EVIWRITE FORMAL REPORT

Synthetic Media and Provenance Adoption Report

C2PA adoption, SynthID and watermarking, platform labelling, camera provenance, verification portals and user-facing signals

Selected public evidence signals show rapid movement in C2PA, SynthID, platform labelling, camera provenance and verification portals, but the core evidential weakness is now survival, interpretation and independent verification after distribution.

REPORT NUMBER	SMPR-2026-01
VERSION	1.0
PUBLICATION DATE	2026-05-27
STATUS	Draft for PDF compilation
DOCUMENT CLASS	Public evidential trend report
REFERENCE	SMPR-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	synthetic-media-and-provenance-adoption-report-2026.md
Status	Draft for PDF compilation
Version	1.0
Period	Evidence horizon to 27 May 2026 (2024-01-01 to 2026-05-27)
Prepared by	EviWrite
Report hash	not-issued
PDF hash	not-issued
Receipt	not-issued

Exclusions

- Forensic determination of any individual media item
- Global incidence of synthetic media
- Complete market sizing
- Legal advice or compliance assessment
- Internal platform telemetry
- Private provider implementation evidence not publicly available

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.

Synthetic Media and Provenance Adoption Report

Executive summary

Synthetic-media provenance has moved past the stage where the serious question is whether C2PA, Content Credentials, SynthID, platform labels, camera provenance or verification portals exist. They now exist in enough public forms to matter. The harder question is whether they survive ordinary distribution and whether the person seeing the content can interpret the signal correctly.

This is not a threat-landscape report, legal update, telemetry report or market-sizing study. It is a proof-landscape report: it examines where public digital claims become evidentially weak and what records would make those claims more defensible.

The selected public signals reviewed here show a structural adoption gap. Provenance is being added at generation, capture and provider layers, while the practical evidence chain often weakens at upload, recompression, platform display, user interface, appeal and later verification. The next serious disputes will not ask only whether content was labelled. They will ask who applied the label, when it was applied, whether it survived, whether it was visible, whether it was correct, and what the signal did not prove.

Source basis for this report

This report combines repeatable international source families with selected public records. Repeatable sources provide continuity; selected sources provide evidence for individual signals. EviWrite classification is applied to identify proof weakness, but selected signal counts are not global incidence unless the underlying dataset supports that claim.

Source group	Used for	Not used for
Technical standards and specifications	C2PA trust model, provenance assertions, signed manifest boundaries	Proving truth, consent, legality, ownership or completeness
Provider and manufacturer notices	Adoption signals, product direction, public verification surfaces	Independent audit of implementation quality
Platform policies and label updates	User-facing disclosure behaviour and platform dependency	Complete enforcement measurement or label accuracy
Regulator and public authority sources	Direction of legal and operational pressure	Legal advice or compliance findings
Technical research and selected investigations	Failure modes, robustness questions, distribution-chain stress	Global incidence or universal platform behaviour

Source basis: C2PA Technical Specification; NIST Synthetic Content Guidance; Content Credentials Adoption Site; OpenAI Provenance Update; Google SynthID Scale Update; EU AI Act Article 50; Ofcom Attribution Toolkit.

Source boundary: These sources support selected public signals and evidential interpretation. They do not prove global adoption, global failure rates, private system behaviour, or legal sufficiency in any individual case.

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

Synthetic-media evidence does not fail only when a file is fake. It fails when the record cannot carry the claim placed on it.

Layer	Failure question	Common weakness in synthetic-media provenance
Event	Did the thing happen?	Content resembles an event without reconstructable origin.
Record	Was it recorded at the time?	The label or provenance record appears after scrutiny begins.
Context	Does the record explain meaning, authority, and limits?	A label, watermark or credential exists but users do not know what it proves.
Custody	Can control and movement be shown?	Provider, platform, CDN, editor or camera workflow controls the strongest record.
Trust Boundary	Did the record leave the system that may later be questioned?	The evidence remains captive inside the generator, platform or account.
Verification	Can the claim be checked without belief?	A public claim depends on private logs or a proprietary detector.
Permanence	Will the proof survive time and challenge?	Metadata, portal results, labels and links can disappear or change.

False confidence patterns

False confidence	Why it is weak	Stronger posture
“The content has credentials.”	Credentials support provenance assertions, not truth, consent, legality, ownership or completeness.	Credential plus source file, edit history, publication record, consent/rights record and independent receipt.
“The platform labelled it.”	A label is a user-interface claim; it may not show trigger, timing, confidence, correction or appeal history.	Label decision record, display capture, trigger basis, appeal/correction trail and verifier state.
“The watermark proves it.”	Watermarks may identify generated or altered content, but robustness and interpretation depend on method and transformation.	Watermark result plus manifest, verifier version, asset hash, chain of custody and human review boundary.
“The camera signed it.”	Capture provenance strengthens origin, but later edit, transfer, compression and publication can still break the chain.	Capture credential, transfer log, edit manifest, approval record, delivery record and public verification capture.
“The detector said synthetic.”	Detector outputs are probabilistic or method-specific and may not explain consent, source or truth.	Detector result with model/version, confidence limits, source file, preservation record and human review.
“The policy says AI content is labelled.”	A policy is not implementation evidence.	Deployment test, sample outputs, exception log, machine-readable marker proof and user-visible disclosure record.

Five findings

1. The adoption story is shifting from provenance creation to provenance survival

C2PA, Content Credentials, SynthID and camera provenance are no longer merely experimental. The harder evidential question is whether the signal survives upload, transformation, platform display and later verification.

Evidence basis: C2PA Technical Specification; OpenAI Provenance Update; Google SynthID Scale Update; Washington Post C2PA Platform Test; GPT-Image-2 Twitter Dataset.

Source boundary: Sources support public adoption and selected breakage signals. They do not measure global retention or label accuracy.

2. A label is not proof; it is a user-interface claim that needs its own evidence trail

Platform labels can inform users, but the evidentially important record is how the label was triggered, displayed, challenged, corrected and preserved.

Evidence basis: YouTube Altered/Synthetic Disclosure; YouTube Visible AI Labels Update; Meta AI Labelling Approach; European Commission AI Labelling Code.

Source boundary: Sources support platform and regulatory movement toward labelling. They do not prove accuracy of any individual label.

3. C2PA and watermarking solve different evidence problems, and neither solves consent or truth alone

C2PA can support provenance statements. Watermarking can support generated-content identification. Neither proves that the depicted event occurred, that the likeness was authorised, or that a rights claim is valid.

Evidence basis: C2PA Technical Specification; NIST Synthetic Content Guidance; Google DeepMind SynthID; FTC Impersonation Rulemaking; Ofcom Attribution Toolkit.

Source boundary: Sources support technical and regulatory boundaries. They do not determine legality or factual truth in any individual media item.

4. Regulation is turning marking into an operational evidence obligation

The EU AI Act and related code work make synthetic-content marking a compliance architecture problem: what was marked, when, how, for whom, and whether the marking was machine-readable and human-understandable.

Evidence basis: EU AI Act Article 50; European Commission AI Labelling Code; Watermark Adoption Under EU AI Act Study.

Source boundary: Sources support the trajectory and implementation pressure. They are not legal advice and do not assess any provider’s compliance.

5. Capture-side provenance is useful but incomplete unless the custody chain continues after capture

Camera provenance can strengthen first-capture claims, especially for newsrooms, but the evidential risk moves immediately to transfer, editing, publishing, platform ingestion and long-term verification.

Evidence basis: Canon Authenticity Imaging System; C2PA Technical Specification; Cloudflare Content Credentials Preservation.

Source boundary: Sources support capture and infrastructure adoption signals. They do not prove end-to-end integrity for every workflow.

Evidence-signal scorecard

The scorecard renders the selected public signals classified in this report.

Dataset basis: Derived from the 20 selected evidence signals classified in this report. See Source-to-Claim Map and Source Register.

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

Visual chart summary

The charts in this report render selected public evidence signals classified by EviWrite. They show distribution by primary failure type, sector, claim category, region and forward evidence-pressure trajectory.

Dataset basis: Derived from the selected evidence signals and forecast signals classified in this report.

Chart boundary: These charts are not global incidence, market sizing, telemetry, legal findings, or probability forecasts.

Visual chart summary

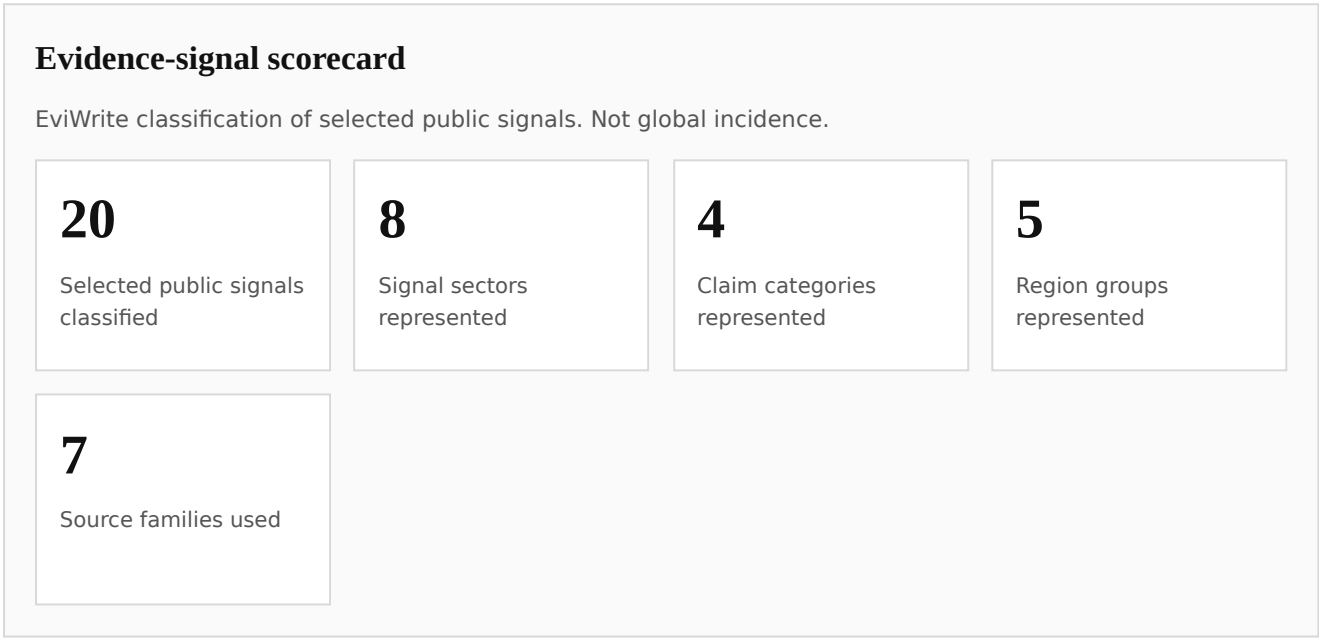

Selected signals by primary evidence weakness

EviWrite classification of selected public signals. Not global incidence.

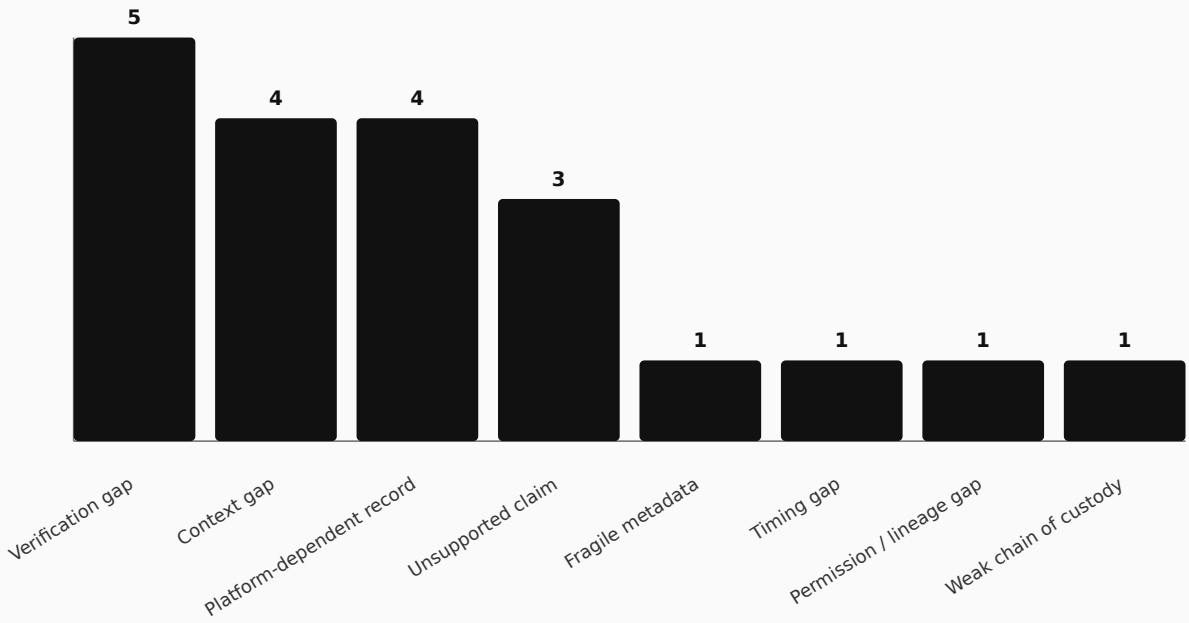

Selected signals by sector

EviWrite classification of selected public signals. Not global incidence.

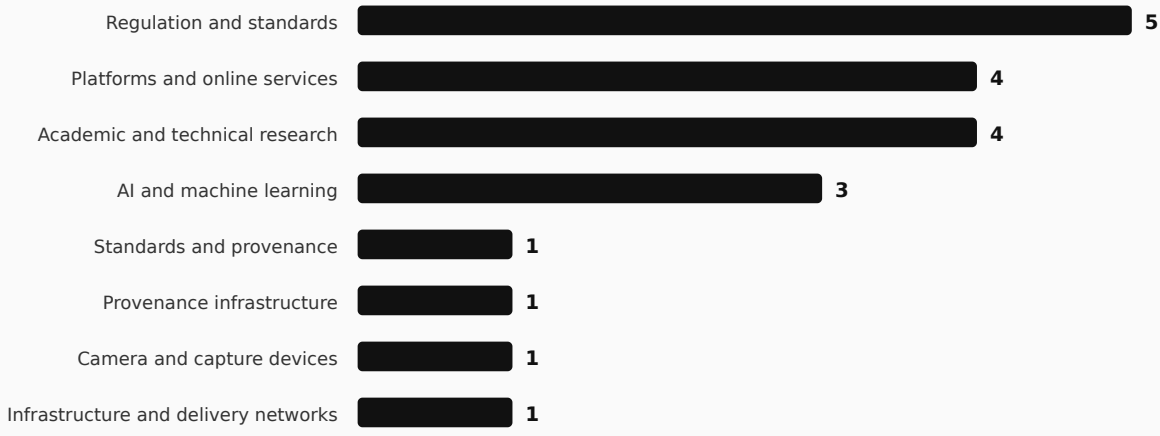

Selected signals by claim category

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

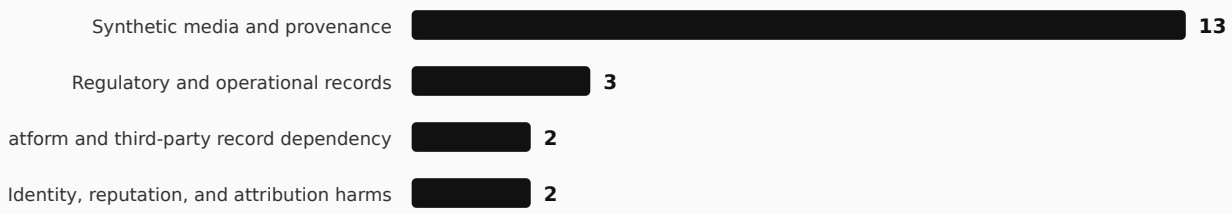

Selected signals by region

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.

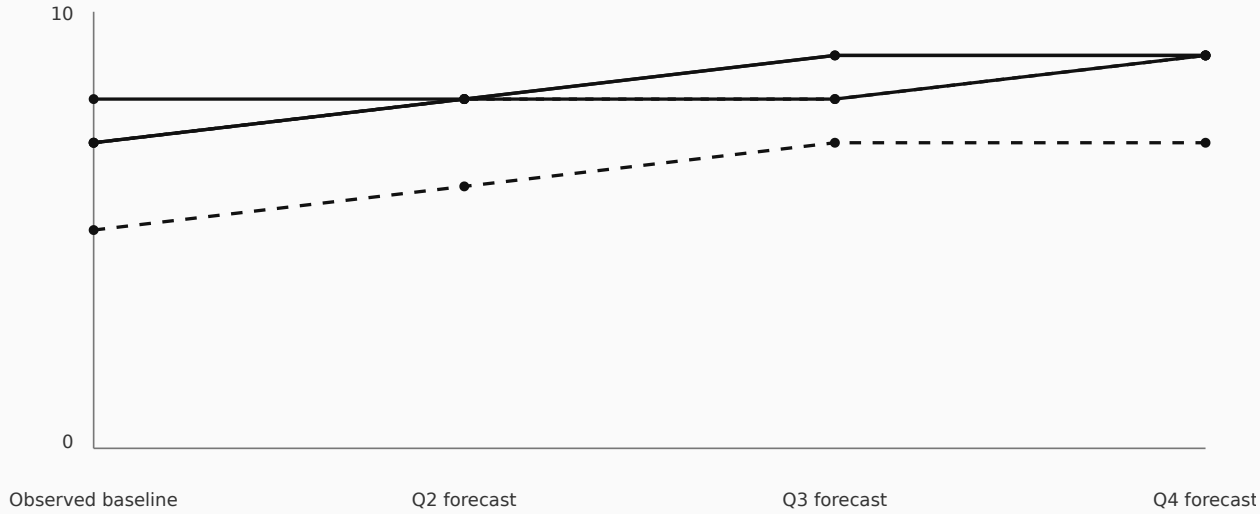

1 Platform retention and display pressure
2 Layered provenance architecture pressure
3 Regulatory marking-to-evidence pressure
4 Capture-side provenance pressure
5 Identity and likeness consent evidence pressure

Evidence failure types by primary weakness

The selected signal set is concentrated in verification gaps, context gaps and platform-dependent records. That pattern matters more than the raw counts. The public adoption layer is strengthening, but the proof layer still breaks when a signal has to travel across systems.

Dataset basis: Derived from `primaryFailureType` classifications in the evidence signal register.

Chart boundary: Primary weakness only. Secondary weaknesses are not counted in the failure-type chart.

Claim-category breakdown

Synthetic-media provenance dominates the selected signal set, but regulatory records, platform dependency and identity harms are already inseparable from the subject. The practical claim is rarely only “is this AI?” It is usually “who made it, who authorised it, what changed, who displayed it, and what can now be checked?”

Dataset basis: Derived from `claimCategory` classifications in the evidence signal register.

Chart boundary: Claim categories are EviWrite classifications for selected public signals, not global claim frequency.

Sector pressure ranking

Regulation, platforms, technical research and AI-provider sources form the main pressure surface in the selected register. That does not mean other sectors are less exposed. It means the visible public records are currently concentrated where synthetic-media systems are built, labelled, challenged and regulated.

Dataset basis: Derived from `sector` classifications in the evidence signal register.

Chart boundary: Sector counts do not represent global prevalence.

Regional selected-signal breakdown

North America is over-represented in this selected signal set because provider, platform, research and journalism sources are concentrated there. EU signals are significant because Article 50 moves marking and labelling toward operational evidence. UK signals are represented through Ofcom attribution work. Asia-Pacific appears through camera-provenance adoption.

Dataset basis: Derived from `region` classifications in the evidence signal register.

Chart boundary: Regional selected-signal breakdown is not a global prevalence measure.

Forecast / forward pressure signals

Forward pressure is strongest in five areas: platform retention and display, layered provenance architecture, regulatory marking evidence, capture-side chain continuity, and identity/likeness consent records.

Source basis: OpenAI Provenance Update; Google SynthID Scale Update; EU AI Act Article 50; European Commission AI Labelling Code; Ofcom Attribution Toolkit; Washington Post C2PA Platform Test; GPT-Image-2 Twitter Dataset.

Forecast boundary: Forecast scores are qualitative EviWrite evidence-pressure scores. They are not incidence predictions, legal predictions, market forecasts or probability estimates.

Forward pressure area	Evidence question
Platform retention and display	Can the signal survive upload, recompression, display, correction and later verification?
Layered provenance architecture	Which signal supports origin, AI generation, edit history, consent or publication state?
Regulatory marking evidence	Can marking be shown in machine-readable and human-facing form at the right time?
Capture-side provenance	Does the chain continue from sensor to edit, publication and verification?
Identity and likeness consent	Can the user prove who authorised likeness, voice, performer, jurisdiction and duration?

Case studies

Case study 1: C2PA video uploaded into social platforms

A public test using AI-generated video with provenance metadata showed that user-facing platform disclosure can fail even when the source file contains credentials.

Source basis: Washington Post C2PA Platform Test; C2PA Technical Specification.

Evidence weakness: The evidence problem moved from generated-file provenance to platform ingestion, retention and display.

What stronger evidence would have required:

Weak record	Stronger evidence
Source file with credentials only	Source file, upload-time preservation check, post-upload credential verification, public label capture, and platform handling record
Buried or absent user-facing disclosure	Visible label placement, timestamped disclosure display evidence, and correction/appeal trail

Evidence boundary: This is a selected platform test, not a comprehensive audit or global incidence claim.

Case study 2: OpenAI moves toward layered provenance for generated images

OpenAI’s public provenance update pairs C2PA metadata with SynthID watermarking and a public verification tool, reflecting a shift away from metadata-only reliance.

Source basis: OpenAI Provenance Update; OpenAI C2PA Help; Google DeepMind SynthID.

Evidence weakness: Layering improves resilience but requires clear boundaries. C2PA, watermarking and verification portals each prove different things.

What stronger evidence would have required:

Weak record	Stronger evidence
Provider says content was generated by its tool	C2PA manifest, watermark verification result, generation account/workflow record, export record, and independent receipt
One verification portal result	Portal result plus source file hash, timestamped verification capture, verifier version, and chain of custody

Evidence boundary: Provider statements describe intended system behaviour. They do not prove every downstream copy remains verifiable.

Case study 3: Capture provenance for news organisations

Camera-provenance systems signal movement toward capture-side evidence for professional news workflows.

Source basis: Canon Authenticity Imaging System; C2PA Technical Specification; Content Credentials Adoption Site.

Evidence weakness: The camera can strengthen the beginning of the chain, but custody and verification must continue through edit, delivery and publication.

What stronger evidence would have required:

Weak record	Stronger evidence
Camera credential only	Device credential, signing-key status, transfer log, edit manifest, newsroom approval record, publication record, and verification page
Final published image only	Linked capture-to-publication chain with hashes, timestamps, edit assertions and archived verification state

Evidence boundary: Manufacturer announcement is an adoption signal, not an independent audit of deployment or newsroom use.

Case study 4: EU synthetic-content labelling moves toward compliance evidence

EU AI Act Article 50 and the Commission’s labelling code process convert synthetic-media labels into records that may need to be demonstrable, not merely visible.

Source basis: EU AI Act Article 50; European Commission AI Labelling Code; Watermark Adoption Under EU AI Act Study.

Evidence weakness: The legal pressure is not simply “put a label on it”; the evidence question is whether marking happened reliably, in the right mode, at the right time, and with preserved proof.

What stronger evidence would have required:

Weak record	Stronger evidence
Product page says AI content is labelled	Marking design record, deployment test, output samples, user-facing disclosure capture, machine-readable marker verification and exception log
Policy-only compliance claim	Versioned procedure, system inventory, audit sample, incident/correction process and retention policy

Evidence boundary: This case study discusses public legal and policy signals. It is not legal advice or compliance assessment.

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, C2PA handling and disclosure visibility	Platform-dependent record	Original asset, upload proof, post-ingestion verification, visible label capture and correction trail
AI and machine learning	Synthetic media and provenance	C2PA plus SynthID and public verification portals	Verification gap	C2PA manifest, watermark result, generation/export record and independent receipt
Regulation and standards	Regulatory and operational records	AI Act transparency and labelling code pressure	Timing gap / context gap	Marking design record, deployment test, machine-readable marker proof and user-facing disclosure capture
Camera and capture devices	Synthetic media and provenance	Capture-side C2PA adoption for news workflows	Weak chain of custody	Device credential, transfer log, edit manifest, approval record and publication verification page
Infrastructure and delivery networks	Platform and third-party record dependency	Image/CDN pipelines preserving or stripping credentials	Platform-dependent record	Delivery pipeline preservation record and post-transform verification capture

Minimum Evidence Records

Area	Minimum Evidence Record	Why it matters
AI-generated image or video	Source file hash, generation tool/model/version, generation timestamp, C2PA manifest, watermark result, export record, verification capture	Separates creation from later display and allows a verifier to check the asset, not only the claim.
Platform-labelled synthetic media	Upload file, upload timestamp, platform label trigger basis, visible label capture, correction/appeal trail, platform policy version	Treats the label as an evidence event, not just a screen element.
Camera-origin photojournalism	Device credential, signing-key status, capture timestamp, transfer log, edit manifest, newsroom approval, publication verification page	Extends provenance from sensor to public trust boundary.
Likeness or voice use	Consent record, scope, jurisdiction, duration, asset linkage, approval trail, revocation/takedown process	Proves authority separately from technical generation.
Regulatory marking compliance	System inventory, marking method, machine-readable marker test, user-facing disclosure capture, exception log, retention policy	Converts labelling into demonstrable operational evidence.
Verification portal result	Asset hash, portal used, verifier version, result timestamp, screenshot plus machine-readable result, limitation note	Prevents a portal result becoming a context-free screenshot.
Publisher provenance workflow	Source asset, edit history, review trail, publication timestamp, CDN/transformation preservation test, public verification URL	Shows whether provenance survived the publication chain.

Recommendations by audience

Audience	Actions
Creators	Preserve original files, exports, edit history and publication evidence before posting. Treat credentials as useful but incomplete.
Businesses	Map synthetic-media workflows from generation to publication. Require vendor evidence for marking, consent, retention and correction.
Legal	Separate origin, truth, consent, authorship and publication claims. Request source files and verification captures, not screenshots alone.
Providers	Document how marks are applied, retained, displayed, removed and corrected. Publish verifier limits plainly.
AI teams	Treat watermarking and C2PA as separate controls. Retain model/version, generation, output, export and verification logs.
Public institutions	Require visible and machine-readable marking evidence for public-facing synthetic content. Preserve procurement and vendor evidence.
Education and research	Preserve provenance through collection, transformation and publication. Do not treat detector output as a complete integrity finding.
Media and publishing	Build capture-to-publication provenance records for high-stakes imagery. Explain credentials without overstating them.

Methodology and limitations

This report uses selected public sources available or materially relevant by 27 May 2026. Signals were selected where they exposed a practical evidential question around source, timing, control, sequence, verification or limits.

Selected signals were classified by sector, region, claim category, source type and evidence failure type. Chart values are derived from the signal register and forecast register. They are not global incidence, telemetry, market sizing, legal findings or probability forecasts.

Private evidence may exist that is not publicly visible. Absence of public evidence is not proof of absence. Provider and platform notices are treated as public statements, not independent forensic audits. Technical papers and journalism are used for selected-signal and limitation analysis, not for universal conclusions.

Deep source appendix

The source register is intentionally structured as an evidential source register rather than a bibliography. Each source is classified by source family, reliability tier, jurisdiction, sector, claim category, evidential relevance and limitation.

Source methodology

Sources were included where they helped answer at least one of the following: what provenance capability exists, what public adoption signal is visible, where the evidence chain can fail, what regulatory pressure is emerging, or what

Minimum Evidence Record would make a claim more defensible.

Core repeatable source spine

The core spine for this report uses C2PA and Content Credentials material, NIST synthetic-content guidance, public provider notices, platform policies, regulator sources, technical research and selected distribution-chain tests.

Report-family source module

The synthetic media and provenance module uses C2PA specifications, Content Credentials adoption material, platform labelling policies, regulator deepfake and attribution materials, identity-harm public records, camera/capture provenance signals, verification portal behaviour and watermarking research.

Source quality tiers

Primary sources ground factual claims about standards, policies, product notices and regulatory direction. Recognised datasets and technical papers support pattern and limitation analysis. Journalism is used as selected public-signal evidence where it tests or reports platform behaviour. None of these sources is used to prove global incidence.

Source group synthesis

The source set points to one conclusion: adoption is no longer the only issue. Synthetic-media provenance is becoming a multi-surface evidence problem across creation, capture, transformation, platform display, public verification and later challenge.

Source-to-claim map

Claim	Support level	Source basis	Boundary
Provenance adoption is moving toward layered manifests, watermarks, verification portals and labels.	Strong	C2PA Technical Specification; Content Credentials Adoption Site; OpenAI Provenance Update; Google SynthID Scale Update	Not a claim of universal adoption.
The weak point is post-generation survival through platform and infrastructure handling.	Medium-high	Washington Post C2PA Platform Test; GPT-Image-2 Twitter Dataset; Cloudflare Content Credentials Preservation	Selected signals only, not global incidence.
AI labels are user-interface evidence claims requiring their own trail.	Medium-high	YouTube Altered/Synthetic Disclosure; YouTube Visible AI Labels Update; Meta AI Labelling Approach; European Commission AI Labelling Code	Label accuracy is not measured.
C2PA and SynthID do not prove truth, consent, legality, ownership or completeness alone.	Strong	C2PA Technical Specification; NIST Synthetic Content Guidance; Independent C2PA Security Analysis; FTC Impersonation Rulemaking; Ofcom Attribution Toolkit	No legal outcome assessed.
Synthetic-content labelling is becoming an operational evidence obligation.	Strong	EU AI Act Article 50; European Commission AI Labelling Code; Watermark Adoption Under EU AI Act Study	Not legal advice.
Capture-side provenance needs continued custody evidence after capture.	Strong	Canon Authenticity Imaging System; C2PA Technical Specification; Cloudflare Content Credentials Preservation	Not proof of end-to-end integrity in all workflows.

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.

The report includes a source register, source-to-claim map, selected signal register, forecast signal register, chart limitations, proof limits and body-level source support. External assurance has not been commissioned for this version.

Version and audit record

Field	Value
Report number	SMPR-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 first public/test version

Future versions may include a report PDF hash, source register hash, EviWrite receipt and public verification link.