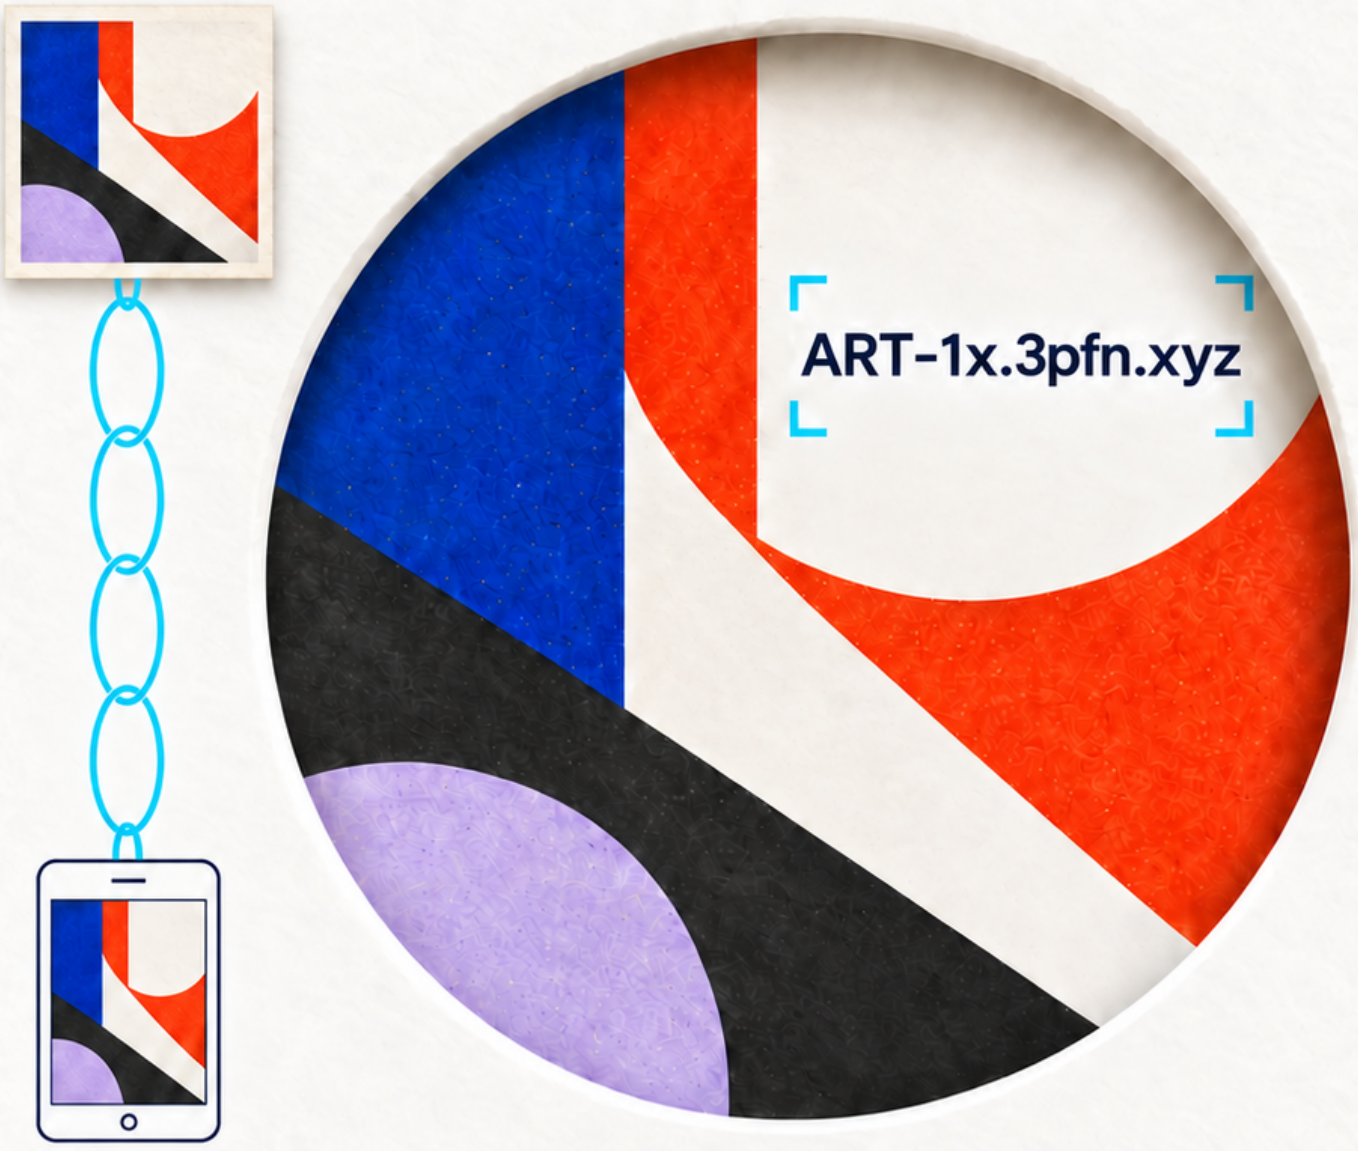
IPFN
International Pinned File Number

A New Publishing Protocol

A web-native identifier for creative work.



TECHNICAL WHITE PAPER / 3 NOVEMBER 2026

A New Publishing Protocol

A web-native identifier for creative work

Kilian Liptrot / Blinding Visuals Press and IPFN Agency

PROBLEM

Creative work travels farther than its creator's name.

SOLUTION

An IPFN lets anyone check a creator's claim web-wide.

ABSTRACT

IPFN gives registered creative work an identifier that works like a web address. A person encountering that identifier can find the recorded claimant and examine the evidence attached to the work.

The system joins controlled registration, an exact-file digital fingerprint, a human-readable certificate on IPFS, and a daily ledger committed to Base, an Ethereum layer 2 network. Many registrations share one blockchain transaction. This paper explains why a publishing business needed that arrangement, places it in the history of copyright and identifiers, and describes its technical operation. A worked photograph demonstrates the complete evidence path. The appendices specify the objects, bytes, states, and comparisons needed to reproduce the checks.

READING GUIDE

Publishing origins and history 3

How the file chain works 6

Two registered works 12

Technical specification and verification 14

1. We were trying to publish and sell art

You ever seen a photograph in the wild and wondered who took that image? You might see it on a wall, in a book, or reposted online. The image is there. The way back to the person who made it often is not.

That was the problem we ran into at Blinding Visuals Press. We were preparing photographs for publication, sale, and what we think of as printable galleries. We could identify the book. We could identify its publisher, edition, and format. But the individual visual works inside that book needed their own way of being identified, attributed, and checked wherever they appeared.

A photograph could be reproduced in a publication, sold as a print, or released digitally. The book's identifier described the publication. The photographer's name in a caption described an attribution. The file carried information that might be lost during an export. We needed those pieces to remain connected through a reference another person could actually use.

IPFN came out of that publishing problem. In ordinary language, it is like an ISBN for images: a registered identifier connects the work to an understood record. IPFN adds an evidence path behind that record. The person opening it can examine who made the claim, what file was registered, and how the certificate and ledger relate to the public blockchain commitment.

We started with the role of a notary

The starting point for FileChain was the role of a notary: keep an account that can later be examined when someone needs to establish what was presented, by whom, and when. We built a provenance query around that function and then made it web-native.

The registration records the claim. The digital fingerprint identifies the exact file. The certificate explains the assertion in language a person can read. The ledger connects the records in an ordered document. The blockchain supplies the timestamped witnessing function by recording that document's fingerprint outside our own editable database. The web address gives people a way to find the account and inspect its parts.

This makes the blockchain useful to a publisher, a photographer, a gallery, or a public institution. Someone can retain the evidence and repeat the comparisons without asking us to pronounce the result. The record answers who claimed what and when the evidence was committed; production declarations and other supporting material explain the claimant's account of how and why the work was made. The architecture separates those recorded assertions from independently measured facts.

European litigation already shows why that evidence can matter. In *AZ Factory v Valeria Moda*, the Tribunal judiciaire de Marseille, First Civil Chamber, 20 March 2025, RG 23/00046, relied on two reports documenting blockchain timestamps when establishing AZ Factory's economic copyright in clothing designs. The judgment also examined publication under the company's name. For IPFN, the practical lesson is clear: a dated commitment becomes useful evidence when it is connected to an identified work and an intelligible account of the claim. [H30]

We built the technology to solve a commercial problem and realized it could have a broader application across the creative world. The institutional proposition follows from that experience:

maintain a consistent way for creators to make an affirmative claim and for everyone else to inspect it.

2. How we arrived at this problem

Copyright and international recognition

Copyright developed around the reproduction and distribution of creative expression. Britain's Statute of Anne took effect in 1710. The first U.S. federal copyright statute followed in 1790, initially covering books, maps, and charts. As new media became important, the law expanded; photographs and photographic negatives expressly entered U.S. copyright legislation in 1865. [H01, H02]

The Berne Convention, adopted in 1886 and revised over the following century, addressed the international problem. A creator should not lose protection simply because a work crosses a border. Its present framework combines national treatment, minimum protections, and protection without formalities for the rights guaranteed outside the work's country of origin. Articles 5(1)–(3) distinguish that international protection from protection governed by domestic law in the country of origin. [H03]

The no-formalities principle serves creators. It also leaves a separate practical job to be done. Having copyright does not automatically give someone encountering a particular image a searchable route to its claimant. WIPO states that there is no searchable international registry of copyright-protected works. It also recognizes that voluntary registration systems can help document claims, resolve disputes, and support transactions. [H04, H05]

IPFN fits within that distinction. A creator can make an affirmative, inspectable claim while retaining the protection that the law supplies independently of registration. The institutional opportunity is to make attribution and its evidence easier to find, use, and challenge when necessary. Berne permits that voluntary work; it does not prohibit international institutions or voluntary registries.

ISBN solved a commercial identity problem

The International Standard Book Number began with the book trade trying to identify what it was buying, stocking, and selling. In 1965, W. H. Smith was preparing for a computerized warehouse and needed books to be identified consistently. The Standard Book Numbering system was devised in 1966 and implemented in 1967. International adoption followed, and ISBN became ISO 2108 in 1970. [H06, H07]

A practical answer to ordering and inventory became publishing infrastructure. An ISBN links a number to a particular publication and its publisher, edition, and format. Booksellers, distributors, libraries, and publishers can use that common reference without inventing a new description every time the book moves. The registration agencies give participants a consistent means of obtaining and using the numbers. [H08]

That history matters to IPFN. A useful institution can emerge from solving an ordinary commercial problem well. The next step is to administer the identifier consistently enough that other people can build useful activity around it. Our proposed authority comes from that work: clear assignment, intelligible records, inspectable evidence, and accountable administration.

Photography was reproducible before it became digital

Photography has long contained a distinction between the originating event and its representations. A daguerreotype produced a direct image on a plate. Talbot's negative-positive process, introduced publicly as the calotype in 1841, made multiple prints possible from a paper negative. Reproduction is therefore part of photography's early history. [H10, H11]

Digital photography changed the mechanics again. In 1975, Steven Sasson built a digital still-camera and playback system at Kodak. His firsthand account describes that development. A photograph could become encoded data that a machine could store, process, and transmit. [H12]

The web made the image easier to distribute independently of the page, caption, or account through which it first appeared. A recognizable image could survive while the route back to its creator disappeared. The practical problem became larger: how could someone encountering the image follow an attribution that remained useful across those settings?

Generative AI adds another change. An image can be synthesized as well as captured. The 2014 generative adversarial network paper described training a generator alongside a discriminator. Latent-diffusion research first published in 2021 demonstrated high-resolution image generation guided by inputs including text. These are milestones in the means of producing visual work, alongside a much longer history of computer-assisted creation. [H13, H14]

Those changes strengthen the case for recording origins. A viewer should be able to understand whether a work came from a camera, a brush, digital tools, a generative model, or a combination of them. The account should be attached to a reference people can find, rather than depend entirely on the survival of a caption.

3. The work and the forms it takes

Consider a novel. The words can be read on a screen, printed in a paperback, or bound in a carefully produced hardcover. A publisher can issue a revised edition. The author can sign an individual copy. That signature adds a fact about the particular object and its relationship to the author; it does not create a different novel every time the author picks up a pen.

ISBN distinguishes the publications involved. Different editions and product formats generally require different ISBNs. A hardcover and an EPUB therefore do not ordinarily share one ISBN simply because they contain the same words. The system is valuable precisely because someone can order the form they intend to receive. [H08, H09]

A photograph has the same practical layers. The visual work might appear as a digital download, a page in an art book, a pharmacy print, a fine-art print, or a signed and numbered edition. Dimensions, paper, print process, edition, and signature describe particular manifestations. The attribution of the underlying work remains useful across all of them.

We believe the work deserves an intelligible identity across those uses. A publisher should be able to identify the photograph inside the book. A gallery should be able to point a visitor toward the recorded artist. A person who sees the image online should be able to inspect the attribution and follow the evidence behind it. The market can distinguish a signed print from another reproduction while preserving the connection to the work they share.

IPFN supplies a reference around which those relationships can be described. Its exact-file fingerprint answers a precise technical question: whether supplied bytes match the file recorded in the evidence. A resized JPEG, a printed photograph, and a scan of the print have different bytes. Their relationship to the work requires additional recorded information or supporting evidence. Keeping these layers distinct lets the system identify a work without pretending every reproduction is the same file.

The distinction also helps with older publications. A work whose copyright has expired can enter the public domain and be published by multiple publishers, each identifying its own editions and formats. Public-domain status differs from a Creative Commons license, which normally grants permissions under continuing rights. New translations, illustrations, and other original additions may have their own protection. [H28, H29]

Visual works as perceptual events

We propose thinking about visual works as perceptual events: images that people encounter, recognize, interpret, and reproduce in different forms. This is our way of explaining why the work's identity matters beyond the physical object or one digital encoding. The viewer recognizes an image; the evidence system must explain what the claim concerns and which particular file has been checked.

That framework accommodates different production methods without ranking their artistic worth. Provenance should explain a work's origins. The person looking at it can then make an informed judgment about its significance, use, or value.

4. Following one work through IPFN

The file chain connects a human account to reproducible comparisons. Each component has a defined job.

1. **The claimant supplies the work and the claim.** A submitter acts through the registration workflow, identifies the claimant, and supplies the file and relevant information. The submitter may be the artist or someone acting for them.

2. **The service fingerprints the original file.** SHA-256 computes a digest from the stored bytes before registration. An independent reader can later compute the same digest from a supplied file and compare the result.

3. **The registration system assigns the IPFN.** The assigned serial gives the registration a consistent reference in the controlled namespace. Its form also indicates the work's series, such as `LENS` or `ART`.

4. **The certificate explains the record.** The human-readable receipt names the identifier, recorded claimant and other assertions, including the available work digest. Its IPFS content address allows its encoded content to be checked.

5. **The daily ledger binds the batch.** An entry connects the IPFN, claimant, certificate address, and work digest. The exact ledger bytes receive their own SHA-256 digest and link to the preceding built ledger.

6. **Base witnesses the commitment.** A signed transaction carries the ledger digest and predecessor reference. A reader can inspect the transaction, compare its data with the ledger, and evaluate its inclusion and finality.

The result is a loop a person can follow: from the identifier to the registration, from the registration to the certificate and ledger, and from the ledger to the commitment. The original file can then be compared with the digest in that same evidence path. The blockchain does not need to store the photograph or assign its web address. It witnesses the document that connects the relevant references.

Controlled registration and independent public checking

IPFN is a controlled registration system. Recognized identifiers are assigned through its rules and registration service. The current authority mapping directs LENS and ART to ARTSSN, and BKS and MAG to Bookend Press. The registry operator maintains the records and namespace; the anchor operator controls the configured signing key. Blinding Visuals Press / IPFN Agency is the operating identity used in this paper. Appointment of additional registration authorities requires separate authorization. [C1, C2]

Public verification is a different authority. A person with the published artifacts can retain them and repeat the comparisons without obtaining permission to issue identifiers. This is the institutional distinction illustrated by ISBN and DOI: an accountable assignment system can support broad use of its references. Describing IPFN's mechanism does not confer registration authority or a software license. [S1, S2]

FIGURE 01

The file chain follows the work



A Light Across the Canyon

Kilian Liptrot / LENS-yc.3pfn.xyz

01	→	02	→	03	→	04
Fingerprint		**Register**		**Pin**		**Anchor**
Hash the original file.		Assign the web address.		Publish the readable certificate on IPFS.		Record the daily ledger digest on Base.

Open the IPFN to find the claim and its evidence. Compare the original file, certificate and daily ledger to check what was recorded.

5. The identifier people use and the fingerprints machines check

Ordinary new identifiers use `TYPE-SERIAL.3pfn.xyz`. The current types are `LENS`, `ART`, `BKS`, and `MAG`. Each type is a fixed series label; contributors do not invent arbitrary alternatives while expecting official assignment. The namespace uses ordinary DNS and HTTPS, so opening a supported identifier can reach its web record. Compatible platforms can display its Open Graph preview. [C1, C4, S8]

We chose `3pfn.xyz` as a short, fixed ASCII namespace that people can type, print, and recognize. The official address has that exact spelling; the system does not substitute Unicode lookalikes. Consistent spelling and the restricted serial alphabet reduce ambiguity when people share an identifier. Other domains can still imitate a familiar address, so the official namespace remains part of what a reader checks.

The normal serial generator selects from the 31-character alphabet `abcdefghjkmnpqrstuvwxyz23456789`. It excludes `i`, `l`, `o`, `0`, and `1`, reducing common transcription confusions. It starts at two characters, grows after each 1,000 unsuccessful attempts, checks existing identifiers and banned patterns, and stops after 10,000 attempts. A case-insensitive unique registry key protects assignment. These choices support a short reference people can share, read, and print. [C1]

The serial is assigned. The SHA-256 digest is computed. The IPFS CID identifies an encoded content graph. Those are three different functions. Knowing the serial does not supply a secret, prove who created the work, or recreate the file digest. The record and file chain connect them in a form another person can inspect.

Historical references also matter. Parsing accepts case variations, legacy dot-separated forms, historical `1pfn.xyz` workspaces, and the legacy dot-form `BOOK` type. The inaugural `LENS-1.3pfn.xyz` is a designated exception to ordinary serial generation. A parser may accept a historical reference that the current generator would never issue. The registry lookup determines whether an assigned record exists, and live DNS reachability is checked separately. [C1]

The first designated ART case in this edition, `ART-1x.3pfn.xyz`, likewise uses a controlled initial assignment. Its `1` is deliberately outside the ordinary generator's alphabet. The assignment is bound to this source artwork and its established personal account, with the same draft used for recovery. It does not open a custom-number field for general registrations. [C14]

6. Declaring how the work was made

The claimant should be able to explain the means of production without being forced into a hierarchy of artistic value. Four useful optional declarations are:

Declaration	What the claimant is describing
Handmade, captured digitally	A physical work made by hand and then photographed or scanned.
Made with digital tools	A person's work produced using tools such as digital drawing or design software.
AI augmented	A workflow in which AI contributes to or transforms material alongside other creative work.
AI generated	A work whose visual material was produced through a generative AI system.

The categories communicate an account of production. A declaration can be supported by process files, source material, or additional evidence; its presence records what the claimant says. It should not be mistaken for an automated determination of how the image was made.

The distinction is especially useful when a physical or digital object carries several layers of production. AI-generated artwork can be printed, placed in a scene, and photographed. The artwork and the photograph depicting it are different objects with different histories. A paper using that photograph on its cover should keep those histories intelligible.

Copyright remains a question about the particular expression and applicable law. The U.S. Copyright Office's 2025 analysis distinguishes human expression created with assistance from material whose expressive elements were produced by a machine. A production label supplies relevant information; it does not decide that legal analysis. [H15]

7. Building on the work that came before

The contribution of IPFN is the relationship among its components: a controlled identifier, an attributed record, a readable certificate, published ledger bytes, and an inspectable commitment. Its components draw on established cryptography, distributed storage, web infrastructure, and institutional identifier practice.

Haber and Stornetta's 1991 timestamping work addressed how digital documents could be dated so that undetected backdating would be difficult. Bitcoin's timestamp-server explanation made the batching idea particularly clear: hash a block of items and publish the hash, then connect timestamps into a chain. IPFN applies that useful witnessing pattern to a ledger of registrations on the existing Base network. Its evidence case checks a complete ledger file rather than a Bitcoin Merkle inclusion proof. [P1, P2]

The Ethereum papers contribute a discipline of named objects and explicit state transitions. A registration can exist before its certificate is pinned, and a ledger can exist before a transaction succeeds. IPFN's public presentation and verifier must preserve those distinctions. Base supplies the chain actually used here, with chain ID 8453 and its own transaction, block, and finality observations. [P3, S6]

Prior system	What it contributes to understanding IPFN
IPFS	Content addressing and linked content graphs explain how a retrieved certificate can be checked against its CID. Maintained pins keep the bytes available. [P4, S5]
Filecoin	Its storage proofs and storage/retrieval markets make the separate job of retaining data explicit. IPFN's demonstrated evidence does not depend on an asserted Filecoin deal. [P5]
DOI	Registration agencies, metadata, resolution, and continuity obligations show why an identifier needs continuing institutional stewardship. [S2]
ISCC	Content-derived metadata, similarity, and instance codes address discovery and comparison. IPFN's assigned serial and claimant registration answer a different question. [S3]
C2PA	Signed assertions and media bindings describe a complementary provenance layer. Preserving a file can preserve credentials; credential validation is a separate operation. [S4, C12]
W3C DID	Its separation of subjects, controllers, methods, and resolution clarifies authority. IPFN uses its controlled DNS namespace rather than a DID method. [S7]

These systems solve related problems through different arrangements. The comparison explains the design choices; a deployed integration or institutional relationship requires its own evidence.

Why this does not require an NFT

An NFT centers on tracking and transferring a token. IPFN centers on registering a claim and making its evidence accessible. Attribution should remain intelligible as a photograph is licensed, printed, or sold. IPFN therefore makes the identifier and registration primary; token ownership is not an input to its verification procedure.

NFTs can also reference human-readable metadata and content-addressed files. ERC-721 permits metadata URIs, and IPFS publishes guidance specifically for NFT data. The case for IPFN therefore rests on the job its complete arrangement performs: consistent agency assignment, an ordinary web address, a claimant record, and a defined evidence path. It does not depend on treating every token as an arbitrary number disconnected from its subject. [P6, P7]

8. Beyond a single image file

The same architecture can organize claims concerning other digital objects. A document, video, advertisement, or publication can be represented by defined bytes and an attributed record. A collection can use a manifest: a document listing its constituent files, their fingerprints, and their relationships. The manifest then becomes a specific object that can itself be fingerprinted and committed.

This is particularly useful for the publishing problem that started the work. A book's production package can connect the publication with the photographs, text, layouts, and other files that make it up. Its ISBN can continue to identify the commercial publication while the file chain supplies a separate evidence arrangement.

Current image registrations support an exact-file `content_digest` scope. Current `BKS` and `MAG` records can instead be `catalog_assertion` records, without a digest of a complete publication package. The repository contains a provisional publication-manifest model. Completing and commissioning its full artifact and signing pipeline is the next technical work, rather than an assumed property of every existing book record. [C2, C3]

9. The economics of maintaining an institution

A registry has to remain useful after the first demonstration. It needs to accept work, retain evidence, resolve identifiers, maintain records, and handle questions or disputes. The design should make those continuing obligations affordable.

Batching is central to that design. A ledger can contain many registrations while one transaction commits its 32-byte SHA-256 digest. Adding another eligible entry to that batch does not require a separate blockchain transaction. The demonstrated 31 October volume contains 113 registrations behind one commitment. The savings concern shared anchoring; storage, retrieval, intake, support, and preservation still require resources. [C9, E1]

U.S. copyright registration illustrates why both price and administration matter. As examined on 3 November 2026, the Copyright Office's published fees were $45 for an eligible Single Application, $65 for a Standard Application, and $55 for a group of up to 750 eligible photographs. Ten thousand eligible photographs can therefore be filed in at least 14 groups for $770 at that rate. Published and unpublished photographs use separate options, groups require the same author and claimant, and published groups must fall within one calendar year. [H24, H25]

IPFN's institutional case is broader than replacing that fee. It makes a recorded claim and its evidence usable in everyday publishing, through a reference someone can open from elsewhere on the web. U.S. registration has statutory consequences under 17 U.S.C. §§410–412; that function remains distinct. Fees are also under review, so this comparison is dated to the published schedule examined. [H26, H27]

The European opportunity

Several European systems demonstrate why protection and public identification are separate jobs. The United Kingdom has no register of copyright works. Ireland has no copyright-registration procedure. Dutch guidance states that copyright is not officially registered and points toward evidence measures such as a notary or BOIP's iDEPOT. France protects copyright from creation while INPI offers a separate e-Soleau dated-evidence service. Germany likewise supplies copyright without constitutive registration, alongside a narrow anonymous-and-pseudonymous-work register concerning the protection term. [H16–H22]

EUIPO's 2026 mapping study identifies copyright information fragmented across national, sectoral, public, and private databases and examines the basis for a possible CopyrightView search service. That is the institutional opening: useful records and evidence that can be made easier to discover and exchange. [H23]

We believe European institutions are well placed to lead that work. IPFN offers a concrete system to evaluate with real creative works and independently repeated checks. A first institutional phase can be organized in months: agree the works and evidence, register a bounded collection, reproduce verification, and test the handling of corrections and challenged claims. Wider adoption should follow the evidence that exercise produces.

10. Two creative works, two useful demonstrations



A Light Across the Canyon
LENS-yc.3pfn.xyz



An AI-generated cover study
ART-1x.3pfn.xyz

A photographed landscape: A Light Across the Canyon

The anchored example is *A Light Across the Canyon*, `LENS-yc.3pfn.xyz`, recording Kilian Liptrot as claimant. A separate read-only download of its original JPEG yielded 1,591,834 bytes. Its SHA-256, `938e065ecc9b7ff0b40ccd604d1c9ddcb4c0f70e03b35e5d89efdd749c1279f2`, matched the digest in the 31 October 2026 production ledger. The certificate address was `QmQEWUWBh6afkwW8ByiBUyCpWrmG9RLxo8QqUHf9R5rYsE`. [E2]

The 113-entry ledger comprised 43,758 canonical JSON bytes and hashed to `94c4ba61af14f74dd15bdc3cd10d210b055671a3172ae9ca58153d49a9f9f9ba`. Its Base transaction was `0x25b1532e64f3af4d363cf21df71ee56b189a5644876f9139ec5c310ba57f8251`, included in block 53,351,386 at `2026-11-01T00:15:19Z`. The retained RPC observation placed that block below its reported finalized head. [E1, E2]

The reproduction reconstructed the certificate CID, checked unique ledger inclusion and the binding between the recorded claim, certificate, and ledger, and compared the transaction input, receipt, and block references. Reconstructing the signed transaction reproduced its hash. Signature recovery yielded `0x20ce5b8aaF38c647A42CfC6dC038a12A12648F6e`, matching the RPC sender and the operator-declared anchor address. [C13, E2]

This demonstrates the useful public result: a reader can follow a named claim and independently compare the original file with its recorded evidence. The timestamp concerns the committed

record. The claimant's account of creating the photograph remains an assertion supported or challenged through other evidence.

The artwork depicted on this paper's cover

The supplied artwork was registered on 3 November 2026 at 10:59:56 UTC as **ART-1x.3pfn.xyz**, under Kilian Liptrot's established personal artist account. Its title is *An AI-generated cover study*. The description explains its use as an ART case study for Blinding Visuals Press and this paper. The unchanged 1,254-by-1,254 PNG comprises 2,488,895 bytes; its original SHA-256 is `66c910899133801a00ecd5b93e27e788504c4bc2652a19e788d4bc8c2173c43a`.

The creator selected the existing **AI-generated** declaration. Live readback found that active label in the linked registration taxonomy, and the public verifier displayed **Fully AI-generated** before its expandable details. This is the creator's current production declaration. It describes how the work was made without ranking its artistic worth or presenting the label as automated AI detection.

The certificate was pinned at `2026-11-03T11:00:02.571023Z` under CID `QmZGjoTgLRZhG9wyCz2mTWpdwGL12qYfpMeWH9qhhDjFnn`. The certification job completed successfully. At the 11:00:29 UTC observation, the public lookup matched the supplied original digest to ART-1x and reported pinned, unledgered evidence. No ledger volume or completed blockchain anchor existed for this artwork at that observation. Its registration and certificate therefore demonstrate the earlier stages of FileChain; the Canyon photograph demonstrates the completed anchored path. The public case is ipfn.net/verify/ART-1x.3pfn.xyz . [C14]

This second case lets the publication demonstrate its own argument. A person encountering the depicted artwork can find its registered attribution and declared production method. Its evidence state is described as observed, separately from the already anchored landscape. The artwork, the supplied cover image, and this publication are distinct objects whose relationships can be explained through their records.

11. What an accountable registry must maintain

The published evidence makes important comparisons possible outside the operator's database. Institutional credibility also depends on the administration surrounding those comparisons.

Responsibility	How it relates to the evidence
Assignment and identity assurance	Record who submitted the claim, for whom, and what assurance was actually obtained. A timestamp preserves a claim's history; it does not make a false claim true.
Signing authority	Maintain an accountable anchor address and incident procedure. The demonstration matches an operator-declared address; a complete public historical rotation record remains an institutional requirement.

Responsibility	How it relates to the evidence
Preservation and resolution	Fund DNS, certificate pins, ledger copies, and evidence export so the referenced material stays reachable.
Corrections and disputes	Keep committed assertions inspectable while distinguishing later corrections, decisions, or current catalog information.
Data protection	Define lawful publication, access, retention, correction, and international processing of claimant information.
Evaluation	Measure availability, latency, cost, and errors before advertising numerical service promises.

The ledger and certificate protect the integrity of the record they bind. The institution must preserve the means of finding it and make clear how a challenged claim is handled. These responsibilities belong together.

The next phase should use a bounded collection with retained originals and independent verification. It should include the difficult cases: a changed file, a missing gateway, an unsuccessful transaction, an unauthorized anchor, a correction, and a disputed attribution. Perceptual-hash and OpenCLIP paths in the repository can support research into related images; a plagiarism-detection beta needs a defined corpus, transformation tests, error rates, and human review before broader claims about performance. [C12]

We understand the publishing problem. We have built a file chain that connects the claim to evidence another person can examine. We want to work with the institutions that can test that arrangement, improve its governance, and put it to use for creators. The following appendices make the technical case concrete enough to inspect.

Appendix A. Objects, actors, and registration states

The claimant is the person or organization whose claim is recorded. The submitter operates the account and supplies the material. The registration authority accepts the registration under IPFN's rules. The registry operator maintains the assigned identifiers and records. The anchor operator controls the configured transaction-signing key. Storage and DNS operators keep the references reachable. The verifier retrieves evidence and performs comparisons. Each role has a separate responsibility. [C1, C2, C8]

Object	Definition
Creative work	The photograph, artwork, or publication about which a claim is made.
Work file `W`	A particular sequence of bytes supplied or preserved for checking.
Registration `R`	Assigned IPFN, object reference, recorded claimant, and evidence links.

Object	Definition
Work digest `hW`	`SHA256(W)`, recorded for an image's exact bytes.
Certificate `C`	Published receipt bytes containing the registration assertions.
Certificate CID `cC`	IPFS content address for the receipt's encoded content graph.
Daily ledger `L`	Exact serialized JSON bytes containing an ordered batch of entries.
Ledger digest `hL`	`SHA256(L)`, carried in the transaction's data.
Anchor `T`	The transaction, receipt, and associated block evidence.
Catalog presentation	Current title, claimant display, caption, and preview, distinct from committed historical bytes.

Original-file custody and the certificate

The image workflow derives the authoritative digest from stored original bytes on the server. The immediate finalize path compares `serverSha256` with a supplied browser value and rejects disagreement with `digest_mismatch`, HTTP 422. The archival path streams and hashes the original and rejects `archive_digest_mismatch`. These are source-observed controls; the worked photograph supplies a separate observation of publicly retrieved original bytes. [C5]

The certificate path includes `IPFN`, `AUTHOR` or `AUTHORS`, `TYPE`, `TITLE`, `PUBLISH_DATE`, `PUBLISHER`, `REGISTRATION_AUTHORITY`, `EDITION`, and `AUTHENTICITY_CODES`. For the image specimen, `AUTHENTICITY_CODES.SHA_256` contains the work digest. Historical receipts include a text heading before their JSON, and every byte belongs to the content-address check. The `AUTHOR` field records an assertion. The receipt is the registration artifact; its integrity is checked through its CID and ledger binding. [C6]

State or event	Meaning
Registered	A registry entry exists.
Pinned	The operator recorded a certificate CID; retrieval and CID comparison can be performed separately.
Included in a volume	The record is linked to a volume; the independent check locates its entry in the committed ledger.
`pending`	The volume awaits anchoring.
`sending`	A transaction identity has been prepared and retained.
`anchored`	The worker recorded validated transaction inclusion.
Finalized observation	A fresh chain observation places the included block at or below the observed finalized head.
`failed`	The anchor operation failed or was refused.

Public lookup derives `anchored`, `ledger_pending`, `link_failed`, and `unledgered`. Production and test evidence remain distinct. This state model gives a reader a useful answer even when one part of the chain is still incomplete. [C7]

Appendix B. Exact ledger bytes and the signed commitment

The daily builder selects eligible unledgered, pinned records before its cutoff. This permits catch-up inclusion; a volume's date must not be read as every work's creation date. Production eligibility uses the recorded registration-time boundary, so repinning an old test record does not automatically turn it into production evidence. [C8]

The current ledger schema is `1.3.0`. Its top-level fields, in order, are:

`schema_version`, `volume_label`, `ledger_date`, `window_start`, `window_end`, `previous_base_tx_hash`, `previous_ledger_sha256`, `base_chain_id`, `ipfn_count`, `standards`, `publisher`, `ipfns`.

The `standards` object orders `timezone`, `locale`, `protocol_language`. Each entry orders `ipfn`, `ipfn_type`, `claimant`, `pinned_cid`, `content_sha256`, `last_pinned_at`, `target_url`. Entries are sorted by the IPFN string using the source's lexicographic comparison. `ipfn_count` equals the array length. `content_sha256` can be null for a publication catalog assertion. Titles are not present in this ledger schema. [C9]

The builder requires explicitly zoned timestamps and normalizes `window_start`, `window_end` and entry pin times to `YYYY-MM-DDTHH:mm:ss.sssZ`. `ledger_date` remains a date. It constructs the fixed-order object, applies ECMAScript `JSON.stringify(ledger, null, 2)`, appends one line-feed character, and encodes UTF-8. It does not apply an additional Unicode normalization pass or use an interchangeable “any canonical JSON” formatter. ECMAScript string escaping remains relevant to reproduction. [C9, S9]

For verification, preserve the downloaded bytes and hash them directly. Parsing and reformatting first can change the digest. Rebuilding the bytes in another language requires the same field order, sorting, timestamp normalization, JSON escaping, whitespace and final newline. That byte profile is the reproducible implementation described here; cross-language conformance needs its own test vectors.

Schema `1.3.0` links to the preceding successfully **built** ledger using `previous_ledger_sha256`, whether that earlier volume is pending, sending or anchored. Selection excludes failed or archived volumes, then orders by ledger date, sequence and creation time. A predecessor lookup error refuses the build rather than silently manufacturing a genesis. The older `previous_base_tx_hash` remains a compatibility reference to an earlier anchored transaction; it is not the authoritative linear ledger link. [C9]

The canonical storage key is create-only: identical existing bytes permit a no-op, while different bytes under the same key cause refusal. Historical bytes are not rewritten to the newest format. Version `1.0.0` had the initial entry fields; `1.1.0` added descriptive and digest fields; `1.2.0` removed title; `1.3.0` added the previous-ledger digest and strict timestamp normalization. [C9]

The transaction payload is UTF-8 text:

`FILECHAIN V2 VOL=<volume> SHA256=<hL> PREV=<previous hL or GENESIS>`

Historical missing-version and `1.0.0-1.2.0` ledgers use `FILECHAIN VOL=<volume> SHA256=<hL>`. The payload helper chooses V2 for other declared versions. A verifier still needs to recognize the ledger schema before interpreting its fields; selecting a payload format is not permission to assume unknown schemas have current semantics. [C10]

The worker checks chain ID 8453 and the served ledger digest. It prepares a signed zero-value self-transaction, retains the exact signed bytes and transaction hash before broadcast, and uses a conditional database transition to claim the volume. Recovery can reconcile or rebroadcast that same transaction identity rather than silently sign a replacement. Receipt validation checks the chain, requested hash, configured sender and destination, value, payload and canonical block. Finality is promoted separately after fresh receipt checks and a finalized-head observation. The worked example exercises the read-only comparisons; fault-injection and independent Base derivation remain separate evaluation work. [C11]

FIGURE 02

Three different byte comparisons

Original file	→	SHA-256	→	Recorded file digest
Certificate bytes	→	IPFS encoding	→	Certificate CID
Daily ledger	→	SHA-256	→	Digest in the Base transaction

The certificate identifies the original file. The ledger lists the certificate. The transaction records the ledger digest.

Appendix C. Independent verification procedure

A public verifier needs the registration identifier `I`, exact ledger bytes `L`, transaction identity `T`, certificate CID and certificate bytes or blocks, and access to Base transaction, receipt and block data. The work bytes `W` are optional but necessary for an exact-file result. An auditor also needs a separately established authorization context: which sender was authorized to anchor for the operator at the relevant time.

The following procedure defines the evidence report used in this paper. The website lookup and an independent verifier are different interfaces; the report below makes the complete comparison requirements explicit.

1. **Locate and classify.** Normalize `I` using the accepted syntax and obtain its registry reference. Retain source URLs, retrieval times, environment and declared schema. A missing record is `not_found`; an unreachable service is `unavailable`. Do not convert either into a finding about copyright.

2. **Hash before parsing.** Calculate `hL = SHA256(L)` over the exact download. Parse the ledger under its declared supported schema. Check field types, count, identifier uniqueness and the selected entry. Report absent or duplicate inclusion explicitly; a hash match cannot repair a malformed or ambiguous record.

3. **Read the transaction.** Confirm the RPC chain is 8453, retrieve the requested transaction and receipt, and compare transaction identities. A missing receipt is pending or unavailable until resolved; a reverted receipt fails the successful-anchor check.

4. **Compare the commitment.** Decode the transaction data as UTF-8 and construct the exact expected V1 or V2 payload from the ledger's own schema, volume label, calculated digest and previous digest. Compare the complete strings, not a substring. Confirm the containing block agrees with the receipt and the canonical block at that height. Report whether the observed finalized head has passed it. Record provider and observation time. [C10, C11, S6]

5. **Check anchor authority separately.** Compare the transaction's sender, self-destination and zero value with the approved operator context for that period. A matching payload from an unknown address can establish a public commitment, but not an authorized IPFN anchor. If the historical authorization context is unavailable, report `authority_unestablished`; do not invent one from the address appearing in the transaction.

6. **Check the certificate.** Validate its CID using the indicated codec, multihash and complete encoded content graph, or a known importer profile that reproduces it. A raw SHA-256 of a gateway file is not generally its CID. Preserve any receipt heading when checking bytes. For the demonstrated small CIDv0 certificate, the verifier reconstructs a single DAG-PB UnixFS File without links or optional metadata. Unknown profiles are `unsupported`, not a successful check. [S5, S10]

7. **Bind the objects.** After certificate byte validation, parse its supported receipt format and compare its IPFN and available work digest with the selected ledger entry. Compare claimed roles as recorded assertions. Check the previous-ledger hash recursively when making a chain-history claim; one volume's success does not prove completeness of the operator's entire history.

8. **Compare the work.** If `W` and a valid recorded work digest exist, calculate `SHA256(W)` and compare them. If either is absent, return `not_checked` or `not_applicable` for a catalog assertion. Never substitute the website's compressed preview for the claimed original.

Return a report with separate results for `record`, `schema`, `ledger_bytes`, `inclusion`, `transaction`, `authority`, `finality`, `certificate`, `record_binding`, `work_file` and `history`. Each check records `match`, `mismatch`, `pending`, `unavailable`, `unsupported`, `not_checked` or `not_applicable` as appropriate, with the observed values. A published dispute notice, if supplied by a responsible authority, is recorded separately with its source. Absence of such a notice is not proof that the claim is undisputed.

The three principal relationships are distinct: `SHA256(W) = hW`; certificate graph validation yields `cC`; and `SHA256(L) = hL` in the examined transaction. A certificate match does not establish a

work-file match. A ledger match does not establish authorized issuance. None of these operations determines legal ownership.

Appendix D. Trust and evaluation

The specimen's transaction signer matches the operator-declared anchor address, supported by independent signature recovery and the authenticated state reference. The examined material supplies no complete public, time-versioned authorization and rotation history. A compromised signing key could create valid new transactions carrying false assertions; an institutional incident procedure must address that possibility. The registration artifact is distinct from an X.509 certificate, a detached claimant signature, or a qualified electronic timestamp. The notary role described in this architecture concerns the public commitment and evidence path; legal effects depend on the relevant law and evidence. [C11, C13]

Challenge	Required verifier or institutional response
Changed original, certificate, or ledger bytes	Recompute the relevant hash or CID and return a mismatch. In-memory changes to all three failed comparison in the retained demonstration.
Right payload from an unrecognized sender	Separate agreement with the commitment from authority to anchor for IPFN.
Missing registration or duplicate ledger entry	Report the inclusion problem; a valid ledger digest does not resolve it.
Pending or reverted transaction	Report pending or failure, preserving the transaction identity.
Unreachable gateway, DNS, or registry	Report unavailable for the affected operation and retain previously retrieved evidence.
Conflicting chain-provider data	Compare transaction, receipt, and canonical block; use additional providers or an independently operated node when needed.
Changed catalog or contested claimant	Preserve committed bytes and identify the later assertion or dispute through a responsible process.

The retained case demonstrates batching and reproducible file, certificate, ledger, transaction, and signer comparisons. Validator fixtures and browser checks also exercise unauthorized-sender, pending/reverted-transaction, and unavailable-service responses. These checks establish the software's responses to supplied conditions. Live signing-key incidents, operator recovery drills, and a completed contested-claim procedure remain institutional evaluation work. The evidence does not measure throughput, median or tail latency, long-term availability, recovery success rate, or total retained-work cost.

FIGURE 03

Registration and evidence states



Filled circles show completed stages at the 3 November observation. The Canyon entry had a confirmed anchor. The cover artwork had a pinned certificate and was awaiting ledger inclusion.

References and implementation evidence

Primary specifications support the mechanisms described. This edition retains the 1 November evidence observations and adds historical and protocol research reviewed on 3 November 2026.

S1 International ISBN Agency, *What is an ISBN?* Edition/format identification, trade use and distinction from copyright protection. isbn-international.org

S2 DOI Foundation, *What is a DOI?* and *What Are Registration Agencies?* Persistent resolution, metadata, controlled agency participation and continuity responsibilities. doi.org ; doi.org . Handbook landing and version of record: doi.org

S3 ISCC, *Specification*, describing ISO 24138:2024 and the distinction between the published core and later extensions. iscc.codes

S4 C2PA, *Technical Specification 2.4*, sections 9-15, content binding, claims, signatures, trust and validation. spec.c2pa.org

S5 IPFS, *Content Identifiers* and *Pin files*. docs.ipfs.tech ; docs.ipfs.tech

S6 Base, *eth_chainId*, *eth_getTransactionReceipt*, and derivation specification. docs.base.org ; docs.base.org ; docs.base.org

S7 W3C, *Decentralized Identifiers*, identifier, controller, method and resolution model. w3.org

S8 *The Open Graph protocol*, page metadata used by compatible consumers. ogp.me

S9 ECMAScript, `JSON.stringify`. tc39.es . SHA-256 definition: NIST, FIPS 180-4. csrc.nist.gov

S10 IPFS, *UnixFS*, DAG-PB file encoding and content graph structure. specs.ipfs.tech

References C1-C13 describe the examined source at repository commit `31e6132d206d4e0e8e0fdd2a7b8144e5f910568c`. The 3 November release adds the designated ART assignment, production-label readback, presentation changes and curated paper publication described in C14. Source review and the dated live specimens are separate evidence. Repository: github.com .

C1 `src/lib/ipfn/format.ts`, lines 13-121; `migrations/MIGRATION_IPFN_DASH_SEPARATOR.sql`, lines 56-94; `migrations/20260603_ipfn_registry_citext_unique.sql`, lines 252-258. Parsing, issuance algorithm and unique-key definition. The historical format document's old ledger-schema description is superseded by current code.

C2 `src/lib/cert/registrationAuthority.ts`, lines 21-70; `src/lib/filechain/verifyResolver.ts`, lines 90-129. Authority mapping, evidence scope, and current catalog projection.

C3 `src/lib/filechain/publicationManifest.ts`, lines 3-24 and 86-104. Provisional publication-package model and explicit unimplemented artifact/signing boundary.

C4 `cloudflare/proof-router-worker/worker.js`, `fetchProofMeta`, `renderProofHtml` and wildcard route; `src/middleware.ts`, host-based resolution. Live reachability and platform rendering are separate observations.

C5 `src/app/api/uploads/finalize/route.ts`, lines 137-211; `src/lib/intake/workerEntry.ts`, lines 255-280; `src/lib/intake/archiveFile.ts`, lines 10-49. Server-derived original digest, browser cross-check and archival custody.

C6 `documentation/hetzner/embed_api.py`, lines 462-517; `src/lib/cert/runIpfsPreview.ts`, lines 58-175. Pinned receipt construction and separately rendered certificate fields; not evidence of a detached claimant signature.

C7 `src/lib/ipfn/anchorStatus.ts`, lines 19-46; `src/lib/filechain/verifyResolver.ts`, lines 31-74. Display vocabulary and persisted anchor states.

C8 `src/app/api/cron/filechain-build-ledger/route.ts`, lines 440-445 and 516-559; `src/lib/filechain/productionEligibility.ts`, lines 1-22. Candidate cutoff, predecessor lookup and production boundary.

C9 `src/lib/filechain/buildLedger.ts`, lines 59-135, 140-237, 256-296 and 303-349. Schema history, fixed-order construction, timestamp rules, previous-volume selection and byte serialization.

C10 `src/lib/filechain/anchorPayload.ts`, lines 3-56. Historical V1 and current V2 payload selection. The worker's matching implementation is `cloudflare/anchor-worker/src/pure.ts`.

C11 `cloudflare/anchor-worker/src/index.ts`, lines 489-582, 653-825, 830-946 and 1279-1408; `cloudflare/anchor-worker/src/pure.ts`, `receiptValidationDecision`. Same-transaction recovery, receipt identity checks, configured signer and staged finality. No complete public historical key-rotation policy is established by these paths.

C12 `src/lib/fingerprint/dispatch.ts`, lines 137-143; `src/lib/fingerprint/driftCompare.ts`; `src/lib/publishPipeline.ts`, `createImageFingerprint`; `src/app/(ipfn)/ipfn/learn/page.tsx`, Originals, exports and C2PA section. Exact original digests, derivative fingerprints and the boundary of C2PA behavior.

C13 `src/app/api/bvpress/filechain/state/route.ts`, line 86, `notary_address`. The authenticated state response is the operator's declared address reference; it does not publish a complete historical key-rotation policy.

C14 Release `7531cb4572e6cd5740e5fa799a74ec8fb1c8ffc0`: `src/lib/ipfn/founderArtOne.ts`, `migrations/20261103101737_p535_founder_art_one_x.sql`, `src/lib/images/publicTags.server.ts`, `src/lib/ipfn/paperEdition.ts`, and `scripts/ipfn/publish-paper-edition.mjs`. Designated initial assignment, existing linked production declarations and immutable paper publication with byte verification. github.com

E1 Retained first-edition read-only verifier and provider responses: `docs/evidence/p531/technical/verify_live_evidence.py`, lines 56-109, and `evidence/verification-summary.json`, observation `2026-11-01T16:53:16.836790+00:00`. Public canonical ledger: images.blindingvisuals.com . Transaction: basescan.org

E2 Revised demonstration, *A Light Across the Canyon*, `LENS-yc.3pfn.xyz`: ipfn.net/verify/LENS-yc.3pfn.xyz . Original-byte observation and the revised certificate/transaction checks are retained with the accompanying revision-2 evidence demonstration. The photograph is a real registered work; it is not a generated illustration.

Historical and institutional primary sources

H01 U.S. Copyright Office, eighteenth-century chronology. copyright.gov

H02 U.S. Copyright Office, nineteenth-century chronology, photographs in 1865. copyright.gov

H03 Berne Convention, authentic Paris Act as amended 1979, Articles 2(2), 5(1)–(3), and 6bis. wipo.int

H04 WIPO, voluntary registration systems. wipo.int

H05 WIPO, copyright FAQ, international registry question. wipo.int

H06 Bowker, U.S. ISBN agency, ISBN history. isbn.org

H07 International ISBN Agency, historical account. isbn-international.org

H08–H09 International ISBN Agency, definition and assignment rules. isbn-international.org ; isbn-international.org

H10–H11 National Museums Scotland, collection account; Getty Museum, original calotype research. nms.ac.uk ; getty.edu

H12 Steven J. Sasson, declaration paragraph 6 and attached résumé, USPTO-hosted Apple Exhibit 1010. ptacts.uspto.gov

H13–H14 Goodfellow et al., 2014; Rombach et al., first submitted 2021. Original generative-model research. arxiv.org ; arxiv.org

H15 U.S. Copyright Office, Copyright and Artificial Intelligence, Part 2, January 2025, summary iii and conclusion page 41. copyright.gov

H16–H19 National offices and government guidance: United Kingdom, Ireland, Netherlands, Germany. gov.uk ; ipoi.gov.ie ; business.gov.nl ; dpma.de

H20–H22 France, Copyright Code Article L111-1 and current INPI guidance on automatic rights and e-Soleau. legifrance.gouv.fr ; inpi.fr ; inpi.fr

H23 EUIPO, 2026 database and metadata mapping, executive summary. euipo.europa.eu

H24–H25 Copyright Office fees and group-photograph eligibility, checked 3 November 2026. copyright.gov ; copyright.gov ; copyright.gov

H26–H27 Copyright Office 2026 fee study; 17 U.S.C. Chapter 4, especially §§410–412. copyright.gov ; copyright.gov

H28–H29 Copyright Office, copyright lifecycle; Creative Commons, its licenses and public-domain tools. copyright.gov ; creativecommons.org

H30 Tribunal judiciaire de Marseille, First Civil Chamber, *AZ Factory v Valeria Moda*, 20 March 2025, RG 23/00046. Judgment text, especially the discussion of copyright title and the two blockchain-timestamp reports; examined 3 November 2026. legalis.net

Protocol prior art

P1 Haber and Stornetta, How to time-stamp a digital document, 1991, publisher record and abstract. doi.org

P2 Nakamoto, Bitcoin, 2008, §3 timestamp server. bitcoin.org

P3 Buterin, Ethereum Whitepaper, historical 2014 account/state model; Wood et al., Ethereum Yellow Paper, §§2 and 4–6. ethereum.org ; ethereum.github.io

P4 Benet, IPFS, 2014, §§3.4–3.6. research.protocol.ai

P5 Protocol Labs, Filecoin, 2017, §§1 and 3–5. research.protocol.ai

P6–P7 ERC-721 specification and IPFS guidance for NFT metadata. eips.ethereum.org ; docs.ipfs.tech