EVIWRITE FORMAL REPORT

Content Credentials and the Provenance Breakage Problem

Broken metadata chains, stripped credentials, contested labels, verification dependency, and the evidence records needed before provenance claims are challenged.

A proof-landscape report on broken metadata chains, stripped credentials, contested labels, verification dependency and the evidence records needed before provenance claims are challenged.

REPORT NUMBER	CCDAR-2026-01
VERSION	1.2
PUBLICATION DATE	2026-05-27
STATUS	Final draft
DOCUMENT CLASS	Public evidential trend report
REFERENCE	CCDAR-2026-01

REPORT POSITION

Proof landscape, not threat landscape.

EviWrite reports identify where public digital claims become evidentially weak, what stronger records would have required, and how similar claims should be evidenced before pressure arrives.

Document control

Source file	content-credentials-dispute-and-adoption-report-2026.md
Status	Final draft
Version	1.2
Period	Evidence horizon to 27 May 2026 (2023-10-01 to 2026-05-27)
Prepared by	EviWrite
Report hash	not-issued
PDF hash	not-issued
Receipt	not-issued

Exclusions

- Forensic determination of whether any specific media item is authentic.
- Legal advice or determinations of liability, admissibility, infringement or regulatory compliance.
- Global incidence measurement of all provenance failures.
- Private platform telemetry, private implementation test results or non-public forensic logs.

Proof limits

- That any named party committed wrongdoing.
- That any individual case would fail in court or before a regulator.
- That private evidence does not exist.
- That a cited source is complete, final, uncontested, or legally determinative.
- That EviWrite has independently audited the underlying private systems, logs, files, datasets or forensic records.
- That selected public signal counts represent global incidence.
- That forecast scores represent probability, incidence, legal outcome or market size.
- That selected international breakdowns represent global prevalence.

Content Credentials and the Provenance Breakage Problem

Broken metadata chains, stripped credentials, contested labels, verification dependency, and the evidence records needed before provenance claims are challenged.

Executive summary

Content Credentials are becoming the default public language of digital provenance. The mistake is to treat that as the end of the evidence problem. It is the beginning of a harder one.

This is not a threat-landscape report, platform trust ranking, legal update, or technical audit of C2PA. It is a proof-landscape report: it examines where public provenance claims become evidentially weak, what records would make them stronger, and how organisations should prepare proof before dispute, incident, publication, regulatory scrutiny or public challenge arrives.

The selected public signals reviewed here point to one central pattern: provenance fails most often after the credential is created. The breakage appears in export, upload, re-encoding, platform labelling, verification portals, user interpretation and later preservation. A credential can be technically valid and still be weak public evidence if the claim placed on it is broader than the record supports.

The next pressure point is not simply adoption. It is reconciliation: C2PA metadata, watermarks, platform labels, verification portals, camera-origin credentials and human explanations will need to agree, or at least preserve their disagreement clearly enough for later review.

Source basis for this report

This report combines technical standards, official guidance, platform notices, provider verification materials, camera provenance adoption pages, specialist research and selected journalism. Repeatable sources provide continuity; selected public records provide signal-level evidence. EviWrite classification is applied to identify proof weakness. Counts and charts are not global incidence.

Source group	Used for	Not used for
C2PA and Content Credentials standards	Technical capability, manifest scope, provenance concepts	Proving truth, consent, legality, ownership or completeness
Platform and provider notices	Public adoption, label logic, stated limits, verification-portal boundaries	Independent proof of real-world label accuracy
Public authority guidance	Governance pressure and operational evidence requirements	Determining compliance by any named organisation
Specialist research	Technical dispute signals and forward evidence pressure	Final peer-reviewed consensus or production incidence
Journalism	Selected public examples of platform-display and public-trust friction	Global failure rates or complete platform measurement

Source boundary: The sources support selected public evidence-pressure signals. They do not prove global incidence, hidden private evidence, legal liability, regulatory breach, or whether any individual media item is authentic.

The EviWrite evidencing lens

Every signal in this report is assessed through six questions:

Lens	Question
Source	Where did the record, content, dataset, action, or decision come from?
Timing	When did the relevant event, creation, access, disclosure, decision, or knowledge occur?
Control	Who controlled the file, system, account, platform, dataset, log, or workflow?
Sequence	What happened before, during, and after the claim being made?
Verification	Can the claim be checked without simply trusting the claimant or the platform?
Limits	What does the record not prove?

The EviWrite Evidence Failure Stack

Layer	Failure question	Common weakness in Content Credentials disputes
Event	Did the thing happen?	Creation, edit, upload, label or verification event is asserted but not reconstructed.
Record	Was it recorded at the time?	The strongest evidence is generated only after public challenge.
Context	Does the record explain meaning, authority and limits?	A label says “AI” or “verified” without explaining the trigger or boundary.
Custody	Can control and movement be shown?	The file passed through tools, platforms or dashboards that changed what survived.
Trust Boundary	Did the record leave the system or party that may later be questioned?	The proof remains inside a provider portal, platform UI or original file only.
Verification	Can the claim be checked without belief?	Public claims depend on private implementation, hidden logs or unsupported assertions.
Permanence	Will the proof survive time and challenge?	Metadata, labels, links, portal results and platform explanations may disappear or change.

The EviWrite Provenance Breakage Chain

Content Credentials do not fail mainly at the point of generation. They fail when a credentialed asset passes through ordinary systems: editing tools, export settings, upload pipelines, platform labels, reposting, screenshots, verification portals and human interpretation.

Layer	Breakage mode	Evidence question
Capture	No credential is created at source, or the source tool is not credential-capable	Can the original capture event be linked to the disputed asset?
Edit	The edit tool changes the asset without preserving or extending the provenance chain	Can each material transformation be reconstructed?
Export	Format conversion, compression or save-as workflows remove or rewrite metadata	Can the exported file be linked back to the source file and export settings?
Upload	The platform strips, rewrites, hides or substitutes provenance data	What did the platform receive, preserve, alter, display and suppress?
Distribution	Screenshots, reposts, previews and recompression detach visible content from the credentialed file	Is the disputed copy the original file, a derivative, or only a visual reproduction?
Label	A simplified UI label replaces a technical record	What rule, signal or policy caused the displayed label?
Verification	The result depends on a current portal, trust list, certificate chain or file state	Can the result be preserved and replayed independently later?
Interpretation	Users treat provenance as proof of truth, consent, legality, ownership or completeness	What exact claim does the record support, and what does it not support?

Evidence basis: C2PA technical specifications and threat model material, Content Credentials verification guidance, OpenAI verification material, NSA content provenance guidance, platform labelling statements and public platform distribution testing.

Source boundary: These sources support the technical and platform-dependency premises. The breakage chain is EviWrite classification, not an official C2PA, platform, regulator or court finding.

Content Credentials dispute classes

The next evidence problem is not simple adoption. Adoption creates new disputes over absence, meaning, authority and replayability.

Dispute class	Evidential question
Absence dispute	No credential is present. Was it never created, stripped, unsupported, hidden, or lost during distribution?
Label dispute	A platform label is visible. What signal, policy rule, tool marker or provider assertion triggered it?
Chain dispute	Which step broke provenance: capture, edit, export, upload, distribution, labelling or verification?
Authority dispute	Who signed the credential, what trust list recognised them, and why should that signer be relied on?
Portal dispute	What did the verification service show at the time, and can that result be preserved or independently replayed?
Conflict dispute	Metadata, watermark, visual context, platform label and third-party analysis disagree. Which record governs the claim?
Meaning dispute	Does the credential indicate camera capture, AI creation, AI editing, software use, export history, or merely tool involvement?

Evidence basis: C2PA specification and threat model material, platform AI label statements, Content Credentials verification guidance, OpenAI verification material and public reporting on platform handling of AI provenance signals.

Source boundary: The dispute taxonomy is EviWrite interpretation based on public technical and platform sources. It does not determine the outcome of any specific dispute.

False confidence patterns

False confidence	Why it is weak	Stronger posture
“No credential means not AI-generated”	Absence may mean stripping, legacy generation, unsupported tool use, screenshotting, recompression, platform hiding or unsupported file format	Preserve the original file, platform journey, tool chain and absence explanation
“Credential present means true”	A credential can support provenance, not factual accuracy, lawful use, consent, ownership or completeness	Bind credential evidence to rights, consent, source, edit and claim-boundary records
“The label tells users enough”	Labels compress technical, legal and evidential distinctions into vague UI language	Preserve the triggering signal, policy basis, label text, timestamp and platform explanation
“The portal verified it”	The verification event itself may depend on current provider infrastructure, trust lists and file state	Preserve verifier result, checked-file hash, timestamp, trust decision and warning state
“Watermark plus metadata solves it”	Multiple signals can conflict and may survive different distribution routes	Maintain a reconciliation record explaining which signal governs which claim
“The platform has the logs”	The affected party may not control, export, preserve or independently verify the platform’s strongest record	Require exportable evidence, retention commitments and independent timestamped captures
“The image looks authentic”	Visual plausibility is not provenance, authorship, consent or custody	Require source, timing, tool, custody and verification records
“We disclosed AI use”	Disclosure is not the same as preserved provenance evidence	Keep the underlying manifest, workflow record, review note and final-publication evidence

Evidence basis: C2PA technical material, platform labelling statements, Content Credentials verification guidance, OpenAI verification material, NSA guidance and public platform distribution testing.

Source boundary: These sources support the existence of technical, platform and verification limits. They do not prove that any specific label, credential or verification result is wrong.

Five findings

1. The weakest point is no longer generation. It is passage through the distribution chain.

Content Credentials can be created correctly and still fail as public evidence once the file moves through export, upload, re-encoding, download, editing, screenshotting or reupload. The practical proof question is therefore not “does the standard work?” but “does the evidence survive the route the public claim took?”

Evidence basis: C2PA Technical Specification 2.4; OpenAI image verification tool; Meta AI labelling approach; TikTok watermark and label update; NSA content credentials guidance.

Source boundary: These sources support visible platform and metadata fragility signals. They do not provide global failure rates.

2. A label is not a credential, and a credential is not a verdict.

The user-facing label is a compressed public statement. It can be triggered by generation, editing, metadata, watermarking, self-disclosure or platform inference. Without a trigger record, the label becomes an ambiguous claim rather than a useful piece of evidence.

Evidence basis: Meta AI Info label update; Content Credentials public site; OpenAI image verification tool.

Source boundary: These sources support label and verification-limit analysis. They do not measure user comprehension globally.

3. Verification portals create evidence, but they also create dependency.

A portal result is useful only if it is preserved with the exact file hash, timestamp, verifier version and limitation text. Otherwise the portal becomes another captive record: useful today, potentially unreproducible tomorrow.

Evidence basis: OpenAI image verification tool; OpenAI provenance update; Content Credentials public site.

Source boundary: These sources support the existence and stated limits of verification portals. They do not prove future portal availability or all-provider coverage.

4. Layering metadata, watermarking and fingerprints improves durability but creates reconciliation duties.

The industry is moving toward layered provenance because metadata alone is fragile. That is correct. But layered systems create a new evidence burden: what happens when C2PA metadata, watermark detection, platform labels and portal results do not agree?

Evidence basis: OpenAI provenance update; TikTok watermark and label update; specialist research on provenance-watermark contradiction.

Source boundary: The provider and platform sources support the layered-provenance trend. The contradiction analysis is specialist preprint evidence and should be tracked as the research matures.

5. Adoption pressure is becoming compliance pressure. Evidence of marking will matter as much as marking itself.

As legal and policy obligations around AI-generated content marking mature, the evidential question shifts from “did we label it?” to “can we prove what was labelled, how, when, where, for whom, and whether the signal survived distribution?”

Evidence basis: EU AI Act Article 50 material; European Commission AI-generated content code material; NIST Generative AI Profile; specialist watermarking adoption research.

Source boundary: These sources support regulatory and governance pressure. They do not determine compliance by any named party.

Evidence-signal scorecard

The report classifies 26 selected public signals. The count is a research register for this report, not global incidence.

Dataset basis: Derived from the selected evidence signals classified in this report. See the Source-to-Claim Map and source register.

Chart boundary: EviWrite classification of selected public signals. Not global incidence, telemetry, legal finding, market sizing or probability.

Visual chart summary

The charts declared in this report render the selected evidence-signal register by primary failure type, sector, claim category and region. A separate line chart renders qualitative forward evidence-pressure trajectories.

Dataset basis: Derived from selected public evidence signals and forecast signals in the YAML register.

Chart boundary: Selected-signal charts are not global incidence. Forecast charts are qualitative pressure scores, not legal predictions, probability forecasts or market sizing.

Visual chart summary

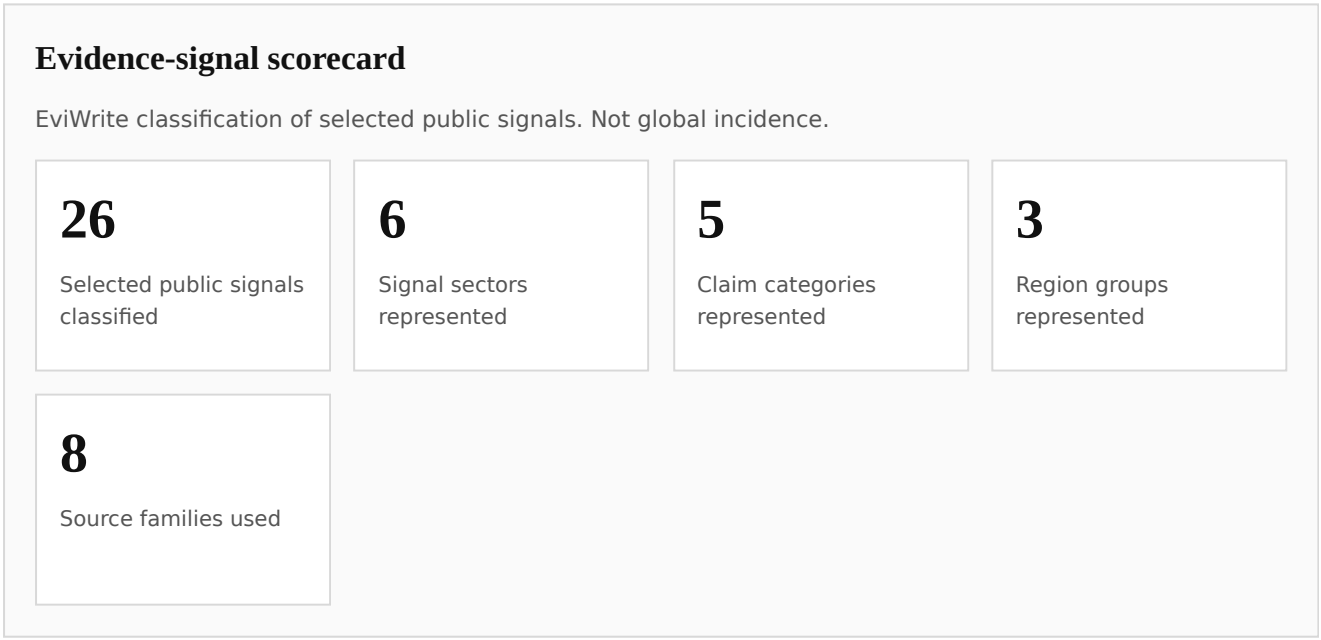

Selected signals by primary evidence weakness

EviWrite classification of selected public signals. Not global incidence.

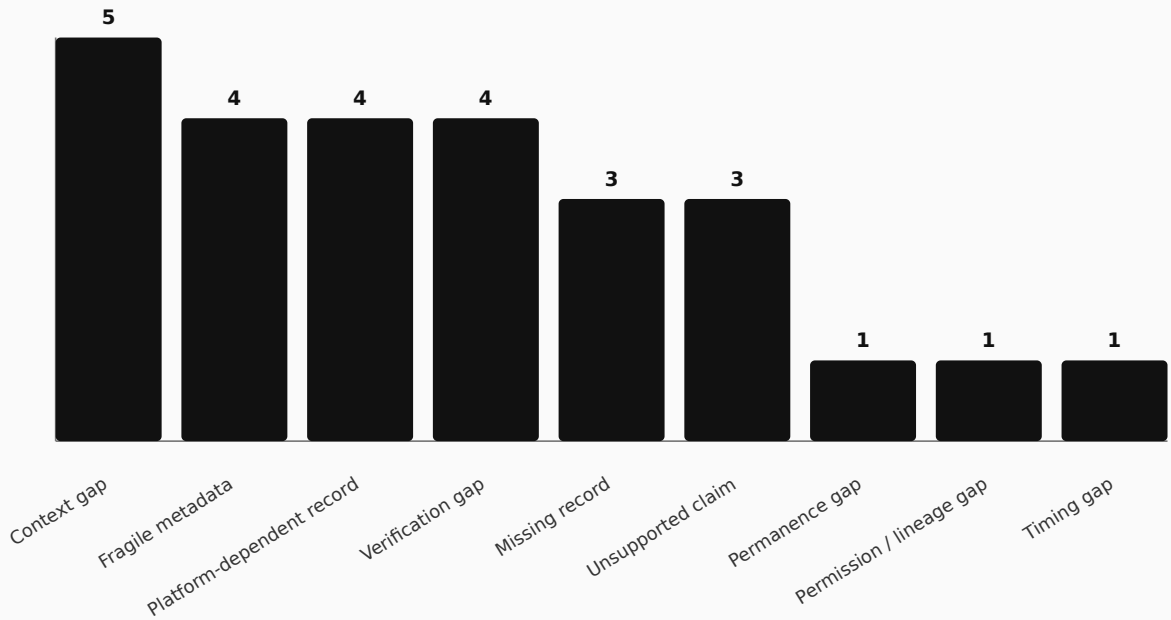

Selected signals by consolidated sector

EviWrite classification of selected public signals. Not global incidence.

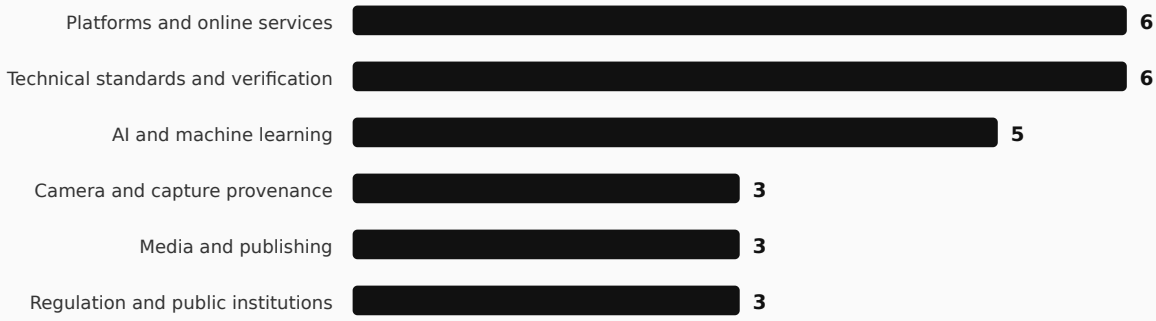

Selected signals by claim category

EviWrite classification of selected public signals by claim category. Not global incidence.

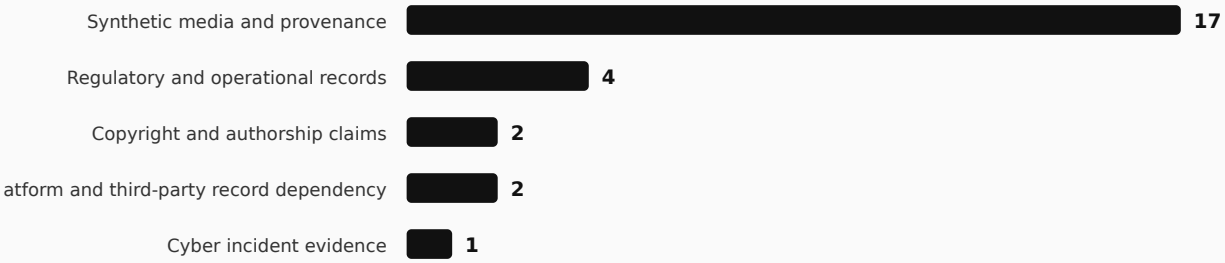

Selected signals by region

EviWrite classification of selected public signals by region. Not global incidence.

Forward evidence-pressure trajectories

Qualitative EviWrite forecast score. Not measured incidence, probability, market sizing, or legal prediction.

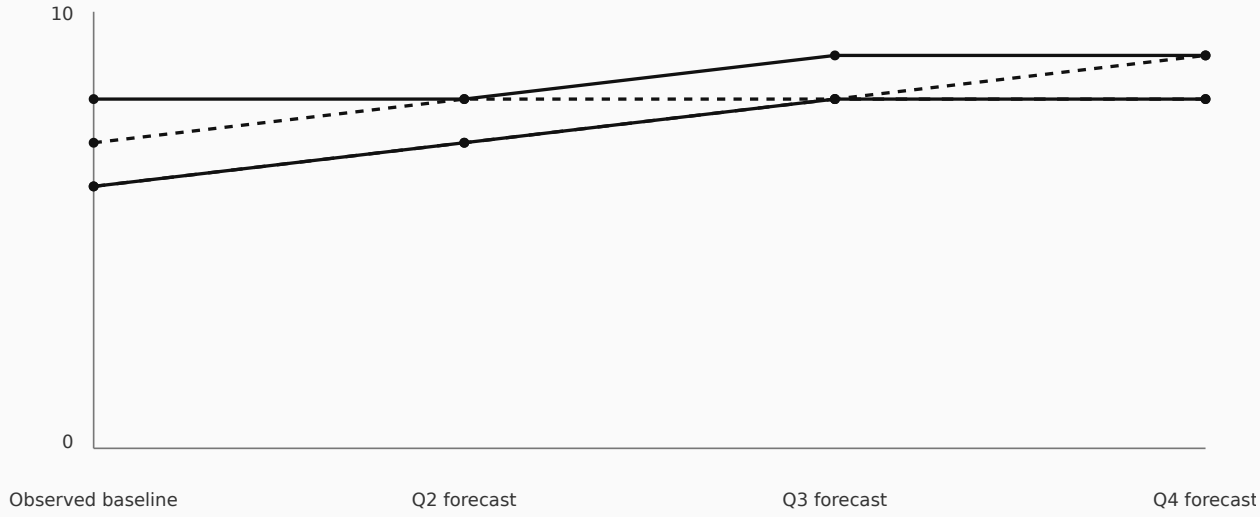

1 Metadata-only reliance declines
2 Label-semantics disputes rise
3 Verification portals become dispute records
4 Cross-layer signal conflict becomes a verification issue

Evidence failure types by primary weakness

The selected signal set concentrates around verification gaps, context gaps, fragile metadata and platform-dependent records. That is the core insight. The public argument is often about whether media is “real” or “AI”. The evidential failure is usually narrower and more operational: the source, label reason, custody path, verification result or survival state cannot be shown cleanly.

Claim-category breakdown

Synthetic media and provenance dominates the selected register because Content Credentials sit at the boundary between technical metadata and public trust. Platform-record dependency is the second pressure area: the strongest proof frequently sits inside distribution platforms, provider portals or tool ecosystems.

Sector pressure ranking

Platform and online services carry disproportionate evidence pressure because they decide whether credentials are retained, stripped, hidden, translated into labels or exposed to viewers. AI providers create the initial signal. Camera and creative-tool providers can strengthen origin. But the platform layer often determines what the public actually sees.

Regional selected-signal breakdown

North America and international/cross-border sources dominate this selected register because many primary provider, platform and standards sources are US or global. European Union sources are critical for forward pressure because transparency obligations and marking guidance are moving provenance into compliance evidence.

Boundary: This is source-location and jurisdictional relevance within a selected register. It is not a measure of global prevalence.

Forecast / forward pressure signals

Four evidence-pressure trajectories are most likely to matter through the rest of 2026:

Forward pressure signal	Direction	Evidence question
Metadata-only reliance declines	Increase	Can provenance still be verified after export, upload, edit, screenshot or reupload?
Label-semantics disputes rise	Increase	What exactly caused the label and what claim can the viewer reasonably infer?
Verification portals become dispute records	Increase	Can a verification result be reproduced, timestamped and linked to the exact uploaded file?
Cross-layer signal conflict becomes a verification issue	Increase	What happens when C2PA metadata, watermarks, platform labels and portal results disagree?

Source basis: OpenAI provenance and verification materials; TikTok watermark update; Meta label update; specialist research on provenance-watermark conflicts.

Forecast boundary: These are qualitative EviWrite pressure scores. They are not incidence predictions, probability forecasts, legal predictions or market sizing.

Case studies

Case study 1: Meta’s AI label wording change

Meta’s move from “Made with AI” toward “AI info” is not just a communications adjustment. It exposes a proof weakness: industry markers can trigger a public label whose meaning is broader, narrower or different from what users infer.

Source basis: Meta AI Info label update; Meta AI labelling approach.

Evidence boundary: This case study does not determine whether any specific Meta label was wrong. It identifies the proof weakness created by compressed labels.

What stronger evidence would have required:

Weak record	Stronger evidence
A visible AI label only	Label trigger reason, source signal, AI role category, edit/generation distinction and user-facing explanation
Platform policy statement only	Applied-label audit records showing actual treatment across representative content states

Case study 2: OpenAI verification portal

The OpenAI verification portal is valuable precisely because it states limits. It can detect supported OpenAI-associated signals in supported image formats. It does not decide whether an image is accurate, misleading, fully unmodified, generated by another provider, or human-made when no signal is found.

Source basis: OpenAI provenance update; OpenAI image verification tool.

Evidence boundary: This case study does not independently validate the tool or generalise it to all AI systems.

What stronger evidence would have required:

Weak record	Stronger evidence
Portal screen result only	Original file hash, upload timestamp, verifier version, captured result, stated limitations and preservation record
No signal found	Explicit inconclusive status plus corroboration from source file, creation workflow and distribution chain

Case study 3: Platform distribution tests

Selected public reporting indicates that credentialed media may still fail at viewer display. This is the distribution-layer problem. A credential invisible to ordinary viewers is not useless, but it is weak public evidence unless independently preserved and surfaced when needed.

Source basis: Washington Post platform C2PA test; The Verge Sora C2PA label analysis.

Evidence boundary: Journalistic tests are selected public signals, not global platform incidence.

What stronger evidence would have required:

Weak record	Stronger evidence
Public post only	Pre-upload original, embedded credential validation, upload receipt, platform processing state and public label capture
Buried label or absent label	UI visibility record showing location, wording, explanation and availability to ordinary viewers

Case study 4: Camera-origin credentials

Camera-origin credentials are a stronger foundation than post-hoc labelling for capture claims. But they do not make the scene true, lawful, consented, complete or rights-cleared.

Source basis: Leica M11-P Content Credentials; CAI Leica launch note.

Evidence boundary: This case study does not determine whether any specific camera-captured image is truthful or rights-cleared.

What stronger evidence would have required:

Weak record	Stronger evidence
Signed capture metadata only	Original in-camera file, device/certificate record, photographer identity basis, assignment/consent record and edit/export chain
Camera-origin claim detached from publication path	Publication-chain log showing where the file moved, who handled it and which derivative was published

---

Sector, claim-category, and failure-type table

Sector	Claim category	Visible pressure	Primary evidence weakness	Stronger record
Platforms and online services	Synthetic media and provenance	Labels, reupload, stripping, viewer display	Platform-dependent record	Pre/post upload state, label trigger, public UI capture, export/download record
AI and machine learning	Synthetic media and provenance	C2PA, watermarking, verification portals	Fragile metadata	Model/version, output marking, watermark status, C2PA validation, portal result
Camera provenance	Synthetic media and provenance	Capture-time credential adoption	Timing gap	In-camera original, certificate chain, assignment, consent, edit manifest
Legal and regulatory disputes	Regulatory and operational records	AI transparency obligations	Missing record	Marking log, exception basis, disclosure state, preservation record
Media and journalism	Platform record dependency	Public trust in authentic media	Context gap	Source chain, editorial review, publication chain, public correction trail
Creative tools	Copyright and authorship claims	Attribution and AI-preference fields	Permission / lineage gap	Work history, rights record, licence/assignment, credential boundary note

Minimum Evidence Records

Area	Minimum Evidence Record	Why it matters
Content Credential creation	Original file hash, manifest export, signer/tool identity, claim boundary, creation timestamp	Separates provenance from unsupported truth claims
Platform upload	Pre-upload file, upload timestamp, platform processing record, post-upload label state, public URL capture	Shows whether the evidence survived distribution
AI output marking	Model/version, output ID, C2PA status, watermark status, visible label, exception basis	Turns synthetic-content marking into demonstrable evidence
Verification portal use	Uploaded file hash, portal result, verifier version, timestamp, stated limits, screenshot/PDF capture	Prevents a transient portal result becoming unverifiable later
Camera-origin media	In-camera original, certificate chain, device model, capture metadata, photographer/assignment record	Strengthens origin and timing without overclaiming scene truth
Rights and attribution	Creator identity basis, licence, assignment, consent, publication authority, credential display settings	Prevents attribution fields being mistaken for ownership proof
Cross-layer conflict	C2PA validation, watermark result, platform label, portal result, reconciliation decision	Handles disagreement between independent signal layers
Archive and permanence	Immutable storage record, hash, credential copy, verification log, retention policy, correction trail	Keeps proof alive after metadata, links or dashboards change

Minimum dispute packet for contested Content Credentials claims

When a Content Credentials claim is contested, the practical issue is not whether one record exists. The issue is whether the claim can be reconstructed across source, file, tool, platform, label and verification layers.

Record	Required detail
Original asset	Hash, format, size, source location, capture/export timestamp and preservation timestamp
Credential manifest	Full manifest, signer, certificate chain, validation result, trust-list context and external-manifest reference where applicable
Tool chain	Capture device, editing software, generation tool, model/tool version, export settings and material transformation log
Platform journey	Upload account, upload time, upload URL, platform processing state, download/repost path and metadata rewriting evidence
Label evidence	Screenshot, page capture, label text, platform policy version, timestamp and explanation of what triggered the label where available
Verification evidence	Verifier used, result shown, timestamp, checked-file hash, trust decision, warning state and preserved result page or report
Preservation record	Independent timestamp, immutable receipt, custody note, holder identity or role and record-retention location
Claim boundary	Plain statement of what the credential supports and what it does not prove: truth, consent, legality, ownership, completeness or unchanged distribution

Evidence basis: C2PA specifications, Content Credentials verification guidance, OpenAI verification guidance, NSA content provenance guidance and platform labelling material.

Evidence boundary: This packet improves dispute readiness. It does not guarantee admissibility, regulatory compliance, legal success or factual certainty.

Recommendations by audience

Audience	Recommended actions
Creators	Preserve originals, hashes, credentials and publication records before upload. Keep rights and consent records separate from attribution metadata.
Businesses	Build a provenance evidence policy covering creation, export, upload, verification and retention. Capture pre-upload and post-upload states for high-risk public material.
Legal	Ask what the credential proves, what it omits and whether the exact file and verification state were preserved. Do not rely on labels without trigger and custody records.
Providers	Expose label trigger reasons and verification limits in machine-readable and human-readable form. Publish conformance scope and known failure modes.
AI teams	Record model/version, marking method, watermark status, C2PA generation, transformation and verification tests. Build cross-layer audit.
Public institutions	Treat Content Credentials as one control inside a broader multimedia integrity record. Require preservation outside platform dashboards and media files.
Education and research	Teach credential interpretation separately from truth evaluation, authorship, citation and consent. Preserve source data and publication-chain records.
Media and publishing	Use capture-time credentials where possible, but pair them with editorial, assignment, rights and publication-chain records. Audit CMS, wire, social and archive workflows.

Methodology and limitations

This report uses selected public sources and EviWrite classification. It is not a global incidence report, legal opinion, forensic audit, regulatory finding or cyber telemetry product.

Signals were selected where a public source revealed an evidential question around source, timing, control, sequence, verification or proof limits. Each selected signal was classified by sector, region, claim category, failure type, severity and confidence. Charts derive from that selected register.

Private evidence may exist. Absence of public evidence is not proof of absence. A cited source may be incomplete, contested, updated or superseded. Specialist preprints are used as technical dispute signals, not final consensus. Journalism is used as selected public-signal support, not as global measurement.

Deep source appendix

This appendix is intentionally detailed because provenance disputes often turn on source type, support boundary and verification dependency. If the PDF is later shortened for executive circulation, the source register should be preserved or issued as a separately controlled evidential appendix rather than removed from the record.

Source methodology

Sources were selected for evidential relevance, not volume. The priority was to identify where a public claim depends on a record that may be missing, fragile, captive, ambiguous, late, stripped, misunderstood or overclaimed.

Core repeatable source spine

The repeatable spine for this report family includes C2PA and Content Credentials standards, platform labelling policies, public authority guidance, provider verification materials, specialist provenance research and selected journalism.

Report-family source module

This report uses the synthetic media and provenance module: C2PA specification material, Content Credentials adoption material, platform labelling policies, verification-portal behaviour, public authority guidance and camera provenance adoption pages.

Source quality tiers

Primary sources ground standards, public policies, provider claims and official guidance. Specialist sources identify technical dispute signals. Journalism identifies selected public failures or friction points. None of these source types alone proves global incidence.

Source group synthesis

The source register shows convergence around a clear pattern: the ecosystem is building provenance signals faster than it is building durable, user-comprehensible and dispute-ready evidence chains around those signals.

Source-to-claim map

Claim	Support level	Source basis	Boundary
Content Credentials strengthen provenance but do not prove truth, consent, legality, ownership or completeness.	Strong	C2PA specification, Content Credentials site, OpenAI verification limits	Does not assess any individual file.
Metadata stripping and platform processing create a major evidence-breakage layer.	Strong	TikTok watermark update, NSA guidance, OpenAI verification limits, selected journalism	Not global incidence.
User-facing labels can be technically triggered but semantically misunderstood.	Strong	Meta label update, Content Credentials site, selected journalism	Not a user comprehension study.
Verification portals are useful but must be preserved as evidence artefacts.	Strong	OpenAI verification tool and provenance update	Not future availability proof.
Layered provenance creates reconciliation duties.	Medium-high	OpenAI and TikTok layered-signal updates; specialist research	Specialist contradiction analysis is preprint evidence.
Regulatory pressure is moving provenance toward demonstrable evidence of marking and disclosure.	Medium-high	EU AI Act material, European Commission code material, NIST, watermarking research	Not compliance advice.

Assurance and review status

This report is a public evidential trend report. It is not an audit, legal opinion, forensic report, regulatory finding, assurance engagement, cyber telemetry report or global incidence report.

Source register, source-to-claim map, selected signal register, forecast signal register and chart limitations have been prepared. External review has not been commissioned for this draft version.

Version and audit record

Field	Value
Report number	CCDAR-2026-01
Version	1.0
Period covered	Evidence horizon to 27 May 2026
Prepared by	EviWrite
Status	Draft for PDF compilation
Report hash	Not issued
PDF hash	Not issued
Source register hash	Not issued
Machine-readable register hash	Not applicable
Receipt	Not issued for this version

Glossary

Term	Meaning
Content Credential	A provenance record based on the C2PA standard that can describe origin, edits, assertions and signatures associated with a digital asset.
C2PA	Coalition for Content Provenance and Authenticity, the body behind an open technical standard for media provenance.
Metadata stripping	Removal of embedded metadata during export, upload, editing, re-encoding, screenshotting or platform processing.
Verification portal	A tool that checks uploaded media for supported provenance signals and returns a result with stated limits.