IPTC Frequently Asked Questions on C2PA Content Credentials

Created by the IPTC Media Provenance Advocacy Working Group

Version: 1.0
Last updated: 8 July 2026

Introduction

Media organisations around the world are exploring C2PA Content Credentials as a way to demonstrate the provenance of their content and to help audiences decide what to trust. This document answers the questions we at the IPTC hear most often from broadcasters, publishers and news agencies who are considering implementing the technology — from the basics of how it works, through practical implementation and cost questions, to security, privacy and what comes next.

This document was prepared by the IPTC Media Provenance Advocacy Working Group and will be updated as the ecosystem evolves. If you have a question that isn't answered here, please contact the IPTC — we would be happy to hear from you.

Table of Contents

Basics

What is C2PA?

The Coalition for Content Provenance and Authenticity (C2PA) is an industry-led organisation with members from media companies to software vendors to AI platforms to social media platforms and search engines.

The C2PA steering committee currently includes Adobe, the BBC, Amazon, Google, Meta, Microsoft, OpenAI, Publicis Groupe, Sony & Truepic.

These organisations, and over 400 other general and contributing members, are all working together to create an open, free technical standard that can be used by anyone to build Content Credentials into their tools and systems.

An organisation does not need to be a member of the C2PA consortium to use C2PA technology; you just need to use tools that support the standard.

What is the problem solved by C2PA Content Credentials?

Giving people more information about where content comes from, and how it was edited, helps them to decide whether or not they should trust it.

In an era in which trust has been attacked on all fronts—by false narratives, declining confidence in authority and the rise of generative AI—C2PA/Content Credentials provides a way to link a piece of content to its publisher, and potentially all the way back to its source. This provides transparency about how the content was made and a way to flag specific issues such as the use of generative AI.

What are "Content Credentials" exactly?

Content Credentials is the user-facing "brand name" used across the C2PA ecosystem to describe the verifiable metadata that can be inserted into a media file (or stored in a database) based on the C2PA's technical standards.

Content Credentials function as a "nutrition label" for digital content. They can contain information about who produced a piece of content, when it was produced, and which tools and editing processes were used.

What is CAWG?

The Creator Assertions Working Group (CAWG), operating under the Decentralized Identity Foundation (DIF), has created additional specifications that add extra information to the C2PA metadata section of a media asset.

One of these is the CAWG Identity Assertion, which builds upon the C2PA standard to allow creators and publishers to assert source and attribution of digital content while also supporting privacy and transparency.

Other additional specs are the CAWG Metadata Assertion, which allows publishers to add additional metadata to assets, such as rights, captions and accessibility information; and the CAWG Training and Data Mining Assertion, which is among the machine-readable opt-out solutions being assessed in the European Commission's process to identify agreed Text and Data Mining (TDM) opt-out protocols under the EU AI Act; a final list is expected in late 2026.

What's the difference between C2PA, CAI and Content Credentials?

These three names are often used interchangeably, but they refer to different things:

- C2PA (the Coalition for Content Provenance and Authenticity) is the standards body. It develops and publishes the open technical specification that defines how provenance metadata is created, signed and verified.
- CAI (the Content Authenticity Initiative) is a community founded by Adobe in 2019 to promote the adoption of provenance technology. It publishes open-source developer tools and SDKs that implement the C2PA specification, and its membership now numbers in the thousands, including many news organisations. Broadly: C2PA writes the standard, CAI promotes it and builds tools for it.
- Content Credentials is the user-facing name for the technology itself — the verifiable metadata attached to a piece of content based on the C2PA standard. It is represented by the "cr" pin icon. When you hear "Content Credentials", think of it as the product; C2PA is the specification underneath it.

Does C2PA detect deepfakes or AI-generated content?

No. C2PA Content Credentials is a way to attach secure metadata describing the origin of any piece of content, whether it was created by a human or a machine.

Software or services that read C2PA metadata from content do not perform any kind of AI check on the content. They simply extract the data that was put there by its creator or publisher, and ensure that the content and its accompanying metadata have not been tampered with since it was published.

As a broadcaster or publisher, why should I use Content Credentials?

There are three main use cases:

- Proving that content has not been modified since it was created or published. The core C2PA specification ensures that if signed content is modified in any way, the content's digital signature will no longer be valid. Detectors and validators can use this to signal to users if a media asset has been manipulated. (Of course, the tools themselves cannot tell whether the manipulation is malicious or innocent—only viewers and readers can tell that, but we can give them as much information as possible to make that decision.)
- Attaching provenance information to a media asset. By using the CAWG Identity Assertions in a C2PA manifest, information can be added that declares who published the asset, signed with the publisher's organisational certificate. Because the C2PA metadata travels with the content, even if that content is surfaced on a third-party site, it can be traced back to the organisation that published it.
- Attaching additional metadata to a media asset. Similarly, the CAWG Metadata Assertion allows publishers to add additional information to the signed metadata component of the asset. This could include information such as caption, source, date created, location, accessibility information such as alt text, copyright, declaration of the use of AI in the content production process, usage rights (including AI opt-out information).

As a social media platform, what's the benefit to me?

Platforms are where most people encounter news content, and where questions about authenticity are hardest to answer at scale. C2PA Content Credentials give platforms a machine-readable signal they can act on automatically, rather than relying solely on detection algorithms or user reports.

There are three practical benefits:

- **Labelling AI-generated content at scale.** Reading Content Credentials on upload lets a platform label AI-generated content automatically and reliably, based on a declaration made at the point of creation rather than an after-the-fact guess. TikTok already does this: it automatically labels content as AI-generated when uploaded files carry C2PA metadata indicating so. This also supports compliance with emerging transparency obligations such as those in the EU AI Act.
- **Surfacing provenance to users.** LinkedIn displays the "cr" pin on images and videos that carry valid Content Credentials, letting users see who signed a piece of content and how it was made. This gives audiences a reason to trust content on your platform rather than elsewhere.
- **Reducing the moderation burden.** Provenance signals help distinguish content from verified publishers from content of unknown origin, giving trust and safety teams an additional, cryptographically verifiable input to their decisions.

As provenance signals become common across the internet, we anticipate that audiences will come to expect them — and content without provenance information may increasingly be treated with suspicion.

Implementation

Who is using C2PA Content Credentials?

Many organisations are working on C2PA support but not all have launched yet. Here are some C2PA implementations that are currently available to the public:

- France Televisions signs its daily news broadcast videos using its C2PA certificate, which is on the Verified News Publishers List. France TV won an EBU Innovation Award for this work.
- **OpenAI**, **Google**, **Adobe** and **Microsoft** all embed C2PA Content Credentials in their AI-generated content. **Samsung** applies C2PA metadata to images edited with AI on the Galaxy S25 phone.
- **Leica** and **Nikon** have released high-end cameras with C2PA capabilities using older versions of C2PA. Canon has released EOS R1 and EOS R5 Mark II cameras with the "Authenticity Imaging System" specifically for news organisations. Sony has released both still image (ILCE-1 and Alpha-9 III) and video camcorder (PXW-Z300) products along with its "Camera Authenticity Solution."

- **Google** has added C2PA Content Credentials support to its Google Pixel range of phones. **Google** also uses C2PA to surface content "created with a camera" on YouTube and to highlight AI edits in Google Photos.
- **Adobe**'s Creative Cloud software suite supports C2PA Content Credentials, including Photoshop, Lightroom and Premiere, so images and video edited with AI have a full provenance trail.
- Image editing and management software from **Camera Bits** (Photo Mechanic), **Fotoware** and others either already support C2PA metadata or are in the process of implementing support.
- **LinkedIn** shows Content Credentials when embedded in images and video files attached to posts.
- **TikTok** highlights AI-generated content by checking for C2PA Manifests on upload.
- For our own part, IPTC signs all images attached to blog posts on iptc.org using an approved certificate on the C2PA Trust List. To do this, IPTC underwent the C2PA Conformance Program. See later questions for more information on this.

What kinds of media can C2PA Content Credentials be used on?

Most early implementations worked with still images but increasingly C2PA Content Credentials can be found in other types of content. Now, C2PA supports almost any type of media: text, images, static video files and streaming video protocols, audio and music files, PDFs, ebooks, even fonts and AI model files.

Some organisations are using C2PA to sign video content: for example, France Télévisions is signing their daily news broadcast in the form of a single mp4 video file.

Specs have recently been published to handle live streaming, but no production implementations yet exist. EZDRM, Sony, Unified Streaming and others are working on support and have demonstrated successful proof-of-concept implementations.

How do I get started? What do I need?

We at the IPTC have prepared a **detailed guide for news outlets who want to get started** using C2PA Content Credentials in both checking content's origins and in signing your own content. The latest version is 1.1, released in June 2025. We are working on a new version to be released very soon. You can always find the current version of the guide at https://iptc.org/std/guidelines/media-provenance/

In summary, as a publisher you can choose to put your own publishing and signing tools through the C2PA Conformance Program, which allows you to sign content with your own C2PA certificate; alternatively, you can use commercially available software to sign content and add your publisher signature using a CAWG Identity Assertion.

Here is a list of commercial software currently available as C2PA Conformant generator products with APIs that can be integrated into your newsroom workflows. If you know of other tools that can be used by media publishers, please let us know; we would be happy to add new tools to the list.

- DigiCert Content Trust Manager
- Encypher platform API
- EZDRM C2PA Video Signature Service
- Lumid from Verify.tech
- Secure Content Engine from Digitality Consulting (Italy)
- Trufo Provenance Platform
- Vbrick Enterprise Video Platform (EVP)

As far as we are aware, only the Trufo Provenance Platform has built-in support for our recommended approach using signed CAWG Identity assertions.

Other providers, including SSL.com's "Enterprise Content Authenticity Solutions" and Noosphere Technologies, offer implementation support to media organisations but do not yet have a tool available via the C2PA Conformance Program.

How much does this cost? How much effort would it be to adopt in our newsroom?

As we explain above, there are two main routes to signing your content: putting your own tools through the Conformance Program, or integrating existing tools into your publishing workflow. Associated costs will depend on which route you choose. Whichever step you choose, most of the cost will be in the effort needed to integrate the signing system into your newsroom production workflow.

Obtaining a certificate that you can use to sign CAWG Identity Assertions ranges from hundreds to low thousands of dollars or euros per year, depending on which certificate authority you use. (The IPTC currently endorses Digicert, GlobalSign, SSL.com and Trufo as Certificate Authorities for Verified News Publisher List certificates used to sign Identity Assertions; see details below.)

What happens when my CDN or CMS strips metadata?

After your signing tools add a signed C2PA Manifest including a signed CAWG Identity Assertion to your assets, it is important that the metadata is not stripped. This will probably require some configuration of your Content Management Systems (CMSs) and Content Delivery Networks (CDNs). Most CDN tools can be configured not to strip metadata (or even specifically not to strip C2PA metadata).

If your CMS generates multiple versions of images for different screen sizes and devices, it is important that all variants of the image are signed.

If you cannot avoid metadata stripping (and of course you as a publisher have no control over what happens to an asset after it leaves your publishing system), then we recommend using "soft bindings" such as fingerprints or watermarks.

Publisher Identity

How is publisher identity handled in the C2PA ecosystem?

The core C2PA specification covers the "how" and "what" aspects of content provenance: the tools and systems that were used to create content, and the actions that were taken by tools in the production workflow (cropping, resizing, transcoding, editing etc). This also includes information on whether AI was used in the creation of the asset, via the Digital Source Type declaration.

The CAWG Identity Assertion extends C2PA to add the "who" aspect: it defines standardized assertions that allow creators to express individual and organizational identity, accountability, text & data mining, and intent.

The CAWG Organizational Identity Profile for Content Credentials sets out exactly what is required to support organisational identity in C2PA: specifically, it requires implementing the CAWG Identity Assertion, using X.509 certificates for signing entities, and also supporting the cawg.metadata assertion which allows for arbitrary metadata. The IPTC recommends using IPTC Photo Metadata properties in the cawg.metadata assertion to store publisher metadata. Our forthcoming implementation guide will contain specific details on which metadata properties can be used.

The C2PA has published the C2PA Recommendation on Human and Organisational Identity which officially recommends using the CAWG Organizational Identity Profile to convey organisational identity information. We expect this to be adopted by all C2PA implementations that need to understand the "who" aspect of content provenance.

Following these guidelines, the IPTC's recommendation to the news industry is to use a conformant tool to sign your content, and add a CAWG Identity Assertion that attaches your publisher identity via a publisher certificate.

How does IPTC's Verified News Publishers List fit into the C2PA ecosystem?

Since early versions of C2PA, the IPTC has maintained the Verified News Publishers List, which allowed us to maintain a list of entities that were verified as publishing journalistic news content (without making any judgement as to the quality or truth of that content). Users could check using IPTC's C2PA validation tool, IPTC Origin Verify, to see whether content was published by an organisation that has been verified as a news publisher.

Since that time, the C2PA trust ecosystem has changed to be more strict about who can sign content, and has introduced the C2PA Conformance Program. This ensures that only tools (hardware and software) can sign C2PA manifests. The role of asserting creator and publisher identity has moved to the CAWG Identity Assertion, as described above.

The CAWG Identity Assertion specification requires a second certificate to be used to sign CAWG Identity Assertions (containing publisher identity) and CAWG Metadata Assertions (containing editorial metadata). Those identity certificates need their own trust ecosystem, and the IPTC Verified News Publishers List has been designated by CAWG as one of the standard mechanisms that every compliant validator must support for verifying trusted identity certificates.

The diagram below shows how both the C2PA Claim and CAWG Identity Assertion work together: the C2PA Manifest is signed by the tool, and the CAWG Metadata is signed by the publisher via the Identity Assertion.

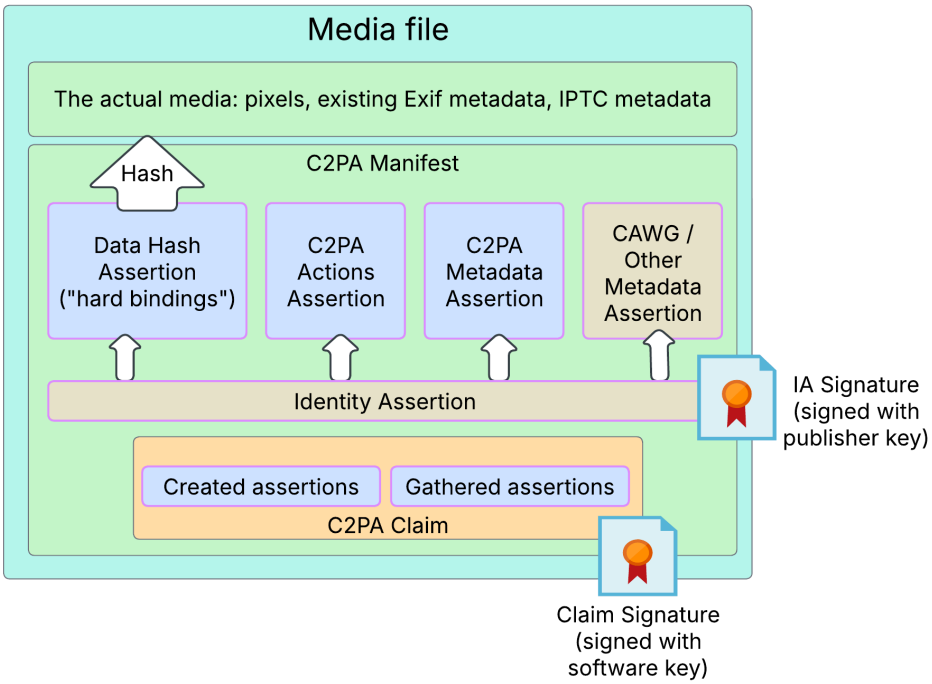

How do I get onto the IPTC Verified News Publishers List?

If you are a news publisher, you can fill out this short form on the IPTC website to begin the process of joining the IPTC Verified News Publishers List. You will need to complete a short questionnaire detailing your news output and company information. All applications are subject to approval by the IPTC, acting on behalf of the IPTC Media Provenance Committee.

Security

If a C2PA signal can be stripped or lost in processing, isn't that an issue?

C2PA metadata is removable by design. Content passes through many tools and platforms that legitimately remove metadata, so enforcing persistence might harm compatibility and user control. In the future we hope that the absence of C2PA metadata will be signal enough to indicate that content might not be trustworthy.

C2PA introduces the concept of "**soft bindings**" (as opposed to "hard bindings" which are mathematical hashes of the content added directly to the asset's C2PA manifest). Soft bindings can take the form of **fingerprints** (calculated from the asset and stored in

a registry) or **invisible watermarks** (identifiers embedded in the content). These can be implemented in such a way that they can survive simple content manipulation such as cropping, colour changes, rotation and screenshots.

By using fingerprints and watermarks, even if the C2PA metadata is stripped from the media file, it can be retrieved from a registry using the watermark.

Fingerprints and invisible watermarks use statistical analysis and all current algorithms have false-positive and false-negative issues. They will never give perfect protection against malicious actors, but they can address some cases.

Can C2PA metadata be hacked?

The core security technology underlying C2PA is public-key cryptography, which is the same technology which has been used for many years to protect internet banking, credit card transactions and more. There were some issues identified in early versions of the specification but they have been addressed in recent versions, for example by requiring secure timestamps to be added to the signed content.

The underlying technology has been reviewed by the National Institute of Standards and Technology and the "Five Eyes" intelligence agencies of USA, UK, Australia, New Zealand and Canada.

Additionally, in order to provide assurance that products adhere to the specification, and fulfill a set of security requirements which ensure they are producing and validating C2PA data correctly, the C2PA Trust List was set up to distinguish valid, known certificates from unknown ones. The Conformance program is responsible for vetting products before they appear on the official C2PA Trust List.

Can metadata be redacted? How does that happen?

Redaction is an important part of the C2PA specification: if we are to enable a full end-to-end C2PA provenance workflow for all content, it is imperative that we are able to redact information such as location and exact capture time which may be used by bad actors to violate the privacy or security of content creators.

As news organisations, we must protect our sources, our subjects and our staff.

The C2PA specification fully describes mechanisms for redacting sensitive images and metadata from the provenance chain. However, that part of the spec is not yet well implemented across all tools.

Therefore, until redaction technology is more widespread, IPTC recommends that news publishers implement a "clean slate stamping" technique, whereby at publish time, all C2PA metadata is stripped from content, and then a manifest is consciously added by the publisher, including only the metadata that the publisher is willing to attest, along with the publisher's Identity Assertion, asserting that the publisher stands by the content and the processes that were undertaken in the path to publishing.

What can be done to protect people at the point of acquisition whose data might be sensitive?

This is an issue we have considered carefully. There is no compulsion to enter metadata which might identify or otherwise endanger the provider of a piece of content.

As consultant Paul Reinitz (ex-Getty Images) explains:

> Here's the reality. You can't have both complete anonymity and verified authenticity. If content is truly anonymous, there's no way to prove where it came from.
>
> But privacy isn't impossible. There are workarounds. A creator can verify their identity to a trusted intermediary like a news organization, which then signs the content publicly. The creator stays protected while the content gets authenticated. This model works well, but it requires the intermediary to use certified, trustworthy tools.

This echoes the approach that we describe above: publishers should remove any existing C2PA manifests, sign content with a "clean slate stamping" mechanism using conformant signing tools, and add their own publisher certificate and publisher metadata via a CAWG Identity Assertion and CAWG Metadata Assertion.

Can I use watermarking alongside C2PA for extra security?

Yes, this is an option. Watermarking and C2PA can both be applied to the same piece of content as separate signals. In fact, the EU AI Act Code of Practice recommends this approach for AI-generated content.

If C2PA metadata is removed from a piece of content, if a copy of the manifest data is stored in an online database, you can use a watermark or a fingerprint to find it again. The C2PA specification refers to watermarking and content fingerprinting as soft bindings, and requires that they be generated using one of the approved watermarking and fingerprinting algorithms.

The Content Authenticity Initiative calls this approach "Durable Content Credentials," where C2PA is used alongside invisible watermarking and/or fingerprinting.

Google's Gemini AI model uses this approach, applying both C2PA metadata and Google's proprietary SynthID digital watermark to its generated images. Recently, OpenAI announced that it would also adopt SynthID alongside C2PA to label its AI-generated content.

One downside is that by necessity, watermarking systems are proprietary technologies: if a watermarking algorithm is published, it can easily be reverse-engineered. Therefore, it is not possible to publish an open standard specification for durable content watermarking. Watermarking systems also require a database of some kind to be used as a repository of watermarks and associated metadata, and these repositories can be costly to maintain. Therefore all watermarking systems are commercial solutions and we do not recommend any individual watermarking provider.

Adobe has published the TrustMark system as an open specification for watermarking. Being an open specification it can easily be reversed, so it can only be used to prevent accidental removal of metadata. It cannot prevent malicious actors from stripping metadata.

It's important to remember that fingerprinting and watermarking algorithms use statistical methods to embed identifiers into media files, and therefore we cannot say with complete certainty that watermarks and fingerprints will match. The technology will never be perfect but it is better than nothing. Users must be careful to look out for both false negatives (files with watermarks that do not match) and false positives (files that purport to have watermarks but do not actually include them). Storing a thumbnail of the original content in a watermark repository is usually a good way to verify that the watermark matches the correct media file.

Due to the technology's proprietary nature, there will always be a multitude of watermarking solutions in use. Currently the C2PA spec requires validators to check the API of each and every watermarking system registered on the "soft binding algorithm

list" to see whether the content contains a watermark from that vendor. To avoid this problem, a "signposting" approach may be used, whereby an open watermark such as TrustMark is used to embed a simple identifier on top of the proprietary watermark. This identifier would then allow validator tools to easily determine which proprietary watermarking system's API should be queried.

This signposting approach is described in the C2PA Implementation Guidance and documented in Adobe's TrustMark implementation; it has not yet been incorporated into the normative C2PA Technical Specification.

Do I need to use blockchain?

No. C2PA is not based on blockchain technology (although it shares some common technological roots, such as public-key cryptography).

Some watermark repositories are based on distributed ledger technology such as blockchains, but most watermarking systems are based on simple online databases.

Audiences and Consumers

As a consumer of content, what do I see?

Currently, content containing C2PA Content Credentials is not surfaced to users by default. If a user has a browser plugin installed (such as the Adobe Content Authenticity Chrome extension), then a "cr" pin is shown.

Some websites such as LinkedIn display a cr pin on images that contain valid C2PA Content Credentials. The 'cr' icon is overlaid on the top left corner of the image: see the screenshot to the right for an example.

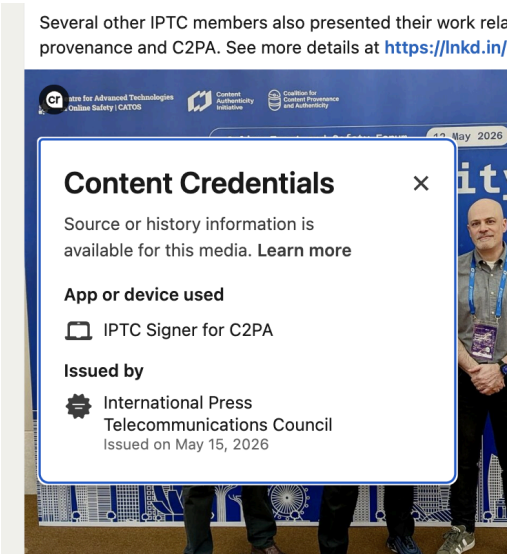

Clicking on the 'cr' pin shows more detail, including the tool used and the owner of the certificate that was used to sign the content.

Unfortunately, the image displayed by LinkedIn has the C2PA metadata stripped out, so users can't load the LinkedIn version of the image into a C2PA validator.

If a user loads the signed image directly into a validator such as IPTC's Origin Verify tool, the [full provenance information can be seen]:

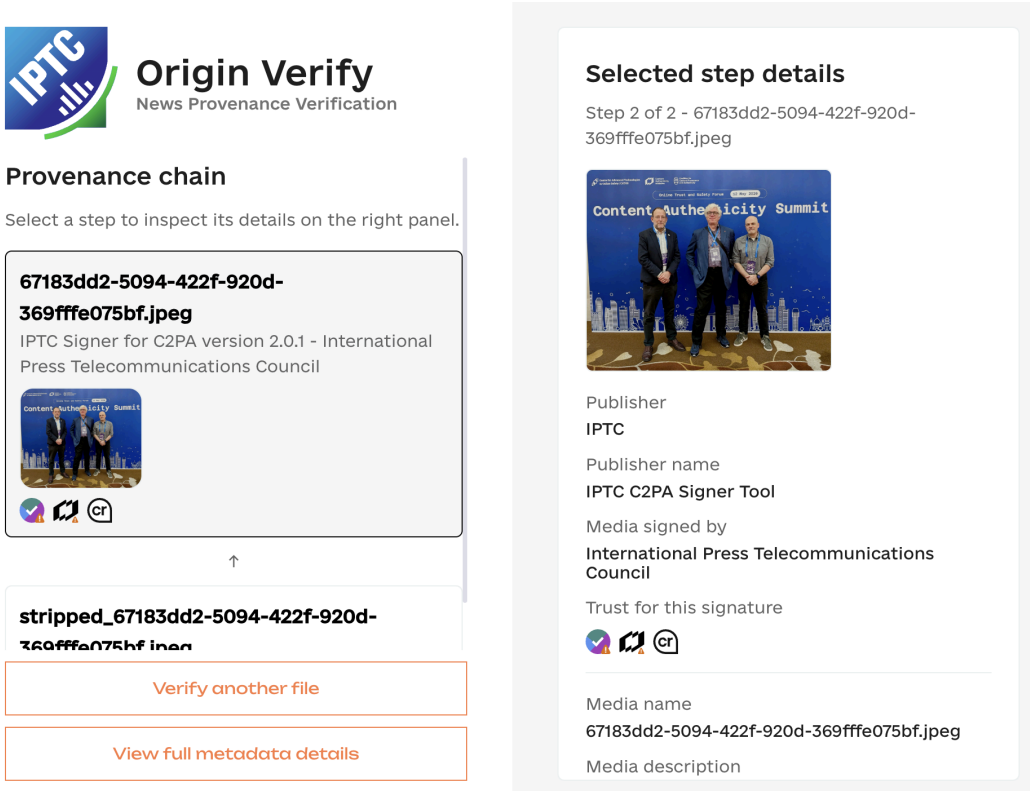

Some IPTC partners have visualised how C2PA and CAWG metadata could be surfaced to users of a news website. The following image shows how C2PA Content Credentials could be displayed on a publisher website:

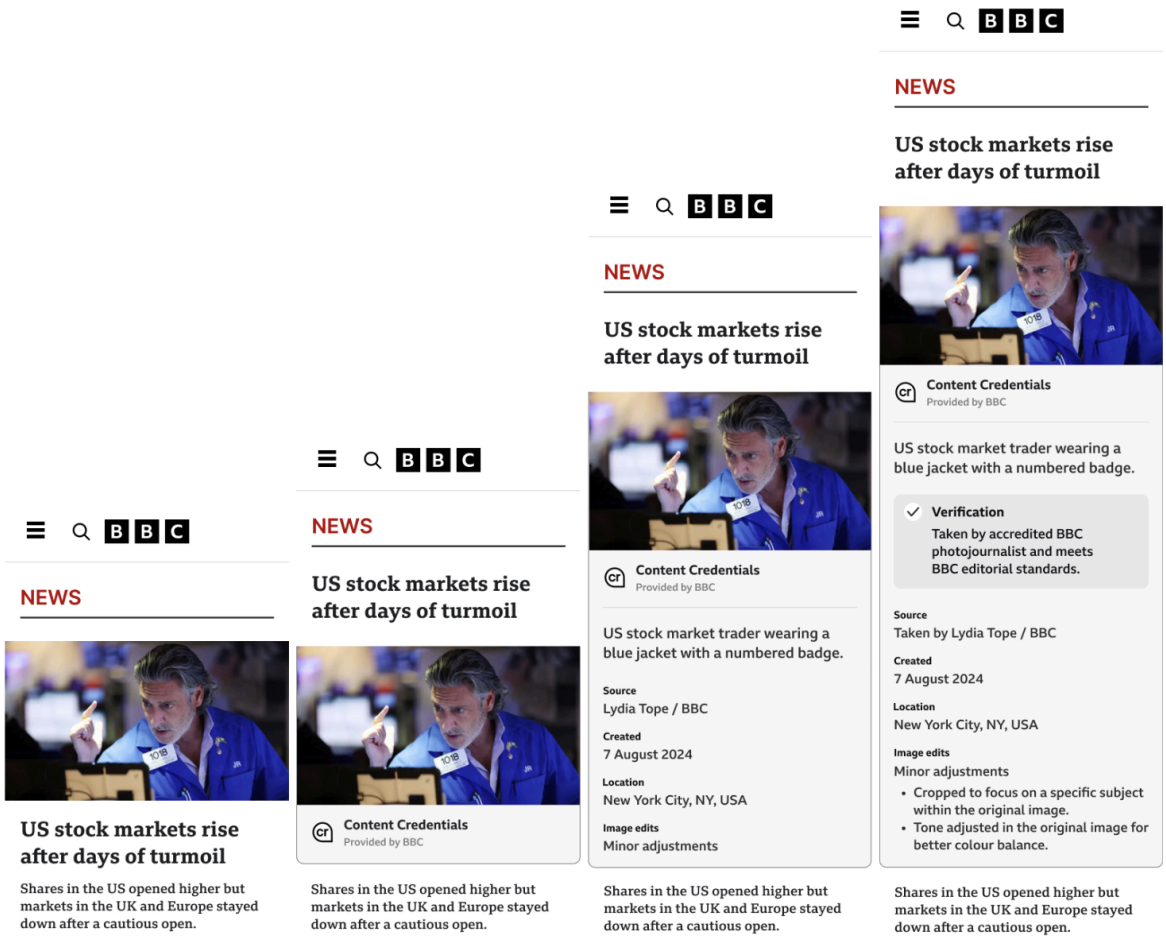

As a consumer of content, what do C2PA Content Credentials tell me, and what don't they tell me?

Content Credentials tell you how a media asset was made, who stands behind it, and whether it has been edited or tampered with since it was published.

They don't tell you whether the content is *true*: Content Credentials make no statement about the factual accuracy of the content itself. And they can't identify AI-generated content that carries no credentials: C2PA is not a deepfake detector.

Armed with the information conveyed by Content Credentials, you can make your own judgement on whether or not to trust a media asset.

(see also: Does C2PA detect deepfakes or AI-generated content?)

Next steps

What are the plans to develop C2PA further?

Since publishing the original C2PA specification, the Technical Working Group has been working on expanding the specification to apply to more media formats such as plain text, HTML, live streaming video and audio, 3D files, AI models and more. Most of these are now resolved.

With most of the technical work complete, work is now moving to the adoption phase, promoting C2PA Content Credentials across the tech ecosystem.

What are IPTC's plans to support media companies with implementing C2PA in their newsrooms?

The IPTC has already created the Verified News Publishers List and support materials such as these guidelines, reference implementations of signing software and best practices guidance on obtaining a certificate, going through the C2PA Conformance Program and more. Some of these materials are publicly available and some are available only to IPTC members.

The IPTC is organising or co-organising events to bring C2PA and media provenance awareness to newsrooms around the world. We have already held an event in Toronto, Canada and will soon be co-organising events in Bergen, Norway (September 2026) and Tokyo, Japan (October 2026).

Who can I talk to for further information?

IPTC members—media publishers and broadcasters, media software and hardware vendors, consultants and service providers—get together regularly to discuss topics related to C2PA implementation and media provenance generally.

The IPTC Media Provenance Committee has three Working Groups that work on specific issues:

- Governance Working Group: working on the governance of the Verified News Publishers List, how it works, how to maintain and develop the list in the future.

- Best Practices & Implementation Working Group: where members share their ongoing work to implement C2PA Content Credentials in their newsroom, and where we get together to decide best practices and discuss tools and services that can help media organisations with their C2PA implementation projects.
- Advocacy Working Group: working on materials to explain C2PA and IPTC's Media Provenance work to the public, including this FAQ.

If you would like to talk to us about your work or potentially joining IPTC, you can contact IPTC using this online form. You can learn more about joining IPTC on our Become a Member page and you can sign up to our Friends of IPTC Newsletter and Media Provenance Updates newsletter at our IPTC Newsletters page.

Links and Further References

- C2PA Technical specification, version 2.4
- C2PA Recommendation on Human and Organisational Identity
- CAWG Organizational Identity Profile
- CAWG Identity Assertion Specification
- CAWG Metadata Assertion Specification
- C2PA Provenance Labels Increase Trust in Digital News Platforms Across Western Countries - a study from Media Futures looking at user understanding of C2PA markings across 6,000 participants in the US, UK and Norway
- Adobe's open-source TrustMark watermarking implementation
- CAWG Technical Specifications
- Content Provenance & Integrity Knowledge Hub from CBC/Radio-Canada