Federal Agencies
Digital Guidelines Initiative

TCR4CAP:
TIERED COMMUNITY RECOMMENDATIONS FOR CONTENT AUTHENTICITY AND PROVENANCE
WITH A FOCUS ON LAM COLLECTIONS CONTENT IMPACTED BY ARTIFICIAL INTELLIGENCE

September 8, 2026

Version 0.2 – Draft for Public Comment

The FADGI Audio-Visual Working Group

http://www.digitizationguidelines.gov/audio-visual/

Federal Agencies
Digital Guidelines Initiative

TCR4CAP:
Tiered Community Recommendations for Content Authenticity and Provenance

With a focus on LAM collections content impacted by artificial intelligence

Version 0.2 (Draft for Public Comment – open until October 15, 2026)

Last Updated: September 8, 2026

Contact point: Kate Murray (kmur@loc.gov)

Version History

Version	Publication Date	Contributors (in alphabetical order)
0.2 (Draft for public comment)	September 8, 2026	Same as below
0.1 (Draft for public comment)	August 26, 2026	**Library of Congress:** Curtis, Rachel Havemeyer-King, Genevieve Hockstein, Dan Kelly, Paul Murray, Kate **Smithsonian Institution:** Sobers, Kira **Community Collaborators:** Hosale, Charlie (MIT Libraries) Krabbenhoeft, Nick (NYPL) Martinez, Jérôme (MediaArea/FADGI contractor) Ng, Yvonne (WITNESS) Rice, Dave (MediaArea/FAGDI contractor)

Change Log

Version 0.2

- Minor update to changed URL for CAP registry location.

Table of Contents

Glossary

Artificial intelligence or AI: A sub-field of computer science research and practice that aims to develop computational models to simulate or approximate human cognition, behaviour, decision-making, and reasoning.

AI-assisted: AI supports a human-directed process but does not intentionally create, remove, replace, or substantially alter the meaningful content of the object. Examples include: AI adjustments for color balance, reduces noise, corrects exposure, or assists transcription/OCR.

AI-modified: AI materially alters existing content, including generating, removing, replacing, reconstructing, or transforming portions of the object, while the resulting object remains substantially derived from an existing source. In other words, did AI change meaningful content while an existing source remains the basis of the result? Examples include: generative fill replaces a damaged region of a photograph; AI removes an object; AI rewrites passages of an article; AI reconstructs missing audio.

AI-generated: AI is responsible for generating all or a substantial portion of the meaningful content of the object from instructions, prompts, data, or other inputs, rather than primarily modifying an existing source. Examples include: A text-to-image model creates an image; an LLM drafts an article; a music model generates a song from a prompt.

CAP: Content authenticity and provenance

Checksum: A value that (a) is computed by a function that is dependent on the contents of a data object and (b) is stored or transmitted together with the object, for detecting changes in the data. (NIST Glossary: https://csrc.nist.gov/glossary/term/checksum)

Content authenticity: the quality of being genuine, not a counterfeit, and free from tampering, and is typically inferred from internal and external evidence, including its physical characteristics, structure, content, and context." (SAA Dictionary: https://dictionary.archivists.org/entry/authenticity.htm)

C2PA: Coalition for Content Provenance and Authenticity (https://c2pa.org/)

Content Credentials: Technical specification for providing content provenance and authenticity; developed by C2PA. (https://spec.c2pa.org/specifications/specifications/2.4/index.html - this is frequently updated as new versions are released)

Cryptographic hash: A function that maps a bit string of arbitrary length to a fixed length bit string and is expected to have the following three properties: 1) Collision resistance (see Collision resistance), 2) Preimage resistance (see Preimage resistance) and 3) Second preimage resistance (see Second preimage resistance). (NIST Glossary: https://csrc.nist.gov/glossary/term/cryptographic_hash_function)

DIP: Dissemination Information Package, which is the information package transferred from the archive in response to a request by a consumer. (Lavoie/DPC Tech Watch Report 14-02 October 2014: https://www.dpconline.org/docs/technology-watch-reports/1359-dpctw14-02/file)

Digital signature: The result of a cryptographic transformation of data that, when properly implemented, provides a mechanism for verifying origin authentication, data integrity, and signatory non-repudiation. (NIST Glossary: https://csrc.nist.gov/glossary/term/digital_signature)

Digital watermark: A pattern of bits, hidden in a digital file, that can be used to provide protection of intellectual property rights. A digital watermark is designed to be imperceptible in ordinary use; they are invisible in images and inaudible in sound recordings. At most, they appear as slight noise in the signal. They are designed so that they cannot be identified or manipulated without special software. (SAA Dictionary: https://dictionary.archivists.org/entry/digital-watermark.html)

Fit for purpose: Used informally to describe a process, configuration item, IT service, etc., that is capable of meeting its objectives or service levels. Being fit for purpose requires suitable design, implementation, control, and maintenance. (NIST Glossary: https://csrc.nist.gov/glossary/term/fit_for_purpose)

Fixity: Property of a digital file or object being fixed or unchanged. …The most widely used tools for creating fixity information are checksums (like CRCs) and cryptographic hashes (like MD5 and various SHA algorithms), but there are other methods such as expected file size and file count that provide basic fixity information. (NDSA 2014: https://perma.cc/8YTV-QQ3U)

Generative AI or GenAI: The class of AI models that emulate the structure and characteristics of input data in order to generate derived synthetic content. This can include images, videos, audio, text, and other digital content. (NIST Glossary: https://csrc.nist.gov/glossary/term/generative_artificial_intelligence)

IPTC: International Press Telecommunications Council

LAM: Libraries, archives and museums

PREMIS: Preservation Metadata: Implementation Strategies. International standard for metadata to support the preservation of digital objects and ensure their long-term usability. (https://www.loc.gov/standards/premis/)

Provenance: Information regarding the origins, custody, and ownership of an item or collection. (SAA Dictionary: https://dictionary.archivists.org/entry/provenance.html)

Tamper-evident: Device or process that makes unauthorized access to the protected object easily detected. Seals, markings, or other techniques may be tamper indicating. (Wikipedia: https://en.wikipedia.org/wiki/Tamper-evident_technology)

Tamper-resistant: Indicator or barrier to entry which, if breached or missing, can reasonably be expected to provide visible evidence to consumers that tampering has occurred. (21 CFR 700.25 - https://www.ecfr.gov/current/title-21/section-700.25)

V&V: Validation and verification

Validation: Process of providing evidence that the system, software, or hardware and its associated products satisfy requirements allocated to it at the end of each life cycle activity, solve the right problem (e.g., correctly model physical laws, implement business rules, and use the proper system assumptions), and satisfy intended use and user needs. (IEEE 1012-2024: https://doi.org/10.1109/ieeestd.2025.11134780). See also verification as these efforts are often linked together.

Verification: Process of evaluating a system or component to determine whether the products of a given development phase satisfy the conditions imposed at the start of that phase. (IEEE 1012-2024: https://doi.org/10.1109/ieeestd.2025.11134780). See also validation as these efforts are often linked together.

Introduction

This document outlines a framework of progressive criteria for evaluating and documenting content authenticity and provenance (CAP), especially for collections content impacted by artificial intelligence (AI). In the libraries, archives and museum (LAM) context, content authenticity is widely defined as “the quality of being genuine, not a counterfeit, and free from tampering, and is typically inferred from internal and external evidence, including its physical characteristics, structure, content, and context” and provenance as “information regarding the origins, custody, and ownership of an item or collection.”

The importance of documenting CAP across the stewardship lifecycle of collections content, especially when there are intersections with AI, is a growing concern. AI, especially generative AI, is fundamentally reshaping how collections are created, modified, described and reused. The abundance of AI-assisted, AI-modified and AI-generated materials has a side effect for the users, often without access or awareness to provenance data, increasingly seeing collections materials with increasing doubt or uncertainty. The white paper “Content Authenticity and Provenance in the Age of Artificial Intelligence” describes the issues of this potential trust erosion: “fast-changing AI technologies can introduce extraordinary opportunities, but also serious risks. AI may automate some workflows, but it can also challenge long-standing methods for documenting provenance, verifying authenticity, and sustaining public trust across the digital preservation lifecycle. In the face of rapid changes in technology and user expectations, LAMs must address how to ensure that digital collections’ content remains authentic, transparent, and verifiable from creation through access in order to meet the central community’s mission of public trust.”

Inspired by the National Digital Stewardship Alliance (NDSA) Levels of Digital Preservation, this framework takes an organizational perspective rather than a strictly technology or implementation perspective. It includes technical examples but the underlying structure asks how an organization decides what it needs, incorporates those needs into operations, establishes consistent requirements and demonstrates that those needs are being met.

By breaking down the complex aspects of CAP into a matrix of six evaluation criteria with actions described along four levels that provide increasing assurance of CAP, the aim of this work is to help guide LAM institutions to better plan for, understand, evaluate, and improve their work in accordance with long-standing archival CAP principles.

There is no one perfect way to meet these goals and much will depend on factors such as institutional readiness, capacity and flexibility. Organizations should consider what is fit for purpose. For some organizations, alliance with industry-based options such as the Coalition for Content Provenance and Authenticity’s (C2PA) Content Credentials or perhaps format specific options like JPEG Trust (ISO/IEC 21617) for photos and videos. But other LAM organizations without the need for external trust centers and digital signatures may want to investigate options to expand their current workflows with established structures such as PREMIS (see Appendix: Conceptual crosswalk of PREMIS to C2PA). These discussions and decisions may also need to take legal obligations and requirements for content made accessible on the web into account. Important rulings to be aware of include Article 50: Transparency Obligations for Providers and Deployers of Certain AI Systems | EU Artificial Intelligence Act and the California AI Transparency Act (CAITA).

The authors of this work, a collaborative group with representatives from a cross section of American federal agencies and cultural heritage organizations, welcome feedback and discussion. Send comments on this draft to Kate Murray, FADGI AudioVisual Working Group Leader (kmur@loc.gov) by October 15, 2026.

Related FADGI Project: CAP Environmental Review and Registry for Formats, Tools and Mechanisms

For point-in-time technical capacity for how CAP is addressed in a variety of file formats, open source tools and mechanisms, see the related FADGI-sponsored CAP environmental review and registry on GitHub. Note that this work is more tightly scoped at this time to audiovisual content but may expand in the future.

Defining the Levels

The four levels describe the development of institutional capacity, not a progression toward increasingly sophisticated technology.

Level 1 — Plan: “We know what we need.”
The institution has identified the CAP issue, its requirements, stakeholders, risks, lifecycle points, and intended outcomes. Implementation may still be partial or exploratory.

Level 2 — Structure: “We have defined how we will do it.”
The institution has translated planning into repeatable structures, workflows, responsibilities, data elements, and system relationships. CAP is moving from an individual or ad hoc activity toward an institutional practice.

Level 3 — Standardize: “We require it to be done consistently.”
The institution has defined the standards, formats, profiles, controlled vocabularies, bindings, or implementation rules that should be used. The emphasis shifts from simply capturing CAP information to making it predictable, interoperable, and consistently implemented.

Level 4 — Validate and Verify: “We can demonstrate that it works.”
The institution can test whether its requirements have actually been met. Validation and verification (V&V) processes, activities, and tasks are repeatable, results are documented, failures or indeterminate results are handled, and the institution periodically reassesses whether its methods remain effective and sustainable. Level 4 is an evidence-based response to these four questions that prove the requirements from Levels 1-3 have moved from policy into practice:

1. What claim or requirement are we testing?
2. What evidence demonstrates whether it has been met?
3. How do we evaluate that evidence?
4. What happens when V&V efforts fail, are incomplete, or produce an indeterminate result?

It’s important to note that Level 4 does not necessarily require external verification. Institutions should implement independent trust services, cryptographic signatures, C2PA trust lists, watermark detectors, and similar mechanisms which can strengthen evidence where appropriate, but the specific validation and verification methods should be in proportion to the claim being made, the risks involved, institutional requirements, and available resources. This aligns with the sustainability and “how much is enough” language in the Metadata Binding and Granularity criterion.

Tiered Community Recommendations for Content Authenticity and Provenance, Especially for AI-impacted Collections Content

Evaluation Criteria	Level 1: Plan	Level 2: Structure	Level 3: Standardize	Level 4: Validate and Verify
Awareness and engagement This criterion documents a range of awareness about CAP data, from documenting/detecting the existence of CAP data in any form at the lower levels to requiring specific types of data (checksums, assertions, logs) to standardized data structures (C2PA manifests) and normalizing non-standard data to standard types at the higher levels. For more information and resources: • Fixity and checksums - Digital Preservation Handbook - Digital Preservation Coalition • Carbon Footprint Toolkit - Digital Preservation Coalition	**Recognize CAP as information that requires active stewardship.** Identify where authenticity and provenance information may enter, be created, modified, retained, or exposed during the content lifecycle and stewardship process. Develop plans for recognizing, storing, managing, and providing access to CAP data. Planning considers staffing, systems, budgets, environmental impacts, legal requirements, policies for file formats and other types of process logging and resource tracking as well as needs for downstream users.	**Define what CAP data the institution expects to capture and maintain.** Establish institutional criteria for the types of authenticity and provenance information needed at identified lifecycle stages. Define responsibility for creating, receiving, retaining, and managing this information.	**Require defined structures, formats, and implementation practices.** Move from accepting CAP information in arbitrary forms to requiring documented standards, schemas, profiles, controlled vocabularies, or institutional templates. Establish procedures for normalizing non-standard CAP information when appropriate. Consider if V&V options need to be local-first/offline-first or cloud-based. Cloud-based dependencies may have an impact on privacy, environmental impact, latency, speed and trust.	**Demonstrate that institutional CAP requirements are being met.** Implement repeatable processes for evaluating whether required CAP information is present, complete, structured as expected, and usable at relevant lifecycle points. Retain evidence of successful, failed, incomplete, and indeterminate results and establish appropriate follow-up actions. Openness about V&V processes sets the expectation to build and sustain trust and confidence. Examples: Repository audit results; schema-validation reports; QC reports; records showing required PREMIS events were created; C2PA manifest validation; reports identifying missing required CAP elements; documented remediation.

Evaluation Criteria	Level 1: Plan	Level 2: Structure	Level 3: Standardize	Level 4: Validate and Verify
Tamper evident to tamper resistant/proof This criterion explores the progressive arc addressing how resilient the CAP data is against change (inadvertent or purposeful) and the ability to identify change when it does occur. For more information and resources: • Examples of tampering • C2PA Implementation Guidance: Guidance on the use of Content Bindings	**Identify where alteration of CAP data presents a risk and determine appropriate protections.** Define acceptable tamper-evident mechanisms and identify lifecycle points where alteration should be detectable. Document what constitutes an expected versus unexpected change and what evidence should be retained when changes occur.	**Integrate tamper detection into routine stewardship workflows.** CAP information is checked and preserved as part of acquisition, ingest, processing, preservation, migration, and/or access workflows rather than through ad hoc procedures. Define how detected changes are recorded, reviewed, and resolved.	**Require standardized mechanisms that strengthen the relationship between content and its CAP information.** Use defined tamper-resistant or cryptographically bound mechanisms where warranted—for example, C2PA hard bindings, signed manifests, or JPEG Trust mechanisms. Document which mechanisms are required for particular content or workflows and how they are preserved through transformations.	**Demonstrate that selected integrity mechanisms perform their intended function.** Conduct repeatable checks to determine whether unauthorized or unexpected changes can be detected and retain evidence of the results. Evidence may include checksum comparisons, binding and signature validation, trust-chain status, or controlled tests showing that alteration is detected. Where signatures or trust mechanisms are used, verify relevant signatures, bindings, certificates, or trust anchors. Establish procedures for failed, expired, unavailable, or indeterminate verification and validation results. Examples: Successful checksum comparison; C2PA signature/binding validation; documented signature validation; evidence of a deliberately altered test object being detected; failed-validation records; documented trust-chain status.

Evaluation Criteria	Level 1: Plan	Level 2: Structure	Level 3: Standardize	Level 4: Validate and Verify
Metadata binding and granularity This criterion considers the relationship of the CAP data to the content itself, including exploring what the CAP is describing at what level of detail (such as partial file, whole file, package level). More mature implementations/higher levels explore specific requirements for association including requiring operational relationships that support digital preservation activities.	**Determine what the CAP information describes and how closely it needs to be associated with the content.** Identify whether provenance or authenticity information needs to apply to a subpart of a file, an entire object, component, region, derivative, package, or collection-level entity. Balance desired granularity against legal, preservation, technical, operational, and sustainability requirements.	**Define explicit relationships between CAP data and the entities it describes.** Establish institutional structures for associating CAP information with files, components, packages, derivatives, agents, events, or other entities. Ensure these relationships can be maintained through institutional workflows.	**Require consistent, operationally meaningful bindings.** Standardize how CAP information is associated with content so relationships remain interpretable during preservation actions such as migration, transformation, replication, and audit. Define what happens to bindings when content changes or derivatives are created.	**Demonstrate that CAP information remains correctly associated with the intended content.** Periodically audit and verify relationships at the required level of granularity, particularly following migration, transformation, repackaging, or metadata extraction. Evidence may include manifest-to-asset checks, package inventory comparisons, derivative-to-source consistency, component-level association tests, and preservation audit records. Retain evidence of successful and failed checks and establish remediation procedures for broken, ambiguous, or lost associations. Examples include: Audit showing a manifest still references the correct asset; verification that component-level provenance remains associated with the correct component; package inventory comparison; derivative-to-source relationship checks; audit records after migration.

Evaluation Criteria	Level 1: Plan	Level 2: Structure	Level 3: Standardize	Level 4: Validate and Verify
AI attribution This criterion records the specific features in the CAP to document AI interactions with the content along the lifecycle of the digital object. Practically, this arc progresses from repurposing existing, more general components to towards purpose built and structured options for documenting AI interactions.	**Identify where AI may interact with content and determine what should be documented.** Plan for AI attribution across creation, acquisition, curation, digitization, preservation, transformation, and access workflows. Access workflows might include requirements of downstream stakeholders (such as streaming platforms or other distributors) who may require or include specific forms of CAP for assets. Define when recording AI involvement is required and whether distinctions such as AI-assisted, AI-modified, and AI-generated are useful. Consider documenting human review and, where meaningful, explicitly documenting that no AI interaction occurred.	**Incorporate AI attribution into existing provenance workflows.** Adapt current CAP mechanisms so AI-related actions can be recorded consistently. At minimum, capture enough information to identify that AI was involved, the relevant action or process, and the responsible human and/or software agent. For example, PREMIS events can identify AI-enabled software/models through linkingAgentIdentifier and describe outcomes through eventDetailInformation and eventOutcomeInformation.	**Use purpose-built, structured mechanisms for AI attribution.** Require standardized fields or assertions capable of documenting the nature and context of AI interaction. Where appropriate, capture the AI system/model, version, action, intent, source/input relationship, output, human involvement, and resulting transformation. Examples include IPTC AI metadata (`AI System Used`; AI System Version Used) and C2PA Content Credentials `c2pa.ai-disclosure`, `c2pa.actions`, and `digitalSourceType`.	**Demonstrate that required AI attribution is present, structured, and supported by available evidence.** Validate AI attribution against defined institutional requirements, including completeness, appropriate structure, and association with the relevant content or action. Where supported, validate signatures, credentials, provenance relationships, or other evidence supporting the attribution. Clearly distinguish between verified information, assertions from an identified source, and information that cannot be independently verified. Retain validation results and document incomplete, conflicting, or indeterminate AI attribution. Examples: Validation that required `c2pa.ai-disclosure` fields exist; PREMIS event checks; confirmation that required model/system identifiers are populated; schema validation; signature validation; QC reports identifying incomplete AI attribution; documented human review.

Evaluation Criteria	Level 1: Plan	Level 2: Structure	Level 3: Standardize	Level 4: Validate and Verify
Substantiation This criterion explores the ability of the CAP data and policies supporting CAP data to be available, transparent and verifiable at points of need including public access. For more information and resources, ● Data mining restrictions	**Determine what CAP information users need in order to understand and assess collection content.** Document organizational and user requirements for disclosure, including when CAP information should be public and when legitimate restrictions—such as privacy, embargo, quarantine, security, or rights considerations—justify limiting access.	**Define processes for extracting, transforming, and disseminating appropriate CAP data while applying access needs and disclosure restrictions.** Maintain relationships to underlying source information. Document how CAP information has been selected, transformed, redacted or otherwise remediated so that the decisions remain transparent and accountable.	**Establish institutional DIP schemas or equivalent dissemination profiles specifying required and optional CAP information.** There may be multiple profiles which vary by audience, collection, content type, purpose or other requirements.	**Ensure that CAP DIPs or equivalent dissemination representations are accessible to users and provides sufficient content for users to assess and confirm the CAP claims.** Document validation results, exceptions, failures, and determine processes for remediation and reassessment. Evidence may include signed assertions, verifiable credentials, provenance records, repository documentation, or other trusted records. Examples: Signed assertions; credentials that can be validated; provenance-chain review; public-facing validation and verification results; evidence linking claims to supporting records; documentation identifying which claims are validated, asserted but unverified, or unsupported.

Evaluation Criteria	Level 1: Plan	Level 2: Structure	Level 3: Standardize	Level 4: Validate and Verify
System interoperability and sustainability This criterion identifies reliance on specific technical support, tools and standards/specification for CAP data. For more information and resources, • Sustainability and risk management	**Plan for CAP information to remain usable beyond the system in which it was created.** Identify dependencies on particular tools, vendors, formats, standards, and services. Establish precedence rules when equivalent CAP information exists in multiple locations—for example, embedded metadata, repository databases, and sidecar manifests—and determine which representation is authoritative when versions diverge.	**Represent and expose CAP information in consistent, accessible ways.** Structure data and interfaces so CAP information can be located, interpreted, exchanged, and used across relevant institutional systems. Define mappings or synchronization mechanisms where multiple systems maintain related information.	**Require open, documented, and exportable representations.** Maintain CAP information in formats and structures that can be extracted from current systems and transferred to replacement systems without losing essential meaning or relationships. Minimize unnecessary dependence on proprietary implementations and document transformations between representations.	**Demonstrate continued interoperability, reliability, and sustainability through testing.** Periodically test whether CAP information remains interpretable, exportable, transferable, and usable for its intended authenticity and provenance purposes. Evidence may include successful export/import or migration tests, parsing by independent software, before-and-after validation, synchronization testing, and verification that essential CAP properties and relationships survive migration. Document dependencies and failures and reassess implementations as systems, standards, external services, and trust infrastructure change. Examples: Successful export/import tests; migration tests; validation before and after system migration; documented round-trip testing; successful parsing by independent software; dependency audits; preservation-watch reports; verification that essential CAP properties survived migration.

Additional Notes

Examples of purposeful or accidental CAP data tampering

- Stripping metadata: The most common form of tampering is the complete removal of CAP metadata, including for example C2PA manifests, from an asset, often through social media re-encoding or taking screenshots. This breaks the "chain of custody," signaling a lack of provenance but not necessarily indicating the content itself is fake.
- Modifying internal data: Any change to the metadata within a metadata manifest (e.g., altering a "CreateDate" field) or the asset's pixels will break the cryptographic hash.
- Spoofing with Stolen Keys: If a signing key is compromised, an attacker can generate "valid" but misleading provenance for any file. These will pass standard validation until the compromised certificate is revoked.

Data mining restrictions

The concepts in Substantiation are focused on data and policies supporting CAP data to be available, transparent and able to be validated. Adjacent to this work but not directly in scope are documenting permissions and restrictions for data mining of collections content by Large Language Models (LLM), Machine Learning (ML) models and other data harvesting and training efforts. This is a fast changing field with expectations, requirements and technology quickly changing but there are several documented options in industry standard mechanisms including purpose built options in the IPTC Photo Metadata Standard 2025.1 - 11.11. Data Mining and Guidance for Artificial Intelligence and Machine Learning :: C2PA Specifications, specifically Training and Data Mining Assertion from the Creator Assertions Working Group.

Sustainability and risk management

The rapid shifts in technology development and ongoing standardization efforts can lead to high ambiguity. It’s hard to make and stick to a plan when the surrounding landscape is so volatile. Practically, this translates into vulnerabilities for sustainable CAP practices. Institutions must weigh the risks involved in committing to a CAP plan and identify escape paths for their CAP data. For example, if a specific CAP method is no longer viable, what is the plan to migrate that data to another form and infrastructure? How can reliance on external dependencies such as a trust center operated by a third party be managed? For each decision point, organizations need to identify, analyze, and control financial, legal, strategic and security risks and have a plan for adjudication. It's essential to be aware of succession planning for data, systems and policies.

Appendix: Conceptual Crosswalk of PREMIS to C2PA

NOTE: This is a draft conceptual crosswalk mapping of how CAP data might be stored in PREMIS and C2PA. This information should be considered an illustration of the concepts and not a final rule. The authors of this work are consulting with the PREMIS Editorial Committee on a more formal outcome.

Preservation/provenance concept	PREMIS	C2PA/Content Credentials
What happened?	`eventType / eventDetail`	`action / description`
Specific operation	Example for adjusted color in a TIFF file: `eventType` [value = image modification]	Example for adjusted color in a TIFF file: `c2pa.adjustedColor`
When?	`eventDateTime`	`dateTime`
AI software	`linkingAgentIdentifier`	`softwareAgent`
AI involvement	`eventDetail`	`digitalSourceType`
Human oversight	`linkingAgentRole` [value = human reviewer] + `eventOutcomeDetail`	description and potentially AI disclosure
Source file (if modified from other parent file)	linked Object with role `source`	ingredient / prior manifest
Modified file	linked Object with role `outcome`	asset carrying the new manifest; `parentOf` ingredient relationship
Result	`eventOutcome`	action represented in successfully generated manifest
Integrity	PREMIS Object fixity	C2PA content binding + signed claim
Detailed AI/model information	Agent/Event metadata or extension	AI disclosure assertion

Resources

Camera & Imaging Products Association. *Exchangeable Image File Format for Digital Still Cameras*. CIPA DC-008-Translation-2026, 2026. https://www.cipa.jp/std/documents/download_e.html?CIPA_DC-008-2026-E. Accessed 12 Aug 2026.

Coalition for Content Provenance and Authenticity. “Trust List and Conformance.” *C2PA*, https://c2pa.org/conformance/. Accessed 12 Aug 2026.

Coalition for Content Provenance and Authenticity. *C2PA Implementation Guidance*. Version 2.4, C2PA, https://spec.c2pa.org/specifications/specifications/2.4/guidance/Guidance.html. Accessed 17 Aug 2026.

Coalition for Content Provenance and Authenticity. *C2PA Security Considerations*. Version 2.4, C2PA, https://spec.c2pa.org/specifications/specifications/2.4/security/Security_Considerations.html. Accessed 17 Aug 2026.

Coalition for Content Provenance and Authenticity. *C2PA Technical Specification*. Version 2.4, C2PA, https://spec.c2pa.org/specifications/specifications/2.4/specs/C2PA_Specification.html. Accessed 12 Aug 2026.

Coalition for Content Provenance and Authenticity. *Conformance Public: Trust List*. *GitHub*, https://github.com/c2pa-org/conformance-public/tree/main/trust-list. Accessed 12 Aug 2026.

Creator Assertions Working Group. *Creator Assertions Working Group (CAWG)*. https://cawg.io/. Accessed 12 Aug 2026.

Digital Preservation Coalition. “Fixity and Checksums.” *Digital Preservation Handbook*, Digital Preservation Coalition, www.dpconline.org/handbook/technical-solutions-and-tools/fixity-and-checksums. Accessed 17 Aug 2026.

Digital Preservation Coalition. *Carbon Footprint Toolkit*. Draft for comment, Mar. 2026, https://www.dpconline.org/digipres/implement-digipres/carbon-footprint-toolkit. Accessed 12 Aug 2026.

Electronic Code of Federal Regulations. “21 CFR § 700.25—Tamper-Resistant Packaging Requirements for Cosmetic Products.” *eCFR*, U.S. Government Publishing Office, www.ecfr.gov/current/title-21/section-700.25. Accessed 17 Aug 2026.

European Broadcasting Union. *Format for the Coding History Field in Broadcast Wave Format Files, BWF*. EBU Technical Recommendation R98-1999, 1999, https://tech.ebu.ch/docs/r/r098.pdf. Accessed 12 Aug 2026.

Federal Agencies Digital Guidelines Initiative. *Guidelines for Embedding Metadata in Broadcast WAVE Files*. Version 3, 26 Apr. 2021, https://www.digitizationguidelines.gov/audio-visual/documents/BWF_Embed_Guideline_v3_2021.pdf. Accessed 12 Aug 2026.

Google DeepMind. “SynthID.” *Google DeepMind*, https://deepmind.google/models/synthid/. Accessed 12 Aug 2026.

IEEE Standard for System, Software, and Hardware Verification and Validation. IEEE, https://doi.org/10.1109/ieeestd.2025.11134780. Accessed 20 Aug 2026.

International Organization for Standardization and International Electrotechnical Commission. *ISO/IEC PRF 21617-1: JPEG Trust—Part 1: Core Foundation*. 2nd ed., https://www.iso.org/standard/91405.html. Accessed 12 Aug 2026.

International Organization for Standardization and International Electrotechnical Commission. *ISO/IEC DIS 21617-2: JPEG Trust—Part 2: Trust Profiles and Reports*. 1st ed., https://www.iso.org/standard/89038.html. Accessed 12 Aug 2026.

International Organization for Standardization and International Electrotechnical Commission. *ISO/IEC CD 21617-3: JPEG Trust—Part 3: Media Asset Watermarking*. 1st ed., https://www.iso.org/standard/90209.html. Accessed 12 Aug 2026.

International Organization for Standardization and International Electrotechnical Commission. *ISO/IEC CD 21617-4: JPEG Trust—Part 4: Reference Software*. 1st ed., https://www.iso.org/standard/93391.html. Accessed 12 Aug 2026.

International Press Telecommunications Council. *IPTC Photo Metadata Standard*. Version 2025.1, IPTC, https://iptc.org/std/photometadata/specification/IPTC-PhotoMetadata-2025.1.html. Accessed 12 Aug 2026. Accessed 12 Aug 2026.

Internet Engineering Task Force. *The BagIt File Packaging Format (V1.0).* RFC 8493, Oct. 2018, https://datatracker.ietf.org/doc/rfc8493/. Accessed 12 Aug 2026.

Joint Photographic Experts Group. "JPEG Trust." *JPEG*, https://jpeg.org/jpegtrust/index.html. Accessed 12 Aug 2026.

Lavoie, Brian. *The Open Archival Information System (OAIS) Reference Model: Introductory Guide.* 2nd ed., Digital Preservation Coalition, Oct. 2014. *DPC Technology Watch Report,* no. 14-02. https://dx.doi.org/10.7207/twr14-02.

Library of Congress. "Exchangeable Image File Format (Exif), Version 3.0." *Sustainability of Digital Formats: Planning for Library of Congress Collections*, https://www.loc.gov/preservation/digital/formats/fdd/fdd000618.shtml. Accessed 12 Aug 2026.

Library of Congress. "Hashing Out Digital Trust." *The Signal*, 4 Nov. 2011, blogs.loc.gov/thesignal/2011/11/hashing-out-digital-trust/. Accessed 17 Aug 2026.

Library of Congress. "PREMIS: Preservation Metadata Maintenance Activity." *Library of Congress*, https://www.loc.gov/standards/premis/. Accessed 12 Aug 2026.

National Digital Stewardship Alliance. "Levels of Digital Preservation." *NDSA*, www.ndsa.org/publications/levels-of-digital-preservation/. Accessed 17 Aug 2026.

National Digital Stewardship Alliance. *Checking Your Digital Content: How, What and When to Check Fixity? Reasons to Collect, Check, Maintain, and Verify Fixity Information*. 2014, PDF via the Library of Congress. Accessed 17 Aug 2026.

National Digital Stewardship Alliance. *Checking Your Digital Content: What Is Fixity, and When Should I Be Checking It?* 2014, NDSA Fixity Guidance Report. Accessed 17 Aug 2026.

National Institute of Standards and Technology. *Computer Security Resource Center Glossary*. NIST, csrc.nist.gov/Glossary. Accessed 17 Aug 2026.

National Institute of Standards and Technology. *Secure Hash Standard (SHS).* FIPS PUB 180-4, Aug 2015, doi:10.6028/NIST.FIPS.180-4.

Rice, Dave. *CAP Registries (on behalf of FADGI)*. MediaArea GitHub, 2026 https://github.com/MediaArea/cap-environmental-review. Accessed 08 Sept 2026.

Society of American Archivists. *Dictionary of Archives Terminology*. https://dictionary.archivists.org/index.html. Accessed 17 Aug 2026.

Sternfeld, Josh, and Kate Murray. *Content Authenticity and Provenance in the Age of Artificial Intelligence: A Call-to-Action for the LAMs Community*. 2026. doi: 10.13140/RG.2.2.22223.96163. Accessed 20 Aug 2026.

"Tamper-Evident Technology." *Wikipedia*, Wikimedia Foundation, en.wikipedia.org/wiki/Tamper-evident_technology. Accessed 17 Aug 2026.

WITNESS. *C2PA Content Credentials and the Surveillance Risk: Adversarial Scenarios and Governance Gaps in the Content Provenance Ecosystem*. WITNESS Library. 8 July 2026. https://library.witness.org/product/c2pa-privacy/. Accessed 24 Aug 2026.

World Standards Cooperation. *Technical Report on AI and Multimedia Authenticity Standards: Mapping the Standardisation Landscape*. International Telecommunication Union, https://s41721.pcdn.co/wp-content/uploads/2021/10/AMAS-Technical-Report-on-AI-and-Multimedia-Authenticity-Standards-Final.pdf. Accessed 12 Aug 2026.