

DRAFT V0.7 11/12/2025 15:45 ET - [this line to be removed once finalized]

Digital Vellum Note (DVN)

DVN #3

Date: 2 November 2025

Author: Daveed Benjamin, Michael Witmore

Organization: Meta-Layer Initiative

Updates: [none]

A Future with Civic Memory Design

This section serves as the conceptual anchor for the document, connecting the philosophical idea of civic memory with the practical objectives detailed in the next section. It outlines how provenance as civic infrastructure directly supports continuity of trust and verifiable digital memory.

Provenance should form a tamper-evident, append-only, AI-traversable substrate of public memory. While content may fade or be deleted, the claims about it should endure as inspectable layers in a shared historical graph, allowing communities and agents to interpret digital history with context and continuity. Digital provenance thus becomes a core civic instrument of long-term trust.

Imagine how differently we might understand our shared past if a provenance-grounded digital noosphere had existed thousands of years ago. Clearer lineage of claims and context might have reduced repeated collective errors, tempered conflicts, and preserved nuance that history often loses. That is long behind us, but we now have an opportunity to accelerate toward a future where truth, context, and continuity are built into the fabric of the digital record.

Context and Purpose

Digital Vellum Note #3 defines how provenance underpins trust and memory in the digital age. It introduces a framework for establishing verifiable, interpretable, and enduring records of origin and transformation. The purpose of this note is to show how provenance can serve as the connective tissue of digital memory, ensuring that knowledge remains verifiable, interpretable, and enduring across generations of technology and cognition.

This note outlines fourteen desirable properties that enable digital content to carry inspectable, verifiable, and context-aware records of its origin, transformation, and endorsement. The aim is not to prescribe a rigid protocol but to articulate design contours and desirable properties that could sustain an interoperable, trustworthy provenance ecosystem. It aligns with Digital Vellum's core mission by linking verifiable origin and long-term durability with economic and civic value for preserved digital objects.

Digital provenance has traditionally referred to the traceable record of authorship, modification, and ownership of digital artifacts. In the age of networked knowledge, AI co-creation, and

dynamic archiving, this definition is no longer sufficient. Provenance now spans multiple layers of participation and preservation — from the original act of creation to the long-term custodianship of meaning.

The next section expands the conceptual and architectural scope of digital provenance to include archival continuity, relational provenance, and artificial participation, all potentially rendered within and/or accelerated by the emerging meta-layer above the web.

Our goal is to identify the contours of this broader system: what it must capture, how it must function across both human and machine agents, and why it must endure as an integral substrate for civic memory.

Contours of Digital Provenance

Each contour of digital provenance contributes to a shared goal: establishing continuity of trust across the entire lifecycle of digital content. The following dimensions—archival, relational, artificial, and meta-layer—work together to ensure that trust persists as content evolves, moves, and is reinterpreted. By clarifying how each supports verifiable continuity, we emphasize that provenance is not a static property but a system of enduring, interoperable trust.

We propose an expanded view of provenance that reflects the full lifecycle of digital content, including the actors, transitions, and interpretations that define its meaning over time.

1. **Archival Provenance** – The documentation of preservation and migration events as first-class provenance records. This includes not only what is created, but how it is maintained, moved, or re-encoded across time. Preservation without provenance risks cultural amnesia.
2. **Relational Provenance** – The mapping of contextual and semantic relationships between works - citations, remixes, critiques, endorsements, and counterclaims - forming a living graph of cultural and informational evolution. Meaning emerges through connection, not isolation.
3. **AI Participation** – The inclusion of artificial systems as accountable agents in the provenance chain. This covers both generative and preservational roles: AI as author, co-creator, verifier, or archivist. Agency now includes software, and its actions must be inspectable.
4. **Meta-Layer Integration** – The visualization and composition layer that makes provenance *legible, inspectable, and dynamic*. The meta-layer augments both the creation and display of digital provenance. It enables provenance to be surfaced contextually across the web, without requiring wholesale protocol adoption in browsers.

Where browser-native implementations may evolve slowly, the meta-layer can accelerate experimentation and deliver composable overlays that enrich provenance for creators, publishers, archivists, and AI systems alike. It bridges between protocol-level standards (e.g., C2PA, W3C PROV, DID) and user-facing interpretation, transforming provenance into a living civic interface rather than a buried metadata field.

Why Expand Provenance Now

The next two years represent a critical inflection point. As AI accelerates content production and the line between original and synthetic blurs, the ability to verify *who, what, when, and how* becomes foundational to digital trust. Provenance is not a luxury—it is the precondition for meaning in an AI-saturated web.

The shift toward AI-assisted creation, ephemeral publishing environments, and decentralized archiving demands a new conceptual scope for provenance. Without it, our digital record fragments faster than we can preserve or trust it.

We face three interrelated challenges:

1. **Continuity of Understanding** – AI systems remix and reinterpret existing content at scale. Provenance must provide interpretive lineage - how meaning and authorship evolve through these transformations.
2. **Durability of Memory** – Content formats, platforms, and institutions will fail. Provenance must ensure that historical accountability and context endure beyond the medium itself.
3. **Civic Accountability** – Provenance is not just a technical layer; it is a civic one. In a world saturated with synthetic media, digital provenance provides a basis for inspectable truth and collective sensemaking.

These pressures demand a broader design lens - one that spans creation, transformation, and preservation across both human and machine agency.

By treating provenance as civic infrastructure, we create conditions for a more transparent, accountable, and memory-rich digital civilization. The challenge ahead is not merely technical but cultural: building a system that invites participation from authors, institutions, and algorithms alike.

Rationale for Inclusion

Element	Why It Matters	What It Solves
Archival Provenance	Records preservation acts, migrations, and re-encodings as part of provenance.	Prevents loss of integrity during digital transitions; supports long-term cultural memory.

Relational Provenance	Maps the contextual relationships between works and claims.	Enables networked understanding, structured disagreement, and traceable remix culture.
AI Participation	Makes artificial agency visible and accountable within provenance graphs.	Clarifies authorship, reduces epistemic drift, and supports responsible synthetic creation.
Meta-Layer Integration	Provides the augmentative layer for creation, display, and composition of provenance.	Enables real-time inspection, modular overlays, and accelerated adoption beyond browser constraints.

Digital provenance should not only describe what was made, but how it is sustained and interpreted over time - by humans, machines, and the institutions that preserve meaning. The Meta-Layer can make this evolution much more visible, composable, and civic.

The Desirable Properties of Digital Provenance

To improve clarity and structure, the properties below are grouped into thematic categories that correspond to core dimensions of trust: Integrity Properties, Interpretive Properties, and Governance Properties.

The Digital Vellum project seeks to ensure digital objects remain authentic and interpretable across centuries. This note identifies the essential properties of digital provenance that underpin that mission and support verifiable preservation through cryptographic, semantic, and institutional means.

Each property reflects a principle of trustworthy digital lineage and may raise trade-offs across usability, privacy, scalability, and incentive design.

Integrity Properties (Truth & Transport)

1. **Verifiability** — Provenance data must be signed or cryptographically anchored to allow independent validation.
2. **Attribution** — Assertions should link clearly to their originator(s), whether individual, institution, or system.
3. **Temporal Anchoring** — Provenance should include durable, tamper-evident timestamps.
4. **Scope Binding** — Assertions must specify *what* they apply to (e.g. URI, hash of a media object, fragment of a page, or a meta-domain).

5. **Addressability** — Provenance data should be directly addressable or retrievable via predictable, well-known locations. This includes support for linking, resolving, or embedding provenance using established web identifiers (e.g. fragment IDs, well-known URLs, or resolvable hashes), enabling easy discovery and inspection
6. **Portability** — Provenance should travel with or be resolvable from content as it moves across formats, contexts, and realities.

Interpretive Properties (Meaning & Context)

7. **Inspectability** — Users and systems should be able to retrieve and interpret provenance in context. Ideally, its presence is intention- or attention-triggered.
8. **Composability** — Provenance data should support layering (e.g. remix chains, endorsements, disputes).
9. **Challengeability** — Third parties should be able to issue counter-assertions or challenges, without deleting or overwriting prior claims.
10. **Legibility** — Provenance must be human-readable and interpretable. Technical validity and machine readability alone are insufficient; users and non-technical stakeholders should be able to understand what is being asserted, by whom, and why it matters. This legibility supports meaningful reuse and elevates the potential historical and cultural value of provenance artifacts. Legibility should support interoperable formats such as JSON-LD to enhance both machine and human interpretation.

Governance Properties (Rights, Incentives & Safety)

11. **Permissionless/Decentralized** — Systems should allow expression of provenance without central gatekeepers; e.g., reputation without monopoly.
12. **Privacy-Respecting** — Provenance mechanisms should avoid passive surveillance or forced identity leakage.
13. **Immutable but Updatable Assertion Stacks** — Provenance systems should allow assertions to be recorded in a **tamper-evident, append-only manner**, while enabling **updatable interpretations** (e.g. withdrawn claims, renamed identities, opt-out of display). This ensures the historical trail is intact, interfaces can honor evolving identity and consent, and the full state is resolved via stack traversal, not overwrite.
14. **Incentivizability and Tradeability**— Provenance systems must offer aligned incentives for adoption across stakeholders, including authors, platforms, and users—and potentially pave the way for new forms of economic participation through provenance artifacts that preserve and reflect the value of verified digital creation. These systems should provide tangible value, low-friction integration, and compatibility with existing editorial and governance workflows. Importantly, provenance should not only facilitate trust; it should **create durable value**. When provenance is verifiable and transferable, it can function as an economic primitive. Those who generate content with strong provenance should be able to **retain and benefit from the intrinsic value** of both the content and its verified history. Comps include fine art, gems, and collectibles whose value is assured by certification of its provenance or fidelity.

Expanding the Framework of Digital Provenance

Each of the following subsections builds upon the previous, tracing a logical progression from preservation (Archival) to contextual linkage (Relational), to the role of intelligent systems (AI Provenance), and finally to the interface that makes all of it visible and usable (Meta-Layer). This connective framing clarifies how these dimensions collectively form a continuum of digital trust and memory.

Provenance as Living Infrastructure

Digital provenance can be more than a static record. Humanity needs a living infrastructure that binds content to its origins, transformations, and contexts over time. Beyond ensuring authenticity, provenance supports continuity of interpretation and meaning. This requires systems capable of:

- Preserving assertions across technological epochs
- Maintaining verifiability through cryptographic continuity
- Providing interpretive legibility to humans and machines alike

Such systems can help prevent the recurring erasures that occur with each digital format transition. Provenance becomes a form of *semantic durability*—a civic and technical memory layer for the internet.

Important uses of digital provenance systems include:

- **Preservation** — Capture and freeze webpages at the moment of publication or over time.
- **Dispute resolution** — Provide authoritative snapshots for legal, journalistic, or civic conflicts.
- **AI grounding** — Supply context and history so models can distinguish old from new, authentic from manipulated, human from machine.
- **Education & research** — Give historians, scholars, and students structured access to evolving knowledge trails.
- **Civic trust** — Enable citizens to see who said what, when, and how claims have been challenged or supported.
- **Economics** — Provide basis for new knowledge-based industries and markets.

The Expanded Field of Provenance Actors

Traditional provenance focused primarily on authors and publishers. However, a resilient ecosystem now includes:

- **Authors and Creators** — who initiate assertions of authorship and contextual intent.
- **Publishers and Endorsers** — who provide distribution, curation, and validation.

- **Archivists** — both institutional and distributed, who ensure temporal persistence and access.
- **AI Systems** — which now act as authors, publishers, and archivists.
- **Attestors and Validators** — who verify claims, check signatures, and mediate disputes.
- **Context Builders** — who form relational provenance by connecting artifacts across contexts.

This expansion recognizes that provenance is not only about *creation*, but also *interpretation* and *preservation*.

AI Provenance: Machines as Participants in Memory

As artificial intelligence systems increasingly act as creators, curators, and preservers of digital content, provenance models must evolve to represent their participation with the same precision afforded to humans and institutions. AI agents can operate in multiple roles:

- **AI as Author** — When an AI system generates content, the provenance should include machine-readable metadata identifying the model, its version, the operator or controller responsible, the generation timestamp, and content hash. This allows transparency in authorship while maintaining accountability for those who deploy or fine-tune the system.
- **AI as Publisher** — Systems that automatically distribute, summarize, or repurpose existing works should include metadata indicating source provenance references and transformation context. Fields should note model lineage, applied prompts or filters, and intent of redistribution.
- **AI as Archivist** — Autonomous or semi-autonomous systems may perform archival actions, such as snapshotting or verifying content integrity over time. These events should be recorded as verifiable provenance acts, including system identity, validation method, and integrity proofs.

Archival as a Provenance Act

Every act of archiving is itself a provenance event: a verifiable statement that *this* version of content existed *then*. Third-party archiving—by entities such as the Internet Archive, Starling Lab, or Harvard LIL’s Scoop Capture Engine—should be considered part of provenance infrastructure. As Michael Witmore notes:

Most important for long term historical interpretability is nearness of provenance event - "signing" with a persistent ID, for example - to the agent who initiates the event. Wax seals on historical documents and indentures (wavy cuts on vellum documents that assist in authentication) are good historical examples.

This nearness suggests the most desired attestor of provenance is the author or publisher, provided attestation at the moment of the provenance event. While third-party is less desirable, it can still provide significant value even if it is a proxy or approximation. And because there are already trillions of webpages (albeit most of which are not worth creating provenance for), for

the sake of meaning making, we need to assess the provenance as best as we can for millions if not billions of pages within any emergent system. To only enable authors and publishers going forward would leave too many gaping holes in our knowledge. It also precludes a huge opportunity to enable the immense creation of value and provenance artifacts for billions of content shards that have already been created.

Archival metadata should include:

- Source URL or content hash
- Timestamp of capture
- Archiving agent or institution
- Storage location or content-addressed identifier (e.g., IPFS CID)

These attestations not only ensure access and durability but also create accountability for what is preserved and how.

Relational Provenance and Bridges

Relational provenance is the connective tissue between information within and across documents for which digital provenance is available. These relational provenance artifacts capture how specific content - such as claims, evidence, or versions - relate to one another. For example, within a document with provenance, a specific assertion may be cited or sourced from, contradicted by, or corroborated by another source, providing the basis for a graph that connects content from different sources. Relational provenance connects artifacts across context, linking commentary, remix, rebuttal, and reuse into an interpretable network.

These *provenance bridges* can be authored by creators, third parties, or AI agents, using standardized schemas to express:

- Reference relationships (“derived from”, “responds to”, “endorses”, etc.)
- Semantic or thematic continuity
- Disagreement or revision

Relational provenance turns provenance from a static property into a dynamic, compositional graph of sensemaking.

The Role of the Meta-Layer

The meta-layer is not the subject of this note but serves as an *accelerant*. It augments digital provenance by:

- Enabling overlays that visualize provenance dynamically across the web
- Providing composable environments for governance and verification
- Enabling users to configure what provenance signals they trust or surface

While the *foundational* system must remain browser-compatible and protocol-based, the meta-layer provides the human and civic interface for interacting with provenance at scale.

An Internet-Draft called *The Meta-Layer: A Coordination Substrate for Presence, Annotation, and Governance on the Web* was published on October 5, 2025 [Mohamed & Benjamin, 2025]. These concepts build on the architectural vision outlined in *The Metaweb: The Next Level of the Internet* (Bridgit DAO, 2023), which introduced the concept of a “meta-layer above the webpage” as a civic and computational trust substrate.

Implementation Considerations

This section moves from conceptual design to practical realization, showing how the principles outlined earlier can be implemented in web-native contexts. As discussed below, JSON-LD provides the connective format linking abstract provenance concepts to concrete, interoperable implementations, ensuring that provenance artifacts remain both machine-verifiable and human-readable.

To demonstrate how these properties might manifest in practice, the following outlines an exploratory architecture and data model. These examples are offered not as fixed standards, but as conceptual scaffolding for feedback, testing, and refinement.

Lightweight Provenance Artifacts

- Structured using JSON-LD for linked data compatibility and web-native integration as mentioned above.
- Include fields for: `@id`, `creator`, `timestamp`, `scope` (e.g., hash or URI), `claim`, `signature`.
- Can be embedded in HTML, stored as standalone documents, or referenced through external resolvers.

Interaction Roles

- Authors: Embed signed assertions of authorship and context at creation time.
- Publishers: Endorse, archive, or contextualize content through signed publication or distribution events.
- Attestors: Independently verify or co-sign provenance assertions on behalf of others, ensuring accountability.
- Archivists: Capture, timestamp, and preserve digital artifacts as verifiable historical records.
- Resolvers: Browser extensions, middleware, or hosted services that retrieve and render provenance overlays.

Key Workflows

- Signed authorship of original content (blog posts, scientific preprints)

- Provenance trails for AI-generated or remixed media. As Vint Cerf rightly pointed out, this could trigger combinatorial explosion when applied to computations on or mixing of multi-source data. While this could be a problem, states are beginning to legislate AI provenance. China requires explicit and implicit labels for AI-generated text, images, audio, video and other virtual content on social media [Feng, 2025]. Absent such a policy, a market could be seeded that ultimately decides what gets provenance.
- Institutional endorsements or disputes
- Third-party challenges layered as overlays or counter-assertions
- Contextual integrity for content fragments or reposts
- Attestors issuing provenance on behalf of others to protect anonymity or ensure credibility
- Select provenance signals
- Review provenance signals

Deployment Surfaces

- Inline HTML Tags (e.g., `<script type="application/ld+json">` or `<link rel="provenance" href="...">`)
- Well-Known Endpoints (e.g., `/.well-known/provenance`)
- Metadata Registries (for search, validation, aggregation, and challenge resolution)

With server access, well-known endpoints can be the single source of truth, linked to from inline HTML tags. Absent server access, the only viable deployment surface is overlays that access metadata registries.

Validation and Display Layers

- Cryptographic Validation: Signature and timestamp checks ensure authenticity and integrity.
- Threshold Consensus: Quorum or weighted systems can surface contested claims, with user-adjustable visibility to explore signals beyond defaults.
- Civic Overlays: Provenance metadata, disputes, endorsements, and challenges can be surfaced via composable overlays in browsers or the meta-layer for public inspection.

JSON-LD as the Backbone of Provenance Artifacts

JSON-LD (JavaScript Object Notation for Linked Data) serves as a lightweight, machine-readable format that links data semantically while remaining compatible with the traditional web stack [W3C, 2017]. In the context of digital provenance, it becomes the throughline for expressing who created what, when, in what context, and with what endorsements or transformations.

This structure doesn't just document what happened - it makes the context of digital actions verifiable, portable, and composable across time and systems.

Trust: Establishing Inspectable Digital History

Just as climate networks use JSON-LD to trace emissions data through supply chains, provenance artifacts encoded in JSON-LD enable web-native trust in authorship, timestamp, and transformation. When authors, institutions, or validators sign these artifacts, they create tamper-evident assertions that:

- Build public memory
- Resist manipulation
- Enable structured disagreement (via overlays or challenge graphs)

Trust, here, becomes a function of transparent lineage, not centralized adjudication.

Commerce: Valuing the Artifact and the Proof

Like carbon credits or verified offsets, provenance artifacts are more than metadata—they're economic primitives. A signed JSON-LD record that links a video, blog post, or AI-generated image to a specific author and context transforms that media into a verifiable, ownable unit of value.

- Creators can monetize the content plus provenance
- Publishers gain durable trust capital and traceability
- Validators and indexers emerge as new economic agents

This creates a feedback loop where embedding provenance becomes a financially rational choice, seeding a new market for digital credibility and cultural inheritance.

Transparency: Public Context as Civic Infrastructure

Provenance isn't just about individual trust - it's about collective legibility. JSON-LD allows assertions to be linked across domains and identities without requiring a single gatekeeper. These artifacts support transparency with pluralism and nuance by enabling:

- Decentralized validators
- Threshold-based visibility (only surface if corroborated)
- Consent-aware display of historical assertions

Over time, this could ground algorithmic governance, content moderation, and archival policy in inspectable context rather than opaque judgment.

A Trust and Memory Rail

In climate, JSON-LD enables a flow of accountability. In provenance, it enables a flow of meaning and memory. As digital content becomes more ephemeral, remixable, and

AI-generated, the provenance artifact becomes a new anchor of truth - not because it enforces truth, but because it invites inspection, context, and interpretation.

This scaffolds a future where the web isn't just a publishing medium - it's a living ledger of cultural, creative, and civic significance.

Incentive Models and Business Cases

This section explores the immediate and emerging incentive structures that support the creation, validation, and maintenance of digital provenance systems. It outlines how these incentives can align with existing publishing, archiving, and AI workflows while opening pathways to new business models.

Emerging Business Models

- **Provenance-as-a-Service:** Toolkits and APIs for embedding, verifying, and managing metadata.
- **Validator Collectives:** Distributed trust networks for semantic or community-based attestation.
- **Reputation Graph Platforms:** Marketplaces aggregating provenance signals.
- **Legal & Archival Services:** For notarized content registration, dispute resolution, and historical memory.
- **Provenance Brokers:** Routing provenance for dynamic media and AI-generated content.

Provenance Artifacts as Economic Instruments

During the Future of Text symposium in Vancouver, Washington in November 2024, Vint said digital preservation needs a business model [Future of Text, 2024]. That thread is alive in this note. The idea of provenance artifacts emerges as signed, portable assertions that can carry not just truth claims, but also historical and cultural value.

Provenance artifacts are not just metadata; they're *verifiable, structured attestations* tied to meaningful digital moments. This could include:

- Original authorship signatures
- Community endorsements or challenges
- Certification of originality, timing, or human authorship
- Traceable remix lineage or interpretive context

They give creators and publishers the incentive to embed provenance by enabling them to earn not just the value of their creations, but also the value of their verified, enduring provenance - creating a feedback loop that supports adoption and trust. Digital provenance sets the stage for durable digital preservation. We're beginning to imagine these as economic primitives, anchoring future markets for memory, integrity, and trust.

There are many benefits across stakeholders including:

- **Authors** gain control over the historical framing of their work.
- **Publishers** acquire durable integrity signals and editorial metadata.
- **Third parties** (validators, communities, historians) can contribute value without needing to rewrite or remove source material.

Digital provenance is envisioned as a foundation for durable digital preservation, trust, and memory. Signed, portable assertions can carry truth claims and historical or cultural value. Creators and publishers gain incentive to embed provenance by earning not only from their creations but also from their verified enduring provenance.

Related marketplace opportunities include:

- **Proof-as-Asset** — Tokenized assertions of authorship or endorsement. Creators or communities might tokenize or license provenance (e.g. as NFTs or verified claims), creating secondary markets for:
 - Early-signed versions
 - Verified origin stories
 - Endorsements by known experts
- **Reputation Bundles** — Curated sets of validated claims for discovery and ranking. Reputation providers could monetize curated, high-trust collections of provenance, e.g. “verified medical claims,” “certified human-created art,” etc.
- **Legacy & Rights Management** — Verified trails supporting digital estates and academic attribution. Provenance artifacts could underpin wills, archives, or copyright claims, creating demand for notarization, brokerage, and arbitration services.
- **Indexing Platforms** — Discoverability services adding contextual value to provenance-rich media. Services that surface high-integrity provenance may become search and navigation layers in themselves—especially useful in AI and education contexts.

From Metadata to Market: A Shift in Ontology

When provenance data moves from being *about* content to being *independent of it*, it crosses into the realm of assets. These artifacts don’t just describe something - they *are* something: signed, time-stamped, relationally positioned, and potentially valuable as units of verified history.

A provenance artifact is thus:

- **Durable:** it persists beyond the lifespan of the content it references.
- **Addressable:** it can be cited, linked, and aggregated.
- **Auditable:** it can be verified by third parties.
- **Compositional:** it can connect with other artifacts to form interpretive chains.

That means provenance itself can become a market primitive, not just infrastructure for trust, but a generator of value in its own right. Provenance artifacts are knowledge-based, real world assets (RWA).

Markets for Provenance Assets

Note on Market-Driven Archival Selection *A functioning market for provenance and archival artifacts removes the burden of deciding what is worthy of preservation. Rather than relying on centralized committees, institutional whims, or arbitrary criteria, markets naturally surface which digital objects accrue value - cultural, historical, legal, or interpretive. High-signal artifacts (or those that become valuable over time) attract attention, validation, and preservation services, while low-signal material fades without requiring deliberate curation.*

This does not diminish public archives; instead, it augments them with a discovery and prioritization mechanism that reflects real-world use, scholarly interest, and cultural evolution. In this model, the market handles selection, allowing provenance ecosystems to scale without agonizing over what to save.

It would be illuminating to give creators the option to immortalize their provenance or archival record for an important cultural moment as a tradable digital asset.

Imagine provenance artifacts functioning like publicly tradable historical anchors. As archives, institutions, or AI systems produce signed records of digital states (“this page existed at this time, in this form”), those attestations become reference points for future knowledge work, legal, academic, cultural, and machine learning applications alike. Provenance artifacts can underwrite the foundation of recorded knowledge, thereby giving us first glimpse at the emerging digital noosphere.

Potential use cases:

- **Verification as a service:** Networks of attestors offering provenance-backed validation of historical digital states.
- **Archival derivatives:** Bundles of verified snapshots or datasets forming historical indexes, offered as research or training corpora.
- **Cultural artifacts:** These are the original works or their verified digital or physical archives that possess cultural, historical, or artistic significance. They serve as enduring records of creative or societal moments, preserved for study, interpretation, and collective memory.
- **Reputation marketplaces:** Provenance-linked endorsement histories that quantify credibility, originality, or authenticity.
- **Provenance collectibles:** Early or canonical provenance artifacts (e.g. “first signed capture of the Wikipedia homepage”) treated as heritage assets or NFTs of record.

This isn't about monetizing data - it's about valuing continuity, integrity, and context. It's also about creating value and impact. It's about harnessing the powers of the market to bring to life a self-funding civic memory apparatus.

Markets for Relational Provenance Artifacts

Relational provenance - the connective tissue between artifacts - introduces network effects and interpretive markets.

For instance:

- **Interpretive Networks:** Scholars, journalists, and AI systems may pay for access to curated provenance graphs connecting events, claims, and sources.
- **Relational Derivatives:** Machine-readable networks of influence (who cited whom, who challenged whom) become assets for sensemaking, knowledge analytics, policy forecasting, and AI training.
- **Trust Graph Curation:** Third parties may curate "provenance trails" that bundle verified relationships into high-value trust fabrics, effectively selling *context as a service*.
- **Cultural Value Markets:** Collectors may buy and trade relational provenance artifacts that have cultural significance.

Each bridge or link becomes a micro-asset: the provenance of provenance. And because they are signed, timestamped, and attributable, these relational artifacts can themselves accrue credibility, reputation, or market value over time.

Legacy and Rights Management

Once provenance becomes durable and tradeable, legacy management evolves from ownership of content to stewardship of context.

Creators and estates might:

- License the right to reference or recontextualize provenance artifacts.
- Establish memory trusts that maintain provenance graphs as part of cultural or civic inheritance.
- Negotiate royalties for derived or recombined provenance graphs (e.g., when AI models use provenance trails to generate insights).
- Buy and sell provenance artifacts based on their market-based cultural value.

In this paradigm, intellectual lineage - not just intellectual property - becomes the asset.

The Civic Layer: Provenance as a Commons

As a "civic substrate of public memory," this is not speculative tokenization, but market-based memory stewardship.

A civic provenance economy would:

- Balance open access with verifiable authorship.
- Support public institutions (libraries, archives, consortia) as memory validators.
- Enable a commons-based infrastructure for cultural continuity where value accrues not to scarcity but to integrity and contextual contribution.

This is how provenance could become the backbone of a public memory economy that monetizes accountability and interpretation, not attention.

Additional Considerations

Open Questions

- How can provenance remain durable in high-churn or low-incentive environments?
- Can integrity be preserved without compromising editorial freedom or anonymity?
- What fallbacks exist when provenance is absent or disputed?
- How should markets for provenance artifacts be governed to prevent abuse?
- What happens when provenance is missing, manipulated, or ambiguous?
- How do we handle name changes, retractions, or pseudonym?

Next Steps

This note invites feedback from collaborators across the Digital Vellum working group. Specific focus areas for further refinement include:

- Provenance metadata schema alignment with C2PA, JSON-LD, W3C annotation standards
- Feasibility of browser-level auto-detection and overlays
- Governance and incentive structures for attestors and archival networks
- Business models for digital provenance

Ultimately, we would like to appropriately limit the scope and submit one or more Internet-Drafts to IETF with respect to digital provenance anchoring a civic memory system, which would require a higher level of specificity. This could include:

- 1. Reference Architecture Sketch**
 - a. Layers: Identity → Assertion → Validation → Resolution → Display
 - b. Components: Provenance artifact, Resolver service, Viewer module, Validator nodes
- 2. Data Model Examples**
 - a. Sample JSON-LD for a provenance claim
 - b. How to embed, overlay, and externally link to it from HTML or PDF
- 3. Flow Diagrams**
 - a. How a claim gets authored, attested, retrieved, and challenged

- b. Trust layering: author + publisher + community + counter-claims
- 4. Interoperability Considerations**
 - a. Canonical structures for time, scope, and authority
 - b. Use of DIDs, content hashes, and well-known URIs
- 5. Roles and Responsibilities**
 - a. Who are the actors? (authors, attestors, resolvers, challengers)
 - b. What accountability trails are expected?
- 6. Risks and Edge Cases (TBD)**

The next step is to do a reference implementation of digital provenance within an application currently under development, to be detailed in another DVN.

Digital Vellum Note #3 marks a step toward a more complete understanding of provenance as both *infrastructure* and *ethic* for the web's enduring civic memory.

This note is formatted and numbered in accordance with DVN #1 and is ready for inclusion in the Digital Vellum series registry.

REFERENCES

1. Mohamed, K.A. & Benjamin, D., The Meta-Layer: A Coordination Substrate for Presence, Annotation, and Governance on the Web, 2025.
<https://datatracker.ietf.org/doc/draft-meta-layer-overview/00/>
2. Bridgit DAO, [*The Metaweb: The Next Level of the Internet*](#), Taylor & Francis / CRC Press, 2023.
3. JSON-LD 1.1 Specification, W3C Recommendation (2020).
4. Feng, C., China's social media platforms rush to abide by AI-generated content labelling law, South China Morning Post, 9/1/2025.
<https://www.scmp.com/tech/policy/article/3323959/chinas-social-media-platforms-rush-a-bide-ai-generated-content-labelling-law>
5. Future of Text Symposium, Vancouver, 2024 (proceedings).

ERRATA

None at issuance. Subsequent corrections or interpretive additions should be issued as DVN #3.x or later.