Authorization Infrastructure for the AI Content Era

Identity-bound provenance, machine-readable authorization state, and verifiable governance artifacts for digital content ecosystems.

Patent status	Patent Pending — USPTO Docket 85498 · Filed January 28, 2026
Version	cgp-1.0 · June 2026
Inventors	Eddie Evans · Jerry Joseph, Esq.
Organization	CreatorGuard Technology LLC · creatorguard.tech

Abstract

The proliferation of AI-generated media and the growth of large-scale content ingestion pipelines have created a fundamental governance failure: digital content moves across platforms, undergoes transformation, and enters AI training datasets without any machine-readable record of who created it, whether its use is authorized, or what rights are attached to it. Existing approaches — copyright law, watermarks, and provenance standards — address fragments of this problem but none produces a verifiable, machine-ingestible authorization state that downstream systems can act on without ambiguity.

CreatorGuard is an identity-bound provenance and authorization infrastructure layer that addresses this gap. The system accepts any digital asset, binds it cryptographically to a verified creator identity, evaluates it against a governance policy across six dimensions, and produces a W3C Verifiable Credential containing the authorization state. This credential is machine-readable, offline-verifiable, standards-aligned with C2PA, and signed with Ed25519. It is not a detection tool, a watermark, or a platform dashboard. It is governance infrastructure.

> The primary output of the CreatorGuard system is a machine-readable boolean: authorizationState.ai_training.authorized. This single signal tells an AI ingestion pipeline, a platform compliance system, or an enterprise legal team whether a specific piece of content has been authorized for training use — cryptographically bound to the creator's identity and verifiable offline.

1. The Authorization Failure

Three structural realities define the current state of digital content authorization:

- **Pixels have no memory.** Once published, a digital asset is a sequence of bytes. Any system can copy, compress, reformat, or redistribute those bytes instantly, erasing provenance. Law cannot physically prevent copying; only architecture can track it.
- **Identity is separate from the file.** A file cannot prove who made it. Signatures embedded in files are stripped by platform compression and format conversion. The only durable proof of origin requires binding the creation event to a verified identity environment at the moment of creation, external to the file itself.
- **Authorization uncertainty is becoming financially material.** Getty Images claimed $1.8 billion in damages from Stability AI over unauthorized training ingestion. The New York Times is in active litigation with OpenAI. These cases represent a category of liability that is growing monthly as AI training pipelines scale and regulatory frameworks tighten under the EU AI Act and emerging US synthetic media legislation.

2. Why Existing Approaches Are Insufficient

Copyright law	Provides legal recourse after the fact. Does not produce a machine-readable signal at ingestion time. Does not scale to billions of assets across thousands of platforms. Enforcement requires human initiation.
Digital watermarks	Metadata embedded in content files. Survives nothing — platform compression, format conversion, screenshots, and recropping all destroy embedded signals. Cannot carry authorization state. Cannot prove identity.
C2PA provenance standard	Records what happened to an asset and when. Does not record who is authorized to use it or for what purpose. C2PA manifests can be created by any party without verified identity. The AI training authorization gap is explicitly outside C2PA's current scope. CreatorGuard extends C2PA rather than competing with it.
AI content detectors	Produce probabilistic scores, not authorization states. A detection score of 87% synthetic does not tell an ingestion pipeline whether training is permitted. Detection and authorization are different problems requiring different infrastructure.

3. CreatorGuard System Architecture

The CreatorGuard platform operates as a four-layer governance infrastructure built above the C2PA provenance baseline. Each layer addresses a specific failure mode in the current content authorization landscape. The architecture is platform-agnostic, content-format-agnostic, and designed for deployment across social media platforms, content marketplaces, enterprise systems, AI training pipelines, and academic institutions.

Asset binding · Manifest stores · Cryptographic signatures · Hard binding (SHA-256 content hash) · Soft binding (decoupled manifest) · C2PA-compliant assertion schema

CreatorGuard is built above the existing C2PA standard. Organizations with existing C2PA infrastructure can add CreatorGuard as an extension layer without re-engineering their provenance pipeline. CreatorGuard consumes C2PA manifests and extends them with identity and authorization payloads.

Multi-modal biometric verification · W3C Decentralized Identifier (DID) anchoring · Creator identity profile generation · Cross-platform identity persistence · Enterprise / government / creator / organization identity sources

The identity binding layer synthesizes biometric and behavioral inputs — including facial geometry, voice spectral analysis, and behavioral fingerprinting — to generate a dynamic trust score and bind the content creation event to a verified, persistent creator identity. GAN detection models identify AI-generated content at this layer. The output is a W3C DID that survives platform migration.

AI training authorization · Commercial reuse permissions · Attribution requirements · Jurisdictional framework resolution · Synthetic media disclosure · Derivative works governance · Policy Evaluation Completeness score (0.0–1.0)

Six governance dimensions are evaluated against creator declarations at issuance time. Each dimension resolves to a machine-readable state. The primary signal — authorizationState.ai_training.authorized — is binary. Authorization state is not a detection probability. It is a declared, cryptographically-bound governance record.

Smart contract enforcement · Field-of-use licensing management · Real-time content tracking · Automated takedown workflows · Audit-ready output records · RESTful API for platform integration · QR code and steganographic embedding for offline verification

The application layer is where authorization state is enforced and licensing revenue is generated. Smart contracts define usage conditions including license durations, geographical constraints, derivative work permissions, and revenue-sharing arrangements. Enforcement is automated and does not require human initiation.

4. Credential Output and Verification

The CreatorGuard system produces a W3C Verifiable Credential in JSON-LD format. This credential is the primary institutional artifact — the object that AI pipelines, compliance tools, legal teams, and audit systems consume. It contains the full authorization state, the content fingerprint, the creator identity binding, and a cryptographic proof.

Credential structure

Field	Content	Purpose
@context	W3C VC + cgp-1.0 schema	Standards interoperability
type	CreatorGuardAuthorizationCred.	System identification
credentialSubject.id	W3C DID (creator identity)	Identity anchor
contentFingerprint.sha256	SHA-256 of raw content bytes	Tamper evidence
provenanceContinuityId	UUID per verification event	Audit trail reference
ai_training.authorized	true / false	Primary governance signal
commercial_reuse.permitted	true / false	Licensing gate
attribution.required	true / false	Rights enforcement
jurisdictional.frameworks	DMCA / GDPR / EU AI Act / etc.	Compliance mapping
synthetic_media.flagged	true / false	Disclosure obligation
governanceConfidenceScore	0.0 – 1.0	Policy completeness
proof.type	Ed25519Signature2020	Cryptographic proof
proof.proofValue	Base64 Ed25519 signature	Tamper detection

Offline verification process

Credential verification requires no network call to CreatorGuard infrastructure. Any party with the credential performs three checks:

- **Step 1 — Content fingerprint match.** Recompute SHA-256 of the content file raw bytes. Compare to credentialSubject.contentFingerprint.sha256. Mismatch means the credential does not correspond to this content.
- **Step 2 — Cryptographic signature.** Verify the Ed25519 signature in the proof block using the public key embedded in the credential. A valid signature confirms the credential has not been modified since issuance.
- **Step 3 — Authorization state.** Read ai_training.authorized. False = NOT AUTHORIZED. Hard stop for training ingestion. Absence of a credential carries the same weight as NOT AUTHORIZED.

Schema endpoint

The JSON-LD context document is served at creatorguard.tech/contexts/provenance/v1.json with Content-Type: application/ld+json. Any JSON-LD processor can dereference this URL to validate and interpret the credential schema. CORS headers allow cross-origin access by browser-based verification tools.

5. Deployment Architecture

The system supports three deployment models depending on the operator's infrastructure requirements and regulatory environment:

Centralized SaaS	REST API integration. Platform operators call the CreatorGuard API to register content and receive signed credentials. Suitable for content platforms, stock media companies, and enterprises requiring rapid deployment.
Blockchain-native	Full decentralization via distributed ledger. Proof-of-creation records, timestamped ownership events, and smart contract-based authorization records recorded on-chain. Suitable for operators requiring full transparency and independence from any single service provider.
Hybrid	Sensitive identity computations occur off-chain within private environments. Critical proof elements — content fingerprints, authorization states, ownership timestamps — recorded on-chain. Suitable for regulated industries and government deployments with data sovereignty requirements.

6. Strategic Position

CreatorGuard occupies a position in the content authorization stack that no existing standard currently fills. The analogy is instructive: PKI (Public Key Infrastructure) became mandatory across every browser, every enterprise, and every regulated industry not because it was a good idea but because the liability of operating without it became untenable. The same dynamic is forming around AI training authorization.

The regulatory pressure tunnel is accelerating. EU AI Act Article 53 creates compliance obligations around training data provenance. FTC synthetic media rules are expanding. Active copyright litigation is establishing case law that will make demonstrable authorization state a legal requirement rather than a best practice. CreatorGuard is positioned as the infrastructure layer that is already built when that mandate arrives.

Licensing targets

Category	Organizations	Licensing track
Stock media	Getty Images, Shutterstock, AP	Issuance
AI platforms	OpenAI, Google DeepMind, Stability AI	Consumption
Creative platforms	Adobe, Canva, Figma	Both
Social platforms	YouTube, Meta, TikTok, X	Both
Government / DoD	Federal agencies — SAM.gov eligible	Issuance
Standards bodies	C2PA, W3C CCG, NIST AI programs	Reference

7. Scope of Verification

CreatorGuard verifies declared authorization state and cryptographic continuity. It does not independently adjudicate copyright ownership, resolve title disputes, or constitute a legal opinion on the validity of the creator's claimed rights. Authorization states reflect creator declarations at time of issuance, bound to the content fingerprint recorded in the credential. Recipients should consult qualified legal counsel for jurisdiction-specific compliance obligations.

Technical Evaluation Access

Full integration documentation, live credential samples, verification library, and detailed use case specifications are available under NDA to qualified institutional parties.

Submit request: **creatorguard.tech/access**
Direct contact: **contact@creatorguard.tech**
Patent counsel: Jerry Joseph, Esq. · (202) 570-4009

implementation details, source code, claim charts, or unfiled embodiments. All technical claims reference the publicly filed provisional application Docket 85498.