EVIWRITE FORMAL REPORT

Quarterly Evidence Failure Report: Q1 2026

Q1 2026

A quarterly EviWrite proof-landscape report examining selected public evidence-failure signals from Q1 2026 across AI training transparency, copyright litigation, synthetic-media labelling, cyber disclosure, AI-cyber governance, platform dependency, and content provenance.

REPORT NUMBER	QEFR-2026-Q1
VERSION	1.0
PUBLICATION DATE	not-issued
STATUS	draft-period-locked-source-reviewed
DOCUMENT CLASS	Public evidential trend report
REFERENCE	QEFR-2026-Q1

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	quarterly-evidence-failure-report-q1-2026.md
Status	draft-period-locked-source-reviewed
Version	1.0
Period	Q1 2026 (2026-01-01 to 2026-03-31)
Prepared by	EviWrite
Report hash	not-issued
PDF hash	not-issued
Receipt	not-issued

Exclusions

- Private incidents not supported by public sources.
- Unverified allegations presented without adequate public record support.
- Legal conclusions about liability, admissibility, infringement, negligence, causation, or regulatory breach.
- Claims requiring access to sealed filings, confidential logs, internal investigations, non-public datasets, forensic images, or private platform records.
- Global incidence measurement, market sizing, cyber telemetry, legal advice, forensic conclusions, or regulatory determinations.

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.

Quarterly Evidence Failure Report: Q1 2026

Executive summary

Q1 2026 shows evidence pressure becoming more specific. AI training disputes, synthetic-media labelling, cyber disclosure, AI-cyber governance, and content provenance all moved toward the same requirement: the record must match the claim placed on it.

This is not a threat-landscape report, legal update, telemetry report, or macro-risk survey. 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 in this report point to one core pattern: general statements are no longer enough. The next evidence-pressure points are likely to include work-level training-source records, AI-generated content labels, cyber known/unknown timelines, platform evidence access, and provenance survival outside controlled workflows.

Source basis for this report

This report combines repeatable international source families with Q1 2026 public records. Repeatable sources provide continuity; quarter-specific 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
Official UK and EU policy material	Timing of obligations, governance pressure, transparency direction, policy evidence questions	Determining compliance by any organisation
Technical standards and technical governance material	Provenance capabilities, AI-cyber record requirements, control boundaries	Proving truth, consent, legality, ownership, or completeness
Public company filings and official cyber material	Incident chronology, disclosure pressure, affected-party dependency, known/unknown records	Full forensic reconstruction
Litigation reporting	Public dispute signals around AI training, source use, output, attribution, and rights pressure	Legal conclusions beyond the reported matter
Stakeholder adoption material	Visible ecosystem pressure around content authenticity	Proof of incidence, effectiveness, or platform survival

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
Event	Did the thing happen?	Event is asserted but not reconstructed
Record	Was it recorded at the time?	Evidence is created after pressure begins
Context	Does the record explain meaning, authority, and limits?	Screenshot, output, label, or notice exists without enough context
Custody	Can control and movement be shown?	Platform, vendor, account, or dashboard controls the strongest record
Trust Boundary	Did the record leave the system or party that may later be questioned?	Evidence remains inside the system being questioned
Verification	Can the claim be checked without belief?	Public claim depends on private logs or unsupported assertion
Permanence	Will the proof survive time and challenge?	Metadata, dashboards, links, or platform records may disappear or change

False confidence patterns

False confidence	Why it is weak	Stronger posture
“We disclosed our AI policy”	A policy does not show which sources, models, prompts, outputs, reviews, or permissions were used	Source register, model/version record, workflow log, permission basis, review trail
“The AI output is labelled”	A label may show notice, not truth, consent, legality, ownership, or source lineage	Label plus technical marking method, generation record, human-review context, and claim boundary
“The platform has the provenance data”	Platform processing can strip, transform, hide, or control the strongest evidence	Exportable credential, source file, independent timestamp/hash, platform-handling record
“The incident was disclosed”	Disclosure is not a reconstructed incident record	Discovery log, known/unknown register, materiality basis, notification and amendment trail
“We use AI securely”	Secure AI cannot be shown without records	AI inventory, data-flow map, security controls, AI-use log, human-review and exception record

Five findings

Evidence pressure has moved from policy intent to record granularity

Q1 2026 signals show that general governance statements are no longer enough. AI, cyber, provenance, and copyright claims increasingly need source-specific, time-specific, role-specific, and system-specific records.

Evidence basis: UK Copyright and AI Report and Impact Assessment.; EU AI-generated content Code of Practice materials.; CIRCIA town hall notice and CISA CIRCIA page.; SEC Form 8-K Item 1.05 and UFP Technologies cyber filing.; NIST Cyber AI Profile workshop material.

Source boundary: These sources support the pattern that public obligations and disputes are becoming record-specific. They do not prove non-compliance, infringement, cyber causation, or private evidence absence.

Implication: Organisations need records that match the claim being made, not generic policies written after pressure starts.

AI training disputes are becoming work-level source disputes

Q1 2026 copyright signals turn on whether specific works were accessed, copied, retained, transformed, reproduced, attributed, excluded, licensed, or reserved. Aggregate dataset language is becoming too blunt for the claim pressure now visible.

Evidence basis: UK Copyright and AI Report.; CourtListener docket records for Britannica/OpenAI and BMG/Anthropic; publicly hosted complaint PDF copies used for pleaded allegation structure; Reuters retained as public-reporting context.

Source boundary: The sources support visible policy and litigation pressure. They do not establish liability, infringement, fair use, or whether any specific private dataset contains a specific work.

Implication: Training-data evidence needs work-level lineage, not merely model-card generalities or broad dataset descriptions.

Transparency without verification is becoming a weak record

Q1 2026 transparency signals show that disclosure, labels, reports, and impact assessments have evidential value only when tied to verifiable underlying records: source, timing, responsible actor, technical method, and limits.

Evidence basis: UK Copyright and AI Report and Impact Assessment.; EU AI-generated content Code of Practice materials.; C2PA Specification 2.4.

Source boundary: These sources support transparency and provenance pressure. They do not show that any specific label, report, or credential is accurate, sufficient, or legally determinative.

Implication: Transparency systems need an evidence layer, not just a disclosure layer.

Synthetic-media labelling is becoming a chain-of-interpretation problem

The strongest Q1 2026 synthetic-media signal is not that labels are useless. It is that labels are easy to overread. Marking and provenance need preserved technical records plus a clear statement of what the signal does and does not prove.

Evidence basis: EU second draft Code of Practice on Marking and Labelling of AI-generated content.; EU Code policy page.; C2PA Specification 2.4.; Content Authenticity Initiative state-of-adoption material.

Source boundary: These sources support labelling/provenance capability and adoption pressure. They do not prove truth, consent, legality, ownership, completeness, or platform survival for any particular item.

Implication: Content authenticity needs evidence survival testing and claim-boundary language, not decorative trust marks.

Cyber evidence is becoming a timed reconstruction asset

CIRCIA rulemaking, SEC incident disclosure, UFP's public filing, NIST's AI-cyber work, and CISA's BRICKSTORM report all point to the same demand: records must reconstruct discovery, control, scope, dependency, known/unknown facts, and response timing while facts are still incomplete.

Evidence basis: CISA CIRCIA page and Federal Register town hall notice.; SEC Form 8-K Item 1.05 and UFP Technologies Form 8-K.; NIST Cyber AI Profile workshop material.; CISA BRICKSTORM analysis report.

Source boundary: These sources support cyber-record pressure and public incident chronology. They do not establish private forensic findings, regulatory breach, attacker identity in any uncited case, or full incident scope.

Implication: Cyber readiness is no longer only containment; it is preservation of the record needed to explain containment.

Evidence-signal scorecard

Metric	Count	Meaning
Public signals classified	12	Selected public-signal count or derived classification count
Primary and secondary failure types referenced	10	Selected public-signal count or derived classification count
Signal sectors represented	6	Selected public-signal count or derived classification count
Region groups represented	4	Selected public-signal count or derived classification count
Claim categories represented	5	Selected public-signal count or derived classification count
Source families used	6	Selected public-signal count or derived classification count

Charts and counts show EviWrite classification of selected public signals. They are not measures of global incidence.

Visual chart summary

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

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

The report charts render the selected signal scorecard, primary failure-type distribution, sector distribution, claim-category distribution, regional distribution, and qualitative forecast trajectories. The chart values reconcile to `evidenceSignals` in YAML.

Visual chart summary

Q1 2026 evidence-failure signals

EviWrite classification of selected public signals. Not global incidence.

12 Public signals classified	**10** Primary and secondary failure types referenced
6 Signal sectors represented	**4** Region groups represented
5 Claim categories represented	**6** Source families used

Evidence failure types by primary weakness

EviWrite classification of selected public signals. Not global incidence.

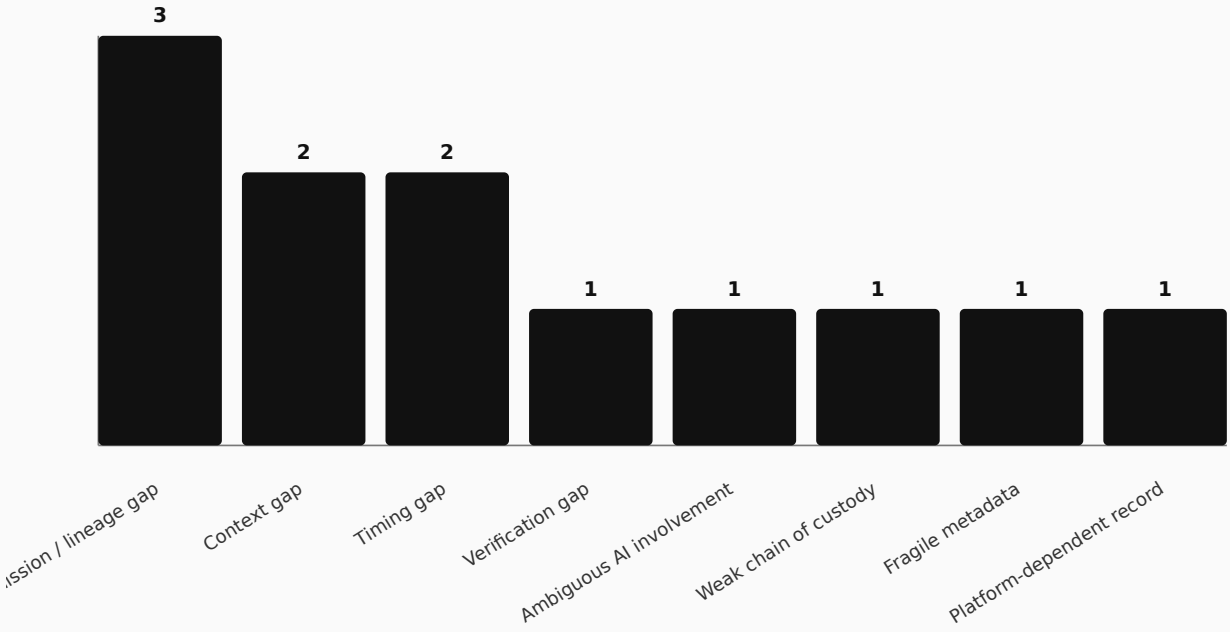

Sectors with visible evidence-failure pressure

EviWrite classification of selected public signals. Not global incidence.

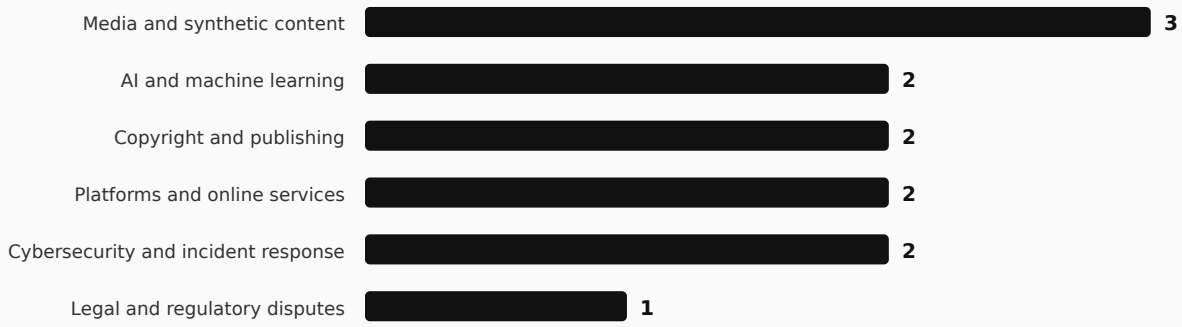

Selected evidence-failure signals by claim category

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

Selected evidence-failure signals by region

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

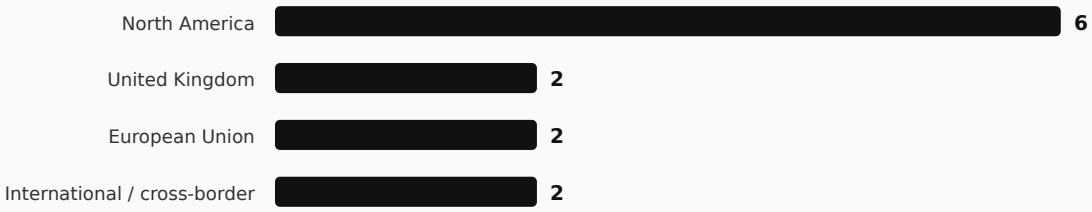

Forecast evidence-pressure trajectories for 2026

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

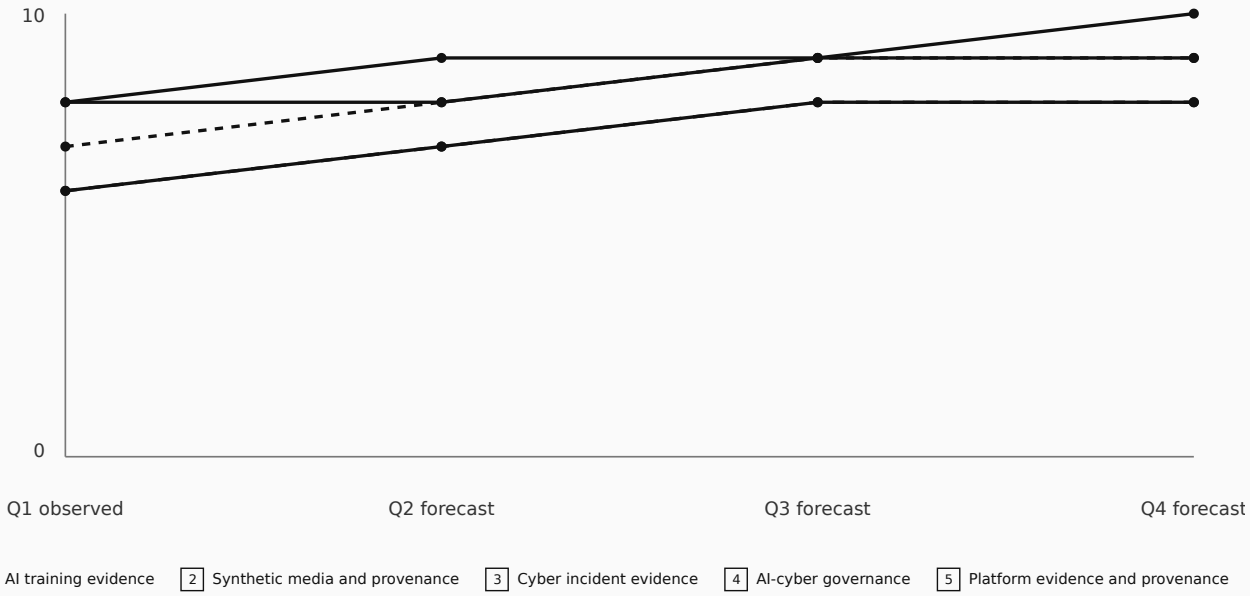

Evidence failure types by primary weakness

Primary failure type	Count	Boundary
Permission / lineage gap	3	Primary weakness count in selected signal set
Context gap	2	Primary weakness count in selected signal set
Timing gap	2	Primary weakness count in selected signal set
Verification gap	1	Primary weakness count in selected signal set
Ambiguous AI involvement	1	Primary weakness count in selected signal set
Weak chain of custody	1	Primary weakness count in selected signal set
Fragile metadata	1	Primary weakness count in selected signal set
Platform-dependent record	1	Primary weakness count in selected signal set

Claim-category breakdown

Claim category	Count	Boundary
AI training-data and dataset lineage	4	Selected public signals only
Synthetic media and provenance	3	Selected public signals only
Regulatory and operational records	3	Selected public signals only
Cyber incident evidence	1	Selected public signals only
Platform and third-party record dependency	1	Selected public signals only

Sector pressure ranking

Sector	Count	Boundary
Media and synthetic content	3	Selected public signals only
AI and machine learning	2	Selected public signals only
Copyright and publishing	2	Selected public signals only
Platforms and online services	2	Selected public signals only
Cybersecurity and incident response	2	Selected public signals only
Legal and regulatory disputes	1	Selected public signals only

Q1 2026 evidence-pressure timeline

Date	Signal	Evidence implication
14 January 2026	NIST Cyber AI Profile workshop material	AI-cyber governance becomes a record problem: systems, data, control, defensive use, and human review.
13 February 2026	Federal Register CIRCIA town hall notice	Critical-infrastructure reporting pressure remains tied to scope, supplier, cloud, MSP/CSP, and open-source dependency records.
24 February 2026	UFP Technologies Form 8-K	Public cyber disclosure shows the known/unknown boundary around affected systems, exfiltration, notification assessment, and ongoing investigation.
5 March 2026	EU second draft Code of Practice on AI-generated content labelling	Synthetic-media labels become evidence objects that need technical method, role, review, and claim-boundary records.
16 March 2026	Britannica and Merriam-Webster v OpenAI reported	AI training disputes extend into source, output, attribution, and alleged endorsement evidence.
18 March 2026	UK Copyright and AI Report and BMG v Anthropic reported	AI copyright pressure centres on work-level source, permission, licensing, transparency, and reproduction records.
31 March 2026	Q1 close	Evidence pressure visible across AI training, cyber disclosure, synthetic-media labelling, content provenance, and AI-cyber governance.

Forecast scoring model

Forecast scores are qualitative EviWrite evidence-pressure scores as at 31 March 2026. They do not measure incidence, probability, market size, legal outcome, or global prevalence.

Each score considers five factors: visible public pressure during Q1, regulatory trajectory, dispute likelihood, evidential difficulty, and actor spread. Scores are comparative within this report only.

Score band	Meaning
0-2	Minimal visible pressure
3-4	Early signal
5-6	Active pressure
7-8	Strong evidence pressure
9-10	Acute evidence pressure

Forecast factor scores

Forecast area	Public pressure	Regulatory trajectory	Dispute likelihood	Evidential difficulty	Actor spread	Rounded score
AI training evidence	9	9	9	10	9	9
Synthetic media and provenance	8	9	9	9	8	9
Cyber incident evidence	8	9	9	9	8	9
AI-cyber governance	7	8	8	8	7	8
Platform evidence and provenance	7	8	8	8	7	8

Forecast evidence-pressure trajectories for 2026

These forecasts are directional EviWrite evidence-pressure assessments based on selected Q1 2026 public signals. They are not incidence predictions, probability forecasts, legal predictions, or market forecasts.

Source basis: UK Copyright and AI Report; UK Copyright and AI Impact Assessment; CourtListener AI litigation docket records, publicly hosted complaint PDF copies, and Reuters public-reporting context; EU AI-generated content Code materials; CISA CIRCIA material; SEC Form 8-K; UFP Technologies Form 8-K; NIST Cyber AI Profile material; C2PA Specification 2.4.

Forecast boundary: This is a qualitative evidence-pressure assessment, not an incidence prediction, probability forecast, legal prediction, or market forecast.

Pressure	Evidence question	Minimum Evidence Record	Direction	Confidence
AI training-source transparency pressure is likely to intensify	Can a rightsholder verify whether a specific work was accessed, copied, licensed, excluded, transformed, reproduced, or cited by a model or AI product?	Work-level source register; Licence and rights-reservation record; Dataset inclusion/exclusion evidence	Increase	High
Synthetic-media labels will create interpretation disputes	Can a viewer or reviewer tell what the label proves, who applied it, what technical method supports it, and what it does not prove?	Marking/labelling method record; Provider/deployer responsibility map; Content credential or equivalent provenance record	Increase	High
Cyber disclosure will turn on timed known/unknown records	Can an organisation show what it knew, when it knew it, who assessed it, what evidence supported disclosure, and what remained unknown?	Incident discovery log; Known/unknown facts register; Materiality and reportability assessment	Increase	High
AI-cyber governance evidence will spread from security teams into enterprise records	Can the organisation evidence which AI systems and cyber tools were used, what data they processed, how they were secured, and how human review occurred?	AI system inventory; Data-flow and leakage-risk map; Model/component security record	Increase	Medium
Platform-controlled provenance will become a portability and contract issue	Can the strongest evidence survive export, upload, editing, compression, screenshotting, platform transformation, account loss, and later dispute?	Evidence-access and export clause; Credential survival test record; Platform handling log	Increase	Medium

Case studies

UK copyright and AI evidence now turns on source, transparency, and licensing records

The UK report and impact assessment show that AI copyright policy depends on records capable of connecting training, transparency, licensing, rights reservation, output transparency, and enforcement.

Source basis: UK Copyright and AI Report; UK Copyright and AI Impact Assessment.

Evidence weakness: A transparency rule is weak if it cannot be connected to specific works, datasets, access events, licences, rights reservations, technical standards, and model-use records.

What stronger evidence would have required:

Weak record	Stronger evidence
General statement that training data was lawfully obtained	Dataset source register, licence basis, rights-reservation handling, and model/version linkage
High-level transparency commitment	Work-level disclosure method, timing record, source categories, exclusion record, and audit trail
Policy impact assumptions	Evidence base separating technical feasibility, economic effect, licensing availability, enforcement practicality, and stakeholder evidence

Evidence boundary: This case study does not determine the legality of AI training, copyright reform, infringement, fair use, or the sufficiency of any proposed UK policy option.

Q1 2026 AI litigation signals expose the limits of aggregate dataset language

Britannica/OpenAI and BMG/Anthropic reporting shows public pressure around work-level source inclusion, output reproduction, attribution, endorsement, lyrics, and alleged unauthorized sources.

Source basis: CourtListener docket records for Britannica/OpenAI and BMG/Anthropic; publicly hosted complaint PDF copies used for pleaded allegation structure; Reuters reporting used as public-reporting context.

Evidence weakness: Public litigation signals increasingly turn on facts that cannot be proved by final outputs or broad dataset descriptions alone.

What stronger evidence would have required:

Weak record	Stronger evidence
Broad dataset description	Work-level source register with ingestion, exclusion, and transformation records
Public denial or fair-use position only	Legal-basis record linked to sources, model versions, output behaviour, and product features
Output example without provenance	Prompt/output log, retrieval path, citation/attribution record, similarity assessment, and review trail

Evidence boundary: This case study treats lawsuits and complaints as public signals. It does not determine infringement, fair use, trademark liability, CMI removal, factual accuracy of allegations, or the content of any private model-training records.

Synthetic-media labels need preserved provenance plus claim boundaries

EU Article 50 code materials and C2PA specification material show the same structural issue: labels and credentials help, but they are not evidence of every claim people may infer from them.

Source basis: EU AI-generated content Code second draft; EU Code policy page; C2PA Specification 2.4; Content Authenticity Initiative state-of-adoption material.

Evidence weakness: A visible label may create false confidence if the underlying generation method, signer identity, edit history, human review, consent, rights, and limitations are not preserved or explained.

What stronger evidence would have required:

Weak record	Stronger evidence
Visible AI-generated label only	Technical marking method, generation record, provider/deployer role record, and preserved label state
Content Credential alone	Credential plus source file, manifest preservation, signer identity, edit history, consent/rights record, and claim-boundary statement
Platform display state	Exportable verification record and evidence of how the platform processed, stripped, transformed, or preserved credentials

Evidence boundary: This case study does not treat provenance metadata as proof of truth, consent, legality, ownership, or completeness.

Cyber disclosure and AI-cyber governance require timed reconstruction, not summaries

CIRCIA rulemaking, SEC Form 8-K incident disclosure, UFP's filing, NIST Cyber AI Profile material, and CISA BRICKSTORM analysis show rising pressure for incident and governance records that can survive incomplete facts and hostile environments.

Source basis: CISA CIRCIA page; Federal Register CIRCIA town hall notice; SEC Form 8-K Item 1.05; UFP Technologies Form 8-K; NIST Cyber AI Profile workshop; official BRICKSTORM Malware Analysis PDF and CISA landing page.

Evidence weakness: Cyber claims often need to be made before the full forensic picture is available, which makes discovery time, known/unknown records, control-layer logs, dependency maps, and amendment triggers critical.

What stronger evidence would have required:

Weak record	Stronger evidence
Incident summary after containment	Discovery log, containment chronology, affected-system map, known/unknown register, and evidence-preservation record
Provider or MSP assurance only	Evidence-access clause, customer-specific logs, supplier-dependency map, and independent preservation rights
AI security policy only	AI system inventory, model/data-flow record, cyber-control mapping, AI-use log, and human review evidence

Evidence boundary: This case study does not determine attack attribution, full incident scope, legal reportability, materiality, regulatory breach, or adequacy of any specific security programme.

Sector, claim-category, and failure-type table

Sector	Claim category	Visible pressure	Primary evidence weakness	Stronger record
AI and machine learning	AI training-data and dataset lineage	UK Copyright and AI report turns training transparency into an evidential design problem	Permission / lineage gap	AI training transparency needs a source register, permission basis, rights-reservation record, and model/dataset linkage before dispute or enforcement pressure begins.
Copyright and publishing	AI training-data and dataset lineage	UK Copyright and AI impact assessment highlights moving evidence base	Context gap	Policy records should distinguish consultation evidence, economic evidence, technical feasibility, rights-management evidence, and operational implementation evidence.
Legal and regulatory disputes	AI training-data and dataset lineage	Britannica and Merriam-Webster v OpenAI adds attribution and hallucination pressure to training-data disputes	Permission / lineage gap	AI providers and rightsholders need work-level source records, retrieval/output records, attribution logs, and claim-boundary records distinguishing training, generation, citation, and endorsement.
Copyright and publishing	AI training-data and dataset lineage	BMG v Anthropic keeps lyric-level training and reproduction evidence under pressure	Permission / lineage gap	Music and text AI workflows need rights-source registers, ingestion records, deduplication/exclusion records, output similarity logs, and licence/exception rationale records.
Media and synthetic content	Synthetic media and provenance	EU AI-generated content Code second draft makes labels evidence objects	Verification gap	Synthetic-media labelling needs content-origin records, marking method, human-review context, user-facing disclosure text, and a claim-boundary statement.
Platforms and online services	Regulatory and operational records	EU Article 50 code process separates provider marking from deployer labelling	Context gap	Operational transparency requires role records, provider/deployer responsibility mapping, platform handling records, and preservation of the marking/labelling state at publication and distribution.
Cybersecurity and incident response	Regulatory and operational records	CIRCIA Q1 2026 town halls keep critical-infrastructure reporting scope unresolved	Timing gap	Critical-infrastructure organisations need incident discovery records, supplier dependency maps, MSP/CSP evidence-access terms, open-source component records, and reporting-decision logs.
Cybersecurity and incident response	Cyber incident evidence	UFP Technologies Form 8-K shows known/unknown cyber disclosure boundaries	Timing gap	Cyber incident evidence should preserve discovery time, containment actions, affected systems, data-scope assessment, legal-notification analysis, board updates, and amendment triggers.
AI and machine learning	Regulatory and operational records	NIST Cyber AI Profile workshop makes AI-cyber controls record-dependent	Ambiguous AI involvement	AI-cyber governance needs system inventories, data-flow maps, model/component records, security-control evidence, AI-use logs, and human review records.
Platforms and online services	Platform and third-party record dependency	BRICKSTORM backdoor report highlights virtualisation-layer custody evidence	Weak chain of custody	Virtualisation incident readiness needs privileged-access logs, management-plane telemetry, configuration baselines, snapshot custody, admin-action records, and external preservation.
Media and synthetic content	Synthetic media and provenance	C2PA 2.4 keeps provenance useful but bounded	Fragile metadata	Provenance systems need source files, manifest preservation, signer identity records, edit history, rights/consent records, and explicit claim limits.
Media and synthetic content	Synthetic media and provenance	Content authenticity adoption becomes a platform-survival problem	Platform-dependent record	Content authenticity programmes need preservation tests, platform handling records, export records, verification UX records, and fallback proof where metadata is stripped.

This table summarises the selected signal set. It does not imply global sector prevalence.

Minimum Evidence Records

A Minimum Evidence Record is the smallest practical record set needed to make a claim more defensible, portable, interpretable, and independently checkable. It does not guarantee legal success, regulatory compliance, or factual certainty.

Area	Minimum Evidence Record	Why it matters
AI training data	Work-level source register; licence/permission basis; rights-reservation and opt-out record; dataset inclusion/exclusion evidence; model/version linkage	Turns transparency into checkable lineage.
AI-assisted outputs	Prompt/output log; source material; human selection, arrangement, edit, review, and approval trail; final publication record	Separates human contribution from machine generation and unsupported authorship claims.
Synthetic-media labelling	Marking method; provider/deployer role record; human-review context; preserved label state; claim-boundary statement	Prevents labels being treated as proof of truth, legality, consent, or ownership.
Content provenance	Source file; C2PA manifest or equivalent; signer identity; edit history; content binding; export/preservation record	Makes provenance checkable and shows where it can fail.
Cyber incident disclosure	Discovery log; known/unknown register; affected-system map; data-scope assessment; materiality/reportability decision; amendment trail	Supports public disclosure while investigations remain incomplete.
AI-cyber governance	AI system inventory; data-flow map; security controls; AI-enabled threat assessment; defensive AI use log; human review	Makes AI security claims auditable instead of rhetorical.
Platform dependency	Evidence-access clause; exportable logs; customer-specific impact record; platform handling record; independent preservation package	Reduces dependence on provider-controlled dashboards.

Recommendations by audience

Not every audience category is equally represented in this quarter’s selected signals. Audience actions identify where the report’s evidence lessons are applicable, not where the quarter produced equal public signal volume.

Audience	Recommended actions
creators	Keep creation, publication, licence, rights-reservation, opt-out, and AI-assistance records before public release.; Preserve source files and evidence packages outside the platforms used for distribution.
businesses	Create claim-specific evidence records for AI use, cyber incidents, supplier dependencies, and public disclosures.; Require exportable provider logs and evidence-access terms in contracts.
legal	Separate allegations, public facts, technical capability, and EviWrite interpretation in evidence reviews.; Demand source-to-claim mapping for AI, cyber, synthetic-media, and platform-dependent disputes.
providers	Design evidence-access, credential preservation, incident export, and customer-specific record capabilities into services.; State what labels, logs, dashboards, and credentials do not prove.
aiTeams	Maintain system inventories, model/version records, source-lineage registers, prompt/output logs, and human-review evidence.; Document AI-cyber risks, data leakage controls, and defensive AI use.
publicInstitutions	Require timed, reviewable records for AI systems, automated decisions, synthetic-media notices, and cyber reporting.; Avoid relying on vendor dashboards as the only authoritative record.
educationResearch	Preserve research/source data, AI-use declarations, platform-export records, and student/researcher contribution trails.; For third-party tools, require logs and breach-impact evidence that can be exported independently.
mediaPublishing	Use provenance and labels with explicit claim boundaries; preserve source files, edit histories, publication timestamps, and rights records.; Treat AI scraping, syndication, and attribution disputes as source-lineage evidence problems.

Methodology and limitations

This report reviews selected public sources materially relevant to Q1 2026. Selection required a clear evidential question around source, timing, control, sequence, verification, or proof limits; relevance to at least one EviWrite audience; and the ability to produce a practical Minimum Evidence Record.

The report is not global incidence analysis. It is not legal advice, a liability finding, a forensic audit, a regulatory finding, or cyber telemetry. Private evidence may exist. Absence of public evidence is not proof of absence. Charts are EviWrite classifications of selected signals. Forecast scores are qualitative.

Deep source appendix

The source register is not a bibliography. It is a controlled evidence map. Official and technical sources carry factual grounding. Court docket indexes carry litigation-existence and chronology support. Journalism and stakeholder sources carry public visibility and context only.

Source reliability matrix

Source family	Strength	Weakness	Report use
Official policy, regulatory, technical, and filing records	Strong support for public facts, formal obligations, technical definitions, filing language, dates, and official framing	Does not prove private compliance, private evidence quality, causation, liability, or incident completeness	Primary factual grounding and timing/control/verification pressure analysis
Court docket records	Strong support for dispute existence, parties, court, docket chronology, and public procedural posture	Does not prove pleaded allegations, infringement, fair use, damages, causation, or model/source behaviour	Primary litigation-signal grounding, paired with strict allegation boundaries
Journalistic litigation reporting	Useful for public visibility, reported allegations, and discovery of disputes	Weaker than pleadings/dockets for record support; must not carry legal conclusions	Context only, not primary proof where court records are available
Stakeholder ecosystem material	Shows visible pressure, adoption narrative, and ecosystem direction	Not independent measurement and not proof of implementation quality or adoption incidence	Forward-pressure and false-confidence analysis, bounded by technical standards and primary records

Source group synthesis

Source methodology

Sources are classified by evidential reliability, not whether EviWrite agrees with their conclusions. Official materials support public facts and regulatory direction. Technical standards support capability and limit analysis. Public filings support public chronology and disclosure boundaries. Journalism supports visible dispute signals only.

Core repeatable source spine

The repeatable source spine used where relevant includes official regulator/government sources, official court or litigation sources where accessible, technical standards, public company notices, specialist technical material, and journalism as public signal.

Report-family source module

This report uses the evidence-failure module with copyright-AI, synthetic-media provenance, cyber-resilience, platform dependency, and technical-standards submodules.

Source quality tiers

Tier	Use
Primary	Factual grounding where source directly supports the claim
Recognised dataset	Pattern support where dataset methodology permits
Specialist	Context and interpretation
Journalistic	Discovery and public-signal support
Stakeholder pressure	Visible pressure, not proof of underlying fact

Full source register

Label	Publisher	Tier	What it is	Used for	Limitations	URL
UK Copyright and AI Report	UK Department for Science, Innovation and Technology / Department for Culture, Media and Sport / Intellectual Property Office	primary	Official UK report to Parliament on copyright works in AI development, covering training, transparency, technical standards, licensing, enforcement, computer-generated works, and digital replicas.	Policy direction and public evidence-pressure context for AI training, transparency, licensing, and enforcement.	Not used to decide whether any particular AI system used any particular work.	https://www.gov.uk/government/publications/report-and-impact-assessment-on-copyright-and-artificial-intelligence/report-on-copyright-and-artificial-intelligence
UK Copyright and AI Impact Assessment	UK Department for Science, Innovation and Technology / Department for Culture, Media and Sport / Intellectual Property Office	primary	Official impact assessment accompanying the UK Copyright and AI report, identifying policy options and areas where further evidence gathering is needed.	Evidence-gathering pressure and policy-option context.	Not used to quantify global economic impact or determine legal outcomes.	https://www.gov.uk/government/publications/report-and-impact-assessment-on-copyright-and-artificial-intelligence/copyright-and-artificial-intelligence-impact-assessment
Reuters Britannica v OpenAI	Reuters	journalistic	Reuters report on Britannica and Merriam-Webster filing suit against OpenAI, including allegations over training use, near-verbatim summaries, traffic diversion, and alleged trademark-related hallucination issues.	Public litigation signal and claim-category classification.	Not used as proof that alleged infringement, trademark infringement, or hallucination-based liability occurred.	https://www.reuters.com/legal/litigation/encyclopedia-britannica-sues-openai-over-ai-training-2026-03-16/
Reuters BMG v Anthropic	Reuters	journalistic	Reuters report on BMG's lawsuit against Anthropic alleging use of copyrighted lyrics in training and reproduction by Claude.	Public litigation signal and evidence-pressure classification.	Not used as proof that Anthropic infringed copyright or that any court accepted the allegations.	https://www.reuters.com/legal/litigation/bmg-sues-anthropic-using-bruno-mars-rolling-stones-lyrics-ai-training-2026-03-18/
EU AI label Code second draft	European Commission	primary	European Commission publication of the second draft of the AI-generated content marking and labelling Code of Practice, with feedback closing 30 March 2026 and Article 50 rules due to apply from 2 August 2026.	Regulatory-development evidence for labelling and synthetic media transparency pressure.	Not used as final law or proof that any provider complied or failed to comply.	https://digital-strategy.ec.europa.eu/en/library/commission-publishes-second-draft-code-practice-marking-and-labelling-ai-generated-content

Label	Publisher	Tier	What it is	Used for	Limitations	URL
EU AI-generated content Code policy page	European Commission	primary	European Commission policy page describing Article 50 transparency obligations for marking AI-generated content and labelling deepfakes and certain AI-generated publications.	Regulatory context for synthetic media and AI-generated content transparency.	Not used as proof of implementation by any specific provider.	https://digital-strategy.ec.europa.eu/en/policies/code-practice-ai-generated-content
CISA CIRCIA page	Cybersecurity and Infrastructure Security Agency	primary	CISA public CIRCIA page describing future cyber incident reporting requirements and Q1 2026 town hall engagement.	Cyber reporting rulemaking direction and evidence-readiness pressure.	Not used as evidence that final CIRCIA rules were operative in Q1 2026.	https://www.cisa.gov/topics/cyber-threats-and-advisories/information-sharing/cyber-incident-reporting-critical-infrastructure-act-2022-circia
Federal Register CIRCIA town halls	Federal Register / CISA	primary	Federal Register notice on Q1 2026 CIRCIA rulemaking town hall meetings, including issues such as managed service providers, cloud service providers, open-source software, and repositories.	Rulemaking process and reporting-scope evidence pressure.	Not used as final rule text or compliance finding.	https://www.federalregister.gov/documents/2026/02/13/2026-02948/cyber-incident-reporting-for-critical-infrastructure-act-circia-rulemaking-town-hall-meetings
SEC Form 8-K Item 1.05	U.S. Securities and Exchange Commission	primary	SEC Form 8-K text requiring material cybersecurity incident disclosures to describe material aspects of nature, scope, timing, and material impact or reasonably likely material impact.	Cyber disclosure obligation context.	Not used as evidence that any listed issuer breached securities law.	https://www.sec.gov/files/form8-k.pdf
UFP Technologies 8-K	U.S. Securities and Exchange Commission / UFP Technologies, Inc.	primary	Public Form 8-K in which UFP Technologies disclosed a cybersecurity incident detected around 14 February 2026, affected IT functions, possible stolen or destroyed data, ongoing investigation, and notification assessment.	Public incident chronology and disclosure-boundary signal.	Not used as independent forensic confirmation of the attacker's identity, full scope, or legal consequences.	https://www.sec.gov/Archives/edgar/data/914156/000162828026011152/ufpt-20260219.htm

Label	Publisher	Tier	What it is	Used for	Limitations	URL
NIST Cyber AI Profile workshop	NIST National Cybersecurity Center of Excellence	primary	NIST workshop material for a Cyber AI Profile supporting cybersecurity programs as they manage AI-related risks, including securing AI systems, minimizing data leakage, defending against AI-enabled attacks, and using AI in cyber defence.	Technical and governance evidence-pressure context.	Not used as final binding standard or proof of any organisation's cyber maturity.	https://www.nccoe.nist.gov/sites/default/files/2026-01/Cyber AI Profile Workshop Slides - January 2026_1.pdf
CISA BRICKSTORM report	Cybersecurity and Infrastructure Security Agency	primary	CISA analysis report on BRICKSTORM backdoor activity affecting VMware vSphere environments including vCenter, ESXi, and Aria Automation.	Cyber technical signal and evidence-readiness classification.	Not used as proof any specific organisation was compromised unless named by the source.	https://www.cisa.gov/news-events/analysis-reports/ar25-338a
C2PA Specification 2.4	Coalition for Content Provenance and Authenticity	primary	C2PA technical specification defining assets, manifests, claims, assertions, content bindings, provenance data, hard bindings, soft bindings, and authenticity properties.	Technical capability and limits of provenance signals.	Not used as proof that any credentialed asset is truthful, lawful, consensual, complete, or owned by a particular party.	https://spec.c2pa.org/specifications/specifications/2.4/specs/C2PA_Specification.html
CAI State of Content Authenticity 2026	Content Authenticity Initiative	stakeholder-pressure	Content Authenticity Initiative article describing 2026 as a turning point for content authenticity, interoperable provenance, and content credentials adoption.	Adoption context and public-signal pressure.	Not used as independent incidence data or proof that content credentials survive all distribution channels.	https://contentauthenticity.org/blog/the-state-of-content-authenticity-in-2026

Source-to-claim map

Claim ID	Claim	Support level	Sources
claim-cross-domain-record-granularity	Q1 2026 public signals show evidence pressure moving from broad governance claims toward source-specific, time-specific, role-specific, and system-specific records.	strong	UK Copyright and AI Report; EU AI label Code second draft; CISA CIRCIA page; SEC Form 8-K Item 1.05; NIST Cyber AI Profile workshop
claim-uk-ai-copyright-lineage	The UK Q1 2026 copyright-and-AI report makes AI training transparency, licensing, rights reservation, output transparency, and enforcement record-dependent.	strong	UK Copyright and AI Report; UK Copyright and AI Impact Assessment
claim-ai-training-litigation-lineage	Q1 2026 litigation reporting shows visible pressure around work-level source inclusion, reproduction, attribution, endorsement, and lyrics/source records in AI systems.	medium	Reuters Britannica v OpenAI; Reuters BMG v Anthropic
claim-policy-evidence-base	AI copyright policy choices are constrained by the quality and granularity of available evidence on technical feasibility, licensing, transparency, and economic effect.	medium-high	UK Copyright and AI Impact Assessment; UK Copyright and AI Report
claim-ai-content-labelling-boundaries	AI-generated content labels and deepfake labelling systems need preserved records that connect provider/deployer roles, technical marking, human review, publication context, and user-facing disclosure.	strong	EU AI label Code second draft; EU AI-generated content Code policy page
claim-provenance-not-sufficiency	C2PA-style provenance can support authenticity and content-binding checks, but does not by itself prove truth, consent, legality, ownership, authorship, or completeness.	strong	C2PA Specification 2.4
claim-provenance-adoption-pressure	Content authenticity adoption pressure is increasing, but evidential value depends on whether credentials survive platform, editing, export, and verification workflows.	medium	CAI State of Content Authenticity 2026; C2PA Specification 2.4
claim-cyber-reporting-records	CIRCIA rulemaking pressure makes incident discovery, scope, supplier dependency, reporting decision, and update records operationally important before final reporting pressure arrives.	strong	CISA CIRCIA page; Federal Register CIRCIA town halls
claim-cyber-disclosure-granularity	Public cyber incident disclosure depends on records of nature, scope, timing, material impact, known/unknown facts, and subsequent investigation boundaries.	strong	SEC Form 8-K Item 1.05; UFP Technologies 8-K
claim-ai-cyber-governance-records	AI use in cybersecurity creates evidence requirements around system inventories, data leakage, AI-enabled attacks, defensive AI use, and human review.	medium-high	NIST Cyber AI Profile workshop
claim-virtualisation-control-evidence	Virtualisation and management-plane threats increase the need for privileged-access, configuration, snapshot, management-log, and custody evidence.	medium-high	CISA BRICKSTORM report

Source-hardening note

A May 2026 source-hardening pass corrected signal-boundary drift, paired AI litigation reporting with court-docket sources, separated stakeholder adoption material from independent adoption measurement, and treated the BRICKSTORM landing page and official PDF as cyber-readiness sources rather than synthetic-media support. This note records the review control; it does not change the Q1 2026 evidence horizon.

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.

Control	Status
Source register prepared	Yes
Source-to-claim map prepared	Yes
Selected signal register prepared	Yes
Forecast signal register prepared	Yes
Chart limitations stated	Yes
Forecast limitations stated	Yes
Legal limits stated	Yes
Global incidence boundaries stated	Yes
Body-level source support included	Yes
External review	Not commissioned for this public evidential trend report

Version and audit record

Field	Value
Report number	QEFR-2026-Q1
Version	1.0
Period covered	1 January 2026 to 31 March 2026
Prepared by	EviWrite
Status	Draft-period-locked-source-reviewed
Report hash	Not issued
PDF hash	Not issued
Source register hash	Not issued
Machine-readable register hash	Not applicable
Receipt	Not issued for this public/test version

Glossary

Term	Definition
Selected public signal	A public record, filing, report, notice, standard, or reported dispute selected for evidential relevance.
Evidence failure	A weakness that makes a claim harder to check, reconstruct, interpret, preserve, or rely on.
Minimum Evidence Record	The smallest practical record set needed to make a claim more defensible, portable, interpretable, and independently checkable.
Provenance	Record of origin, history, and interaction of an asset or content item.
Claim boundary	A statement of what a record proves, supports, does not prove, and should not be used to imply.