system may appear to be merely administrative, informational, or supportive, yet may quietly rank applicants, prioritize residents, recommend outcomes, or otherwise shape consequential decisions. Likewise, while a chatbot may seem like a routine customer-service feature, it may evolve into an open-ended "AI companion" as it grows more adaptive, conversational, and personalized.

Ultimately, however, the practical response to Colorado's new AI laws is not to avoid AI tools altogether, but to govern their use deliberately. In preparation for the Acts' implementation date, municipal officials should implement or update existing policies to preserve meaningful human authority, demand sufficient information from vendors, avoid unexplained automated rejection, and maintain clear limits on public-facing chatbots. Municipal officials should also inventory existing systems, identify where automated outputs affect consequential decisions, establish meaningful human-review procedures, and require sufficient vendor transparency.

Click [here](#) for a practical checklist municipalities can use to begin evaluating current systems and otherwise prepare for SB 26-189's January 1, 2027, effective date. Because application of the Act is fact-specific—and Attorney General rulemaking may provide further guidance—municipalities should consult with their attorneys before implementing automated technology that may deny, delay, rank, or materially affect an individual opportunity, service, or benefit.

This article is intended for general informational purposes only and is not intended or to be construed as legal or professional advice on any specific issue. Municipal officials should consult with their entity's own counsel concerning application of the Acts to particular technologies and municipal processes.

1. SB 26-189 repeals and replaces the framework originally enacted through Senate Bill 24-205. The new Act is narrower in some respects and more operationally focused. Rather than broadly regulating “high-risk artificial intelligence systems,” it applies to “automated decision-making technology,” or “ADMT,” that is used to “materially influence” a “consequential decision” in a “covered domain.”
2. SB 26-189's developer and deployer requirements apply beginning January 1, 2027. HB 26-1263's principal operator obligations likewise become effective January 1, 2027.
3. While there may be some question about the applicability of SB 26-189 to municipal organizations (e.g., due to the Act's placement in the Colorado Consumer Protection Act of Title 6 of the Colorado Revised Statutes, and the Title 6 definition of “person”), the Act does not contain an express, general exemption for municipalities or other government entities. Moreover, the Act expressly applies to domains within which municipalities have significant authority (e.g., essential government services and public benefits). Accordingly, unless and until a court weighs in on the applicability issue, municipal officials should seek to determine whether and how the Act applies to its technologies instead of assuming that public-entity status exempts them from the Act.
4. “Personal data” means information that is linked or reasonably linkable to an identified or identifiable individual; it does not include de-identified data or publicly available information. “Publicly available information” means information that is lawfully made available from federal, state, or local government records and information that a controller has a reasonable basis to believe the consumer has lawfully made available to the general public. C.R.S. §§ 6-1-1701(16) & 6-1-1303(17).
5. A product need not be marketed as “artificial intelligence” or an “ADMT” to satisfy this definition. Applicant-tracking software, eligibility-scoring systems, automated recommendation tools, and other familiar platforms may qualify based on their actual functions.
6. These exclusions are function specific. A tool initially used only for summarization or administrative processing may become covered if the municipality later configures or uses it to generate an inference that materially influences a consequential decision.
7. See *supra* endnote 3.
8. Required notices and disclosures must also be reasonably accessible to individuals with disabilities and individuals with limited English proficiency. Municipalities should also consider the applicability and requirements of Colorado's technology accessibility laws and regulations to notices and disclosures when hosted or transmitted digitally.
9. Developers must provide deployers with documentation addressing intended uses, known limitations and inappropriate uses, categories of training data, appropriate-use instructions, monitoring, and meaningful human review. Developers must also notify deployers of material updates and intentional and substantial modifications to their technologies.
10. Because appropriate post-adverse outcome disclosures may differ by subject area, the Act directs the Attorney General, by January 1, 2027, to clarify the required content and format of the post-adverse outcome disclosure through sector-specific rulemaking while ensuring that consumers receive meaningful and understandable information.
11. The Act does not define or otherwise provide guidance as to what “commercially reasonable” may mean in this context, but it's possible the issue will be addressed by the Attorney General in the Act's implementing rules and regulations.
12. Denver's “Sunny” chatbot provides a useful Colorado example. Denver describes Sunny as an AI-powered virtual assistant that answers questions regarding city services, assists users in reporting municipal issues, communicates in numerous languages, and provides informational, but not official, responses. Denver also states that Sunny cannot access personal accounts or information.

Published on 07/30/2026