Public Comment on Proposed Rules for the Colorado Automated Decision-Making Technology Act

Re: 4 CCR 904-6 — Training Data Transparency Through Independent Certification

Submitted by: Box Commons 30 N Gould St Ste N Sheridan, WY 82801 Contact: Brice Love, Founding Catalyst Email: info@boxcommons.org Web: boxcommons.org

Date: September 23, 2026

Submitted via: comments.coag.gov

Docket: Proposed Rules for the Automated Decision-Making Technology Act (SB 26-189) and the Chatbot Safety Act (HB 26-1263), 4 CCR 904-6

I. Executive Summary

Box Commons respectfully submits this comment on the proposed rules implementing the Automated Decision-Making Technology Act (SB 26-189) and the Chatbot Safety Act (HB 26-1263), scheduled for formal rulemaking hearing on October 26, 2026.

Box Commons is a 501(c)(6) standards body that develops technology-agnostic certification standards for data used in AI systems. We are governed by a three-chamber structure (Industry, Civil Society, Academic/Research) modeled on the Forest Stewardship Council to prevent capture by any single interest group. Our intellectual property policy defaults to royalty-free licensing. We do not advocate for or against regulation; we build the measurement and certification infrastructure that makes compliance verifiable.

Our central recommendation: The proposed rules should recognize independent, third-party data certification as a structured mechanism by which developers can satisfy their training data documentation obligations under Proposed Rule 4 and deployers can satisfy their vendor oversight duties. Specifically, we recommend that the final rules establish that a developer's use of datasets certified by an independent standards body — one that meets defined governance criteria — creates a rebuttable presumption that the developer has provided adequate documentation of "the categories of data, including personal data, used to train the Covered ADMT" as required by Section 6-1-1706(1)(a)(IV) of the Act.

This comment draws on our experience operating a live audio data certification pipeline across a 22-station broadcast radio network, which demonstrates that the training data transparency SB 26-189 contemplates is technically feasible, operationally affordable, and independently verifiable today.

II. The Training Data Documentation Gap

SB 26-189 requires developers to provide deployers with documentation including "the categories of data, including personal data, used to train the Covered ADMT." The proposed rules elaborate on the format, accessibility, and delivery of this documentation. However, neither the statute nor the

proposed rules address a structural problem: **how the deployer — or the Attorney General — can verify that the developer's training data documentation is accurate.**

A. The Verification Asymmetry

The training data documentation requirement creates a familiar information asymmetry. Developers know what data they used to train their systems. Deployers and regulators do not. The documentation requirement creates a disclosure obligation, but disclosure without verification is attestation — the developer's word, unaudited.

This asymmetry is not theoretical. The insurance industry has already encountered it. When ISO released endorsement forms CG 40 47 and CG 40 48 in January 2026 — excluding AI-related bodily injury, property damage, and advertising injury from commercial general liability coverage — the exclusions exposed a consent verification gap that the insurance market has been unable to resolve through attestation alone. Gallagher Re documented a 978.1% increase in generative AI-related lawsuits from 2021 to 2025, with over 700 cumulative filings between 2020 and 2025. Insurers in the emerging affirmative AI coverage market — Armilla (up to $25 million, Lloyd's-backed), Munich Re's aiSure program (up to €15 million, performance cover), and Testudo (up to $9.25 million, mid-market) — now require data provenance audits as a condition of coverage because self-attestation proved insufficient to quantify risk.

The pattern is clear: wherever AI training data intersects with legal liability, attestation alone fails. Independent certification fills the gap.

B. The Colorado-Specific Urgency

Colorado is uniquely positioned to address this gap because of the relationship between SB 26-189 and the existing insurance-specific algorithmic discrimination framework under SB 21-169 (C.R.S. § 10-3-1104.9) and DOI Regulation 10-1-1 (3 CCR 702-10).

Under the deeming provision in Section 6-1-1708, insurers complying with SB 21-169 are deemed compliant with SB 26-189 "regarding the practice of insurance." However, Regulation 10-1-1 requires a documented governance framework for the use of External Consumer Data and Information Sources (ECDIS), algorithms, and predictive models. Insurers must file attestations confirming their frameworks are operational — life insurers since December 2024, and auto and health insurers by July 1, 2026.

These governance frameworks demand that insurers document and control the data flowing into their algorithmic decision tools. But when an insurer uses a third-party vendor's AI system, the insurer's governance framework is only as good as its ability to verify the vendor's training data practices. This is the same vendor oversight bottleneck we identified in our April 2026 comment to the NAIC Big Data and Artificial Intelligence Working Group regarding the AI Systems Evaluation Tool.

Independent data certification resolves this bottleneck for both the insurance deeming pathway and the general ADMT compliance pathway.

III. How Independent Data Certification Works

A. The Three-Layer Architecture

Independent data certification interposes a trusted intermediary between the data producer, the ADMT developer, and the regulator:

Layer	Role	Function
Standards Body	Publishes technology-agnostic certification standards; accredits third-party certifiers; maintains a public registry of certified datasets	Sets the rules
Accredited Certifiers	Conduct independent audits of datasets against published standards; issue certifications with defined scope and validity	Applies the rules
Certified Data Producers	Submit datasets for audit; receive certification marks that travel with the data through the supply chain	Follow the rules

Under this architecture:

1. The **data producer** compiles training data with documented consent chains, content separation controls, and a complete audit trail.
2. An **accredited certifier** independently verifies these claims through sampling, testing, and documentation review.
3. The **ADMT developer** incorporates certified datasets and documents the certification scope, certifier identity, and validity dates as part of its training data documentation.
4. The **deployer** relies on the certification to satisfy its vendor oversight obligations, with verification available through a publicly accessible registry.
5. The **Attorney General** can verify compliance by checking the same public registry — no inspection of proprietary models or trade secrets required.

B. Mapping to SB 26-189 Developer Obligations

SB 26-189 Developer Obligation	How Independent Certification Satisfies It
Categories of data used to train ADMT (§ 6-1-1706(1)(a)(IV))	Certification manifest documents all data categories, source types, and scope

SB 26-189 Developer Obligation	How Independent Certification Satisfies It
Description of intended uses (§ 6-1-1706(1)(a)(I))	Certification specifies permitted downstream uses as a condition of the standard
Known limitations and risks (§ 6-1-1706(1)(a)(II))	Certification audit identifies data quality limitations and exclusion gaps
Instructions for monitoring (§ 6-1-1706(1)(a)(III))	Certification tier defines ongoing surveillance and recertification requirements
Three-year record retention (§ 6-1-1706(1)(b))	Certification audit trail is retained for the greater of active use plus three years or seven years

C. Mapping to Regulation 10-1-1 Governance Framework

For insurers operating under the SB 21-169 deeming provision, independent data certification also maps to the DOI's governance framework requirements:

Regulation 10-1-1 Requirement	How Independent Certification Satisfies It
Governance framework for ECDIS use	Certification standards provide the governance structure for training data
Risk management for algorithms and predictive models	Certification tiers impose escalating audit rigor aligned to risk
Annual attestation (life; auto/health by July 2026)	Certification status is publicly verifiable; surveillance audits provide continuous assurance
Prohibition on unfair discrimination via protected characteristics	Consent documentation and audit trails enable traceability from output back to input

---

IV. Proof of Concept: Audio Data Certification in Operation

We do not offer this recommendation abstractly. Box Commons currently operates a live audio data certification pipeline — the BC-Certified Audio Data Standard — validated against a 22-station broadcast radio network.

A. What BC-Certified Covers

BC-Certified certifies audio datasets compiled from broadcast radio, podcasts, ministry programming, and community radio for downstream uses including AI model training, speech recognition development, and content analysis. The standard requires:

1. **Consent Documentation.** Every audio segment in a certified dataset is traceable to a documented consent authorization from the applicable rights holder, specifying permitted downstream uses. Consent records follow a structured schema covering rights holder identity, content scope, permitted and prohibited uses, grant and expiration dates, and revocation terms.

2. **Copyrighted Music Exclusion.** Certified datasets demonstrate that copyrighted musical content has been identified and separated using at least one automated method (audio

fingerprinting, playout log reconciliation, or classifier-based detection), with defined minimum accuracy thresholds. Broadcast music is licensed for over-the-air transmission under ASCAP/BMI/SESAC blanket licenses — those licenses do not authorize reproduction in AI training datasets.

3. **Audit Trail.** A complete, tamper-evident record of all operations from source ingestion through dataset finalization, including source hashing (SHA-256), consent verification events, music scan results, segment creation, and metadata attachment.

B. The Operational Proof

This is not a specification on paper. Box Commons operates the following tools against live broadcast streams:

- **Stream capture** with ICY metadata and hash computation at ingestion
- **Playout log parsing** (StationPlaylist) for scheduled-versus-actual reconciliation — validated at 95.5% match rate across three days of continuous broadcast, with zero missed advertisements and zero unauthorized insertions
- **Audio fingerprinting** (Chromaprint) for copyrighted music identification against a reference database
- **Schedule reconciliation** for discrepancy flagging and gap detection
- **Proof-of-performance reporting** with compliance metrics

These tools are open-source. Any deployer, developer, or regulator can inspect, reproduce, or adapt them.

C. Why Audio Matters for ADMT

Audio data is a significant and growing input to ADMT systems deployed in Colorado. Voice-based AI systems — customer service automation, call center analytics, insurance claims processing via voice, and content moderation — use audio training data that carries consent, copyright, and privacy risks that text and image data do not. The voice rights legislative wave (the federal NO FAKES Act, Tennessee's ELVIS Act, California AB 2602 voiding broad AI consent boilerplate, and Nevada NRS 597.790) is creating new liability exposure for ADMT developers who cannot demonstrate consent-verified provenance for their audio training data.

The proposed rules' requirement that developers document "categories of data, including personal data" used to train ADMT will apply to audio training data. BC-Certified provides the structured certification framework that makes that documentation independently verifiable.

V. The HITRUST Precedent

This regulatory pattern — independent certification as a compliance mechanism — is proven, and the academic case for it in the AI context is now established. Marchant and Buckwald have proposed that states provide a liability shield for AI systems that conform to recognized, comprehensive risk-management standards adopted by established standard-setting organizations. Gary E. Marchant & Jordan Buckwald, *Private Standards as Liability Shields: A Pro-Innovation Artificial Intelligence Regulatory Approach for States*, 27 Minn. J.L. Sci. & Tech. 61 (2026). This comment demonstrates

how that principle applies concretely to one provision of SB 26-189: the training data documentation obligation under Section 6-1-1706(1)(a)(IV).

The operational precedent is equally compelling. The HITRUST Common Security Framework has achieved dominant adoption in healthcare and insurance cybersecurity compliance — over 81% of hospitals and health systems, and 83% of health plans now utilize the HITRUST CSF. HITRUST certification is widely accepted by state insurance departments as evidence of compliance with NAIC Model Law #668 cybersecurity mandates. The NAIC Model Bulletin on the use of AI by insurers, adopted in December 2023, follows the same logic: governance frameworks work when they are independently verifiable, and states are now in various stages of implementing its principles.

BC-Certified replicates this architecture for AI training data:

HITRUST → Cybersecurity Compliance	BC-Certified → Training Data Compliance
HITRUST CSF mapped to NAIC Model Law #668	BC-Certified mapped to SB 26-189 § 6-1-1706 training data documentation
SOC 2 + HITRUST accepted as compliance evidence	BC-Certified accepted as training data documentation evidence
Insurers rely on vendor HITRUST certification	Deployers rely on developer's use of certified training data
Regulators verify via HITRUST public registry	AG verifies via BC-Certified public registry

---

VI. Proposed Rule Language

We respectfully recommend that the final rules include the following guidance (or substantially similar):

> **Training Data Certification Safe Harbor.** A developer may satisfy its obligation under Section 6-1-1706(1)(a)(IV) to provide documentation of "the categories of data, including personal data, used to train the Covered ADMT" by documenting that the training datasets incorporate data certified by an independent standards body that meets the following criteria:
>
> (a) The standards body develops certification standards through a balanced, multi-stakeholder process that prevents dominance by any single interest group;
>
> (b) The standards body's intellectual property policy defaults to royalty-free or fair, reasonable, and non-discriminatory terms;
>
> (c) The standards body separates standards development from certification assessment;
>
> (d) The standards body maintains a publicly accessible registry of all certified datasets, certification scope, and validity periods;
>
> (e) Certification is conducted by accredited third-party assessors demonstrating independence from the data producer; and
>
> (f) The certification scope covers the categories of data, consent documentation, and data quality measures relevant to the ADMT's deployment.

> Reliance on such certification does not relieve the developer of its responsibility for the accuracy of its documentation, but creates a rebuttable presumption of adequate training data documentation when the certification scope aligns with the ADMT's operational use.

This language is deliberately technology-agnostic and does not name any specific standards body. It establishes structural criteria that any qualifying organization could satisfy — including Box Commons, Fairly Trained (which certifies AI models trained on licensed data), and future entrants.

VII. Standing Offer

Box Commons offers to participate, at no cost, in any pilot program, working session, or technical advisory process the Attorney General convenes to develop training data documentation standards. We can demonstrate our certification pipeline against live data, share our open-source tooling, and provide technical input on what training data transparency looks like in operational practice.

We are also prepared to present at the October 26 rulemaking hearing to address technical questions about data certification feasibility.

VIII. About Box Commons

Box Commons is a 501(c)(6) standards body headquartered in Wyoming. We develop technology-agnostic, outcome-based certification standards for data used in AI systems — certifying observable data quality, consent documentation, and audit trail integrity, not mandating specific technologies or architectures.

Governance: Three-chamber system (Industry, Civil Society, Academic/Research) with equal voting weight, modeled on the Forest Stewardship Council. No single entity controls the standards process.

IP Policy: Royalty-free default. All AI credentialing patents from the founding portfolio are pledged RF. Standards never create preferences for any patented technology.

Standards alignment: ISO/IEC 5259-2 and 5259-4 (data quality for ML), C2PA (content provenance), EBUCore (broadcast metadata), IPTC Digital Source Type, IAB Podcast Measurement Guidelines, NIST AI RMF.

Regulatory engagement: Public comments submitted to NAIC (AI Systems Evaluation Tool — Exhibit B/D vendor oversight, April 2026), CPPA (California opt-out mechanisms), Federal Reserve (reputation risk), FAR Council (semiconductor AI clauses), NHTSA (autonomous vehicle incident reporting), and Singapore IMDA (agentic AI governance). NIST AI 800-2 comment acknowledged by the CAISI team (March 2026).

Respectfully submitted,

Brice Love
Founding Catalyst
Box Commons

info@boxcommons.org
boxcommons.org

AI Disclosure: This comment was prepared with AI assistance. Methodology: regulatory analysis of SB 26-189, SB 21-169, and DOI Regulation 10-1-1; synthesis of insurance market research (12 Gemini Deep Research reports, verified against primary sources); structural modeling from HITRUST, SOC 2, and NAIC Model Law #668 precedents; drafted by Claude (Anthropic) under human direction and review. All factual claims were independently verified before inclusion. The analytical framework, strategic positioning, and policy recommendations reflect the judgment of the human author.