FCC Robocall Mitigation Database Comment — Content Provenance as Complement to Call Authentication

Public Comment on Further Notice of Proposed Rulemaking: Improving the Effectiveness of the Robocall Mitigation Database

Re: FCC 26-49, WC Docket Nos. 24-213 and 17-97; CG Docket No. 17-59

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: ECFS (Electronic Comment Filing System), regulations.gov

Docket: Improving the Effectiveness of the Robocall Mitigation Database; Call Authentication Trust Anchor; Advanced Methods To Target and Eliminate Unlawful Robocalls, 91 FR 57454 (September 9, 2026), Document ID 2026-18366

I. Executive Summary

Box Commons respectfully submits this comment on the Further Notice of Proposed Rulemaking (FNPRM) concerning the Robocall Mitigation Database (RMD), call authentication, and robocall elimination methods.

Box Commons is a 501(c)(6) standards body that develops technology-agnostic certification standards for audio data used in AI systems. We are governed by a three-chamber structure (Industry, Civil Society, Academic/Research) with equal voting weight, modeled on the Forest Stewardship Council, to prevent capture by any single interest group. We currently operate a live audio data certification pipeline across a 22-station broadcast radio network, demonstrating that real-time audio provenance is technically feasible, operationally affordable, and independently verifiable.

Our central observation: The FNPRM appropriately strengthens the Robocall Mitigation Database's filing integrity and provider accountability mechanisms. However, the Commission's call authentication framework — including the RMD and STIR/SHAKEN — addresses only the *signaling plane*. It authenticates who is calling and whether the caller has the right to use a telephone number. It does not, and structurally cannot, authenticate the *media plane* — the actual audio content of the call. A call that is fully authenticated under STIR/SHAKEN, properly signed with an A-level attestation, can still deliver audio that is entirely AI-generated, cloned, or manipulated. This media blindness is the architectural vulnerability that AI-enabled robocall fraud exploits.

Our central recommendation: The Commission should recognize independent, third-party content provenance certification as a complement to signaling-layer authentication within the RMD framework. Specifically, we recommend that the Commission:

1. Encourage voice service providers to include content provenance mechanisms in their robocall mitigation plans filed with the RMD;
2. Initiate a proceeding to develop standardized content provenance metadata for voice calls, building on the existing PASSporT extension architecture (RFC 8225) and aligned with C2PA content credential standards;
3. Recognize that independent certification of content provenance practices by a qualifying standards body provides auditable evidence of compliance with AI voice disclosure requirements under the Commission's February 2024 Declaratory Ruling (FCC 24-17) and the pending NPRM on AI-generated calls (FCC 24-84); and
4. Condition any future content provenance requirements on governance criteria that ensure independence, multi-stakeholder participation, and freedom from capture by any single industry interest.

II. The Media Provenance Gap: STIR/SHAKEN's Structural Blind Spot

A. Signaling-Layer Authentication Does Not Equal Content Authentication

STIR/SHAKEN, as defined in IETF RFC 8224 and RFC 8225 and operationalized through the ATIS-1000080 governance framework, is a significant achievement in combating caller ID spoofing. The framework cryptographically authenticates the right of a voice service provider to use a telephone number by generating a PASSporT — a signed JSON Web Token carried in the SIP Identity header — that asserts originating and destination identities, issuance time, and attestation level.

However, STIR/SHAKEN's scope is explicitly limited to the signaling plane. The PASSporT specification defines claims for `orig` (originator identity), `dest` (destination identity), and `iat` (issuance time), but it contains no representation of the audio payload, no cryptographic binding to the media stream, and no mechanism to verify the nature or provenance of the voice content transmitted over the RTP bearer channel. A PASSporT that carries a full A-level attestation confirms that the caller is authorized to use the asserted telephone number. It says nothing about whether the voice on the call is human, AI-generated, cloned from a real person, or manipulated in real time.

This architectural separation was a reasonable design choice when STIR/SHAKEN was developed. Call authentication and media authentication are distinct engineering problems. But the emergence of AI-generated voice fraud has transformed this separation from a design boundary into a systemic vulnerability.

B. The Scale of the Problem

Americans received an estimated 52.5 billion robocalls in 2025, with telemarketing and scam calls accounting for approximately 57 percent of total volume — a 15.4 percent increase over 2024. The FTC reported that consumers lost more than $12.5 billion to fraud in 2024, a 25 percent increase over the prior year. The Commission itself levied a $6 million forfeiture against a political consultant who used AI-generated voice cloning and caller ID spoofing to distribute deepfake robocalls before the 2024 New Hampshire primary (FCC-24-95).

These figures demonstrate that even after broad STIR/SHAKEN deployment and the creation of the RMD, illegitimate traffic remains pervasive — precisely because current controls address the conduit but not the content.

C. The Self-Certification Problem in the RMD

The FNPRM correctly identifies weaknesses in the RMD's self-certification model. Providers file narrative descriptions of their mitigation practices without standardized technical evidence of effectiveness. The Commission has already removed over 1,200 voice service providers from the RMD for deficient certifications, demonstrating that self-certification alone is insufficient to ensure compliance.

As we noted in our April 2026 comment to the Commission regarding UAS trusted supplier frameworks (GN Docket No. 26-74, ECFS 20260422570517252), self-certification faces a structural limitation: it is static at the moment of filing, while the systems it aims to characterize are dynamic. We proposed there that credential-based trust models operating under ANSI/ISO accreditation principles provide a scalable alternative to self-certification. The same principle applies to the RMD: a provider's robocall mitigation plan, described in narrative form and filed at a single point in time, cannot adequately represent the provider's continuously evolving technical capabilities for detecting and blocking illegal traffic.

D. The AI Voice Detection Gap

The Commission's February 2024 Declaratory Ruling (FCC 24-17) correctly classified AI-generated voices as "artificial or prerecorded" under the TCPA, and the pending NPRM (FCC 24-84) proposes disclosure and consent requirements for AI-generated calls. These are necessary regulatory steps. But they create an enforcement asymmetry: the Commission has declared AI voices illegal without consent, yet neither the Commission nor voice service providers possess a standardized, verifiable mechanism to *detect* whether a given call uses AI-generated voice content.

STIR/SHAKEN cannot fill this gap because it operates on the signaling plane. The RMD cannot fill this gap because it collects narrative attestations, not technical evidence of media-plane controls. The result is a regulatory requirement — disclosure of AI voice usage — that depends entirely on the good faith of the caller, with no independent verification infrastructure.

III. Content Provenance as a Complement to Call Authentication

We do not propose replacing STIR/SHAKEN or restructuring the RMD. We propose extending the existing architecture to include content provenance — a media-layer complement to the signaling-layer authentication already in place.

A. C2PA Content Credentials for Telephony

The Coalition for Content Provenance and Authenticity (C2PA), whose members include Adobe, Microsoft, Intel, BBC, and Truepic, has developed technical specifications for cryptographically verifiable content credentials — tamper-evident manifests that describe the origin, history, and transformations of digital assets, including audio. C2PA manifests use CBOR-based assertions, COSE cryptographic signatures (RFC 8152), and trust lists to bind provenance metadata to content.

For static media, C2PA employs "hard binding" — a SHA-256 hash computed over the exact bytes of a completed file. For live audio streams, C2PA supports "soft binding" using perceptual hashes or imperceptible digital watermarks that survive transcoding, codec compression, and packet loss. C2PA's live video specification already supports segment-based manifests for real-time content. This architecture is directly applicable to live telephony audio.

A "content credential" for a voice call would assert:

- Whether the voice content is human-generated, AI-generated (TTS, voice cloning, real-time conversion), or a hybrid;
- The identity of the generating entity or device;
- Whether AI voice detection has been performed, and the detection result (with confidence score, detector identity, and version);
- Whether AI disclosure was delivered to the called party, and at what point in the call; and
- A cryptographic binding (via watermark or perceptual hash) to the actual audio stream, enabling verification that the manifest corresponds to the received audio.

B. Integration Pathways with STIR/SHAKEN

The PASSporT specification (RFC 8225) is explicitly extensible. Custom claims and extension types can be defined within the IETF STIR Working Group. Three viable pathways exist for integrating content provenance into the existing call authentication chain:

Pathway 1 — PASSporT Extension Type. A new `ppt: c2pa` extension type could carry a content provenance claim within the PASSporT payload. This claim would include a URI pointing to the C2PA manifest (hosted on a Call Placement Service or dedicated manifest store) and a cryptographic digest of the manifest for integrity verification. The extended PASSporT would be signed with the provider's existing STI certificate, binding content provenance assertions to the same cryptographic identity as the caller identity assertion.

Pathway 2 — Rich Call Data Integration. RFC 9795 (`draft-ietf-stir-passport-rcd`) extends PASSporT to deliver rich metadata — including caller name, brand logos, and cryptographic digests (`rcdi` claim) — to the terminating handset. Content provenance indicators could be nested within the RCD object, enabling a C2PA content credential to be displayed alongside the caller's brand identity on the consumer's device. This pathway has the advantage of immediate consumer visibility: a "Verified Human Voice" or "AI-Generated — Consent Verified" indicator displayed on the handset alongside the caller's authenticated name and brand.

Pathway 3 — Media Plane Binding via RFC 8862. RFC 8862 (BCP 228) defines best practices for securing RTP media signaled with SIP, including the `mky` (media key) claim in PASSporT, which provides a cryptographic signature over DTLS-SRTP key fingerprints negotiated in SDP. By binding C2PA soft-binding watermark parameters into the SDP and protecting them via the `mky` claim, the architecture establishes a verifiable cryptographic chain from the caller's identity to the acoustic properties of the voice payload.

These pathways are complementary, not mutually exclusive. A phased deployment could begin with Pathway 1 (PASSporT extension for out-of-band manifest reference), advance to Pathway 2 (consumer-facing RCD display), and mature to Pathway 3 (full media-plane binding).

C. Out-of-Band Integration Solves the TDM Problem

A known limitation of STIR/SHAKEN is that SIP Identity headers are stripped when calls traverse legacy TDM network segments via SS7 signaling. RFC 8816 addresses this through Out-of-Band (OOB) SHAKEN, where the originating provider publishes the PASSporT to a cloud-based Call Placement Service (CPS) and the terminating provider queries the CPS to retrieve it.

OOB SHAKEN is the optimal vehicle for content provenance integration. C2PA manifests — containing complex cryptographic chains and extensive metadata — may exceed standard SIP UDP MTU limits. By publishing manifests alongside PASSporT tokens to the CPS, voice service providers can deliver full content provenance data without the size constraints of in-band SIP signaling and without the TDM stripping problem. The terminating provider retrieves both the signaling-plane authentication (PASSporT) and the media-plane provenance (C2PA manifest) from a single query.

IV. The Insurance Forcing Function

The Commission's robocall mitigation framework operates against a backdrop of rapid change in the insurance market that creates independent economic pressure for content provenance certification.

A. ISO CG 40 47/48 and AI Call Analytics

As we documented in our September 2026 comment to the Colorado Attorney General regarding ADMT rulemaking (4 CCR 904-6), Verisk/ISO released exclusionary endorsements CG 40 47 and CG 40 48 in January 2026, broadly excluding AI-related liability from standard commercial general liability (CGL) coverage. CG 40 47 is categorical: it strips Coverage A (bodily injury and property damage) and Coverage B (personal and advertising injury) for any claim "arising out of" generative AI. Industry estimates indicate 95 percent carrier adoption of these forms.

These exclusions apply directly to voice service providers deploying AI-based call screening, analytics, or robocall detection systems. Providers are increasingly required by the Commission to deploy advanced analytics to identify and block illegal robocalls. Yet their insurance carriers are simultaneously excluding coverage for liabilities arising from those exact AI systems.

B. False-Positive Blocking Creates Uninsured E&O Exposure

Under Sections 201 and 202 of the Communications Act, common carriers face liability for unjust or unreasonable discrimination in the provision of services. If a provider's AI analytics system generates a false positive — blocking a legitimate healthcare appointment reminder, bank fraud alert, or flight cancellation notice — the provider faces exposure for tortious interference and breach of regulatory duties.

W.R. Berkley's PC 51380 endorsement extends AI exclusions to D&O, EPL, and fiduciary coverage lines. The result: providers deploying AI call analytics to comply with FCC mitigation requirements operate with diminished insurance coverage for the errors those systems inevitably produce.

C. Independent Certification as the Insurance Solution

Specialty AI insurers — including Armilla (Lloyd's-backed, up to $25 million), Munich Re's HSB (up to $10 million E&O), and Testudo (Lloyd's Lab, mid-market) — require auditable governance

and data provenance verification as conditions of coverage. Independent certification of content provenance practices provides the auditable evidence these markets demand.

A voice service provider that certifies its AI call analytics against published standards — and cross-references traffic against verified C2PA content provenance manifests rather than relying solely on black-box heuristic detection — establishes a defensible standard of care. This certification provides the documentation necessary for insurance brokers to negotiate affirmative coverage add-backs or secure standalone AI liability policies.

V. Independent Certification as RMD Compliance Mechanism

A. The Four-Element Conformance Pattern

In our April 2026 comment to the Commission on UAS trusted supplier frameworks (GN Docket No. 26-74, ECFS 20260422570517252), we proposed a credential-based trust model consisting of four elements:

1. **Published technical criteria** — version-controlled standards against which systems can be assessed;
2. **Independent conformance assessment** by accredited bodies operating under ISO/IEC 17065 and ISO/IEC 17011;
3. **Machine-readable, cryptographically verifiable credentials** — records that a specified system, as of a specified date, meets specified criteria as confirmed by a specified assessor; and
4. **Continuous conformance** through periodic reassessment or continuous reporting, ensuring credentials reflect current system configurations.

This pattern is proven in adjacent federal contexts: FedRAMP for cloud security, FIPS 140-3 for cryptographic module validation, and Common Criteria for product certification. We propose its application to the RMD:

RMD Element	Current Model	Proposed Enhancement
Mitigation plan filing	Narrative self-certification	Independently certified against published criteria
STIR/SHAKEN implementation	Binary compliance declaration	Continuous conformance reporting with version attestation
AI detection practices	Not required	Certified AI detection with published accuracy metrics
Content provenance	Not addressed	C2PA-aligned manifest generation and retention

B. Content Provenance Authorities

The current STIR/SHAKEN governance chain — STI-GA (ATIS) → STI-PA (iconectiv) → STI-CAs → providers — is entirely focused on the signaling plane. It vets the provider's right to use telephone numbers but has no mandate or capability to vet the nature, safety, or provenance of the traffic generated by those numbers.

We propose that the Commission consider a complementary governance layer for content provenance: recognized "Content Provenance Authorities" (CPAs) that certify providers' content provenance practices. These authorities would:

- Develop published technical criteria for content provenance in telephony, aligned with C2PA specifications;
- Conduct independent assessments of providers' content provenance capabilities;
- Issue machine-readable credentials attesting to certified status; and
- Operate under multi-stakeholder governance to prevent capture by any single industry interest.

This does not require the Commission to build or operate new infrastructure. It requires recognizing independent certification as evidence of compliance within the existing RMD framework.

VI. Proof of Concept: Audio Data Certification in Operation

Box Commons currently operates a live audio data certification pipeline — the BC-Certified Audio Data Standard — across a 22-station broadcast radio network.

A. What BC-Certified Demonstrates

The pipeline captures live broadcast streams, parses playout metadata, performs audio fingerprinting for copyrighted music identification, reconciles scheduled-versus-actual content, and generates tamper-evident audit trails with SHA-256 hashing at ingestion. Consent documentation follows a structured schema covering rights holder identity, content scope, permitted uses, and revocation terms. These tools are open-source and inspectable.

B. Why Broadcast Radio Validates Telephony

Broadcast radio presents a more technically demanding certification environment than telephony:

Dimension	Broadcast Radio	Telephony
Volume	22 stations, continuous 24/7 streams	Individual calls, finite duration
Latency requirements	Real-time stream capture and processing	Call-duration processing windows
Content diversity	Music, speech, ads, station IDs, weather	Speech (human or AI-generated)
Codec variability	Multiple streaming codecs, bitrate changes	G.711, G.722, AMR (standardized)

If audio data certification works for the higher-volume, higher-complexity broadcast radio environment, it works for telephony. The certification infrastructure that processes 22 simultaneous continuous broadcast streams can process individual voice calls with substantially less computational overhead.

VII. Regulatory Cross-References

A. Alignment with FCC 24-84 (AI Voice NPRM)

The pending NPRM on AI-generated calls proposes that callers disclose AI voice usage and identify the responsible entity. Content provenance credentials provide the technical mechanism to operationalize these requirements. A PASSporT content provenance claim carrying `ai_voice: true` and a C2PA manifest referencing the AI model, consent authorization, and disclosure timestamp would provide cryptographic proof of compliance — evidence that is machine-readable, auditable, and non-repudiable.

B. State-Level Requirements

Multiple states impose requirements that content provenance can help satisfy:

- **California (PUC § 2874):** Requires a live, natural-voice announcement disclosing AI voice use before operating an automatic dialing-announcing device. Content provenance manifests could record the time and manner of this disclosure.
- **California (SB 942, AB 2013):** Requires AI developers to embed machine-readable provenance data in generated audio and disclose training data categories. C2PA manifests satisfy both requirements.
- **Washington (RCW 80.36.400):** Flatly prohibits automated AI voice solicitation. Content provenance metadata provides evidence that providers abstained from AI voices in jurisdictions with absolute prohibitions.

C. International Precedent

Article 50 of the European Union's AI Act requires providers of AI systems generating audio to mark outputs in a machine-readable format, effective August 2, 2026. C2PA is emerging as the *de facto* standard for satisfying this requirement. A telecommunications network that cannot transport or verify C2PA manifests will strip legal compliance from enterprise traffic as it crosses international network boundaries.

D. Pending Legislation

S. 2666, the Foreign Robocall Elimination Act, passed the Senate unanimously on August 3, 2026, and would require voice service providers to post surety bonds of up to $100,000 before filing RMD certifications. This legislation, if enacted, will further raise the stakes of RMD compliance. Independent certification of content provenance practices — providing auditable evidence that a provider's mitigation plan is not merely narrative but technically operational — would complement the bonding requirement by giving the Commission substantive, verifiable criteria against which to evaluate filings.

VIII. Recommendations

We respectfully recommend that the Commission:

1. **Recognize content provenance as a component of robocall mitigation.** The final rules should acknowledge that effective robocall mitigation requires both signaling-layer

authentication (STIR/SHAKEN) and media-layer provenance. RMD filings should permit — and the Commission should encourage — providers to describe their content provenance and AI detection capabilities as part of their mitigation plans.

2. **Initiate a technical proceeding on content provenance for telephony.** The Commission should seek comment on standardized content provenance metadata for voice calls, including PASSporT extension types for content provenance claims, integration with the Rich Call Data framework (RFC 9795), and alignment with C2PA content credential standards. This proceeding could be coordinated with the IETF STIR Working Group, ATIS, and the SIP Forum.

3. **Condition any content provenance recognition on governance criteria.** Any standards body whose certifications are recognized within the RMD framework should meet defined governance criteria: multi-stakeholder participation with balanced voting weight, intellectual property policies defaulting to royalty-free terms, separation of standards development from certification assessment, and publicly accessible certification registries.

4. **Coordinate with the AI voice rulemaking.** The Commission should consider how content provenance mechanisms could operationalize the disclosure and consent requirements proposed in FCC 24-84. Machine-readable content provenance metadata provides the technical infrastructure to enforce AI voice disclosure requirements that currently depend entirely on caller self-attestation.

5. **Engage international standards alignment.** As the EU AI Act's machine-readable marking requirements take effect, and as C2PA adoption expands globally, the Commission should ensure that U.S. telecommunications networks can transport and verify content provenance metadata. A proceeding on content provenance standards would position U.S. infrastructure for interoperability with emerging international requirements.

IX. Standing Offer

Box Commons offers to participate, at no cost, in any technical advisory process, pilot program, or working group the Commission convenes to develop content provenance standards for telephony. We can demonstrate our audio data certification pipeline against live broadcast radio streams, share our open-source tooling, and provide technical input on how C2PA content credentials can be integrated into the STIR/SHAKEN framework.

X. About Box Commons

Box Commons is a 501(c)(6) standards body headquartered in Wyoming. We develop technology-agnostic, outcome-based certification standards for audio 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.

Prior FCC engagement: Public comment on Unleashing American Drone Dominance (GN Docket No. 26-74, ECFS 20260422570517252, April 22, 2026), proposing credential-based trust models for autonomous systems.

Broader regulatory engagement: 21 public comments filed across 14 federal, state, and international agencies, including NIST (AI agent security, identity and authorization, AI 800-2), NHTSA (autonomous vehicle safety), NAIC (AI insurance governance), Colorado AG (ADMT training data certification), and Singapore IMDA (agentic AI governance). Acknowledged by NIST CAISI team (March 2026) and cited in IMDA Model Governance Framework v1.5 (May 2026) alongside Google, Microsoft, IBM, PwC, and AWS.

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 FCC 26-49 (91 FR 57454), FCC 24-17, FCC 24-84, and the TRACED Act; technical analysis of STIR/SHAKEN architecture (RFC 8224, RFC 8225, RFC 8226, RFC 8862, RFC 8816, RFC 9795); review of C2PA v2.2 content credential specifications; synthesis of insurance market research on ISO/Verisk exclusionary endorsements (CG 40 47/48); economic data from YouMail Robocall Index and FTC fraud reports; AI voice detection technology assessment; 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.