

# Meaning That Survives Change

## *The Rosetta Pasigraphy Protocol: Conservation, Semantic Continuity, and Compositional Assurance*

Crates McDade | Entif.AI

ETR-2026-04 | Version 0.2 | September 9, 2026

**Architecture position paper and research agenda, with an executable finite reference model. Not peer reviewed. This publication is not a Rosetta protocol release or a production certification.**

### Abstract

Information reused by AI systems crosses representations, institutions, policies, and time. Preserving bytes and recording transformations do not, by themselves, preserve the distinctions a later decision requires. The Rosetta Pasigraphy Protocol proposes an open semantic and provenance substrate that separates observations, interpretations, operational activity, evaluations, and authority. This paper sharpens its research proposition: transformations should not silently manufacture meaning, identity, causality, permission, security, effects, corroboration, certainty, or currency. We organize these obligations as proposed conservation invariants and distinguish local validity from compositional assurance. To make the continuity argument falsifiable within a stated boundary, we specify a finite mapping fragment with set-valued interpretations, joint task observables, partial migration functions, explicit loss, and path-relative confluence. A companion implementation supplies 48 passing bounded checks, including 16 migration fixtures, while a source-extracted probe identifies a canonicalization mismatch in the inspected prototype. These are limited artifact-level results, not validation of Rosetta's broader architecture. A candidate Text-Core research profile freezes a small exchange surface without ratifying new Core semantics. We relate the proposal to provenance, research-object packaging, schema mapping, abstract interpretation, identity reconciliation, and access-control literature, and connect three companion research papers to distinct prospective use cases. The evaluation agenda compares Rosetta with disciplined compositions of existing standards and specifies results that would weaken its case. The objective is not a truth oracle, but independently inspectable knowledge whose transformations, uncertainty, and permitted use remain explicit as systems change.

**Keywords:** semantic interoperability; provenance; schema evolution; compositional assurance; AI systems; research infrastructure; rights-aware retrieval; formal reference models.

## 1. The proposition and its limits

A laboratory records a result as *not confirmed*. Years later, an analysis expects one of three states: *measured negative*, *inconclusive*, or *unmeasured*. The old file remains intact. Its hash still verifies. Every transfer has a receipt. Yet assigning the old result to any one of the new categories invents information. A protocol can preserve this uncertainty, report that a requested comparison is unsupported, or conceal the distinction behind a convenient default. Only the first two are defensible without additional evidence.

This example expresses Rosetta's central research problem: **what must remain invariant, and what must be declared as changed, when knowledge is represented, transformed, exchanged, and reused?** The proposed answer combines exact artifact identity, explicit interpretation, provenance, scoped authority, and version-aware compatibility. The Core Spine is the project's semantic design authority; it is not evidence that every specified capability is implemented. [1]

Rosetta and Entif.AI have different roles. Rosetta is the open protocol and representational commons. Entif.AI is the research and engineering effort developing implementations and higher-level systems around it. An independently useful public contract must not require an Entif-hosted service or undisclosed decision procedure merely to determine what an artifact means or whether it satisfies that contract. The public governance doctrine makes this separation explicit. [27]

This revision makes four bounded contributions. It gives an architecture-level account of the separation between integrity, identity, semantics, justification, and authority. It supplies a finite, executable instantiation of semantic preservation and loss, with assumptions that another implementation can challenge. It organizes the newly documented composition risks into ten proposed conservation obligations. Finally, it provides a candidate text exchange profile and a comparative evaluation agenda, together with source, claim, revision, and test records.

The mathematical fragment is deliberately modest. Its set inclusion arguments are not claimed as new mathematics. The broader hypothesis is that a useful combination of public contracts can make unsupported conclusions and inappropriate reuse harder to conceal across system boundaries. That hypothesis remains unmeasured. A simpler combination of existing standards may achieve the same benefit with less cost.

Rosetta does not establish objective truth, guarantee complete capture of events, expose a model’s hidden neural computation, eliminate disagreement, or turn a compatible operation into an authorized one. A correct record can describe a mistaken interpretation or an unjust decision. A reversible representation transform cannot reverse a settled external transaction. These are limits on the proposition itself, not implementation details to be postponed until deployment.

2. Architecture: distinctions that must not collapse

2.1 Five different questions

The architecture is easiest to understand as five questions that require different evidence. Integrity asks whether a specified representation changed. Identity asks which continuing entity or source a record denotes. Semantics asks what a representation asserts under a particular interpretation. Justification asks what warrants accepting that assertion. Authority asks who may make a decision or perform an action. Answering one does not answer the others. [1]

| Layer         | A defensible question                               | An inference it does not license             |
|---------------|-----------------------------------------------------|----------------------------------------------|
| Integrity     | Do these canonical bytes match this digest?         | The statement is true.                       |
| Identity      | What entity and exact revision are referenced?      | Two similar records denote the same entity.  |
| Semantics     | What is asserted under this Profile and vocabulary? | The interpretation is empirically justified. |
| Justification | What evidence supports this assessment?             | The assessor is authorized to act.           |
| Authority     | Is this actor allowed this action now?              | Every downstream use remains permitted.      |

## Different evidence for different questions

| INTEGRITY                   | IDENTITY                         | SEMANTICS                          | JUSTIFICATION                    | AUTHORITY                         |
|-----------------------------|----------------------------------|------------------------------------|----------------------------------|-----------------------------------|
| Did the bound bytes change? | Which entity and which revision? | What does this representation say? | What evidence supports it?       | Who may do what, now?             |
| Not: the claim is true.     | Not: similarity proves sameness. | Not: its meaning is justified.     | Not: evidence grants permission. | Not: every future use is allowed. |

Architectural reading guide. Not a deployed verification stack.

**Figure 1. Five non-substitutable questions. The relationships are an architectural reading guide, not a deployed verification stack.**

A Tile is an addressable artifact. A content identifier (CID) identifies the exact representation inside a declared canonicalization boundary. An RID supplies continuing logical identity across revisions. External identifiers remain qualified anchors to external authorities, not Rosetta-created replacements for those authorities. Sharing an RID is an identity assertion; it is not a proof that an entity-resolution decision was correct. [1]

The phrase *exact representation* needs care. The inspected prototype hashes a canonical body containing kind, pack, version, parents, and payload. It excludes the enclosing `createdAt`, `cid`, and `canonical` fields. Consequently, its CID does not authenticate every byte or every metadata field in the envelope. Changing `createdAt` alone leaves that CID unchanged. Security-relevant time claims must therefore be bound inside an appropriate authenticated subject rather than inferred from that convenience field. Section 8 distinguishes this inspected implementation boundary from the protocol’s intended semantics. [4]

Two historical Tiles can retain a logical association while having different CIDs. A correction should create an attributable successor or derived projection, not silently reinterpret an older CID. Revision lineage and derivation lineage answer different questions: what replaced this record, and what evidence or inputs produced it? Preserving both permits later reviewers to reconstruct a past interpretation without treating it as currently authoritative. [1,3]

## 2.2 From observations to interpretations

The Core Spine separates observed material from its interpretation. Observations retain encountered signals. Forms locate structured occurrences such as tokens or image regions. Lexemes represent language-bound lexical senses. Pasigrams, Concepts, Frames, and typed relations represent language-neutral semantic structure. Conjectures retain candidate interpretations and their support when the evidence does not justify a single choice. [1]

An occurrence of *bank* is not already a dictionary sense, and a selected dictionary sense is not already a particular institution or riverbank. A defensible pipeline keeps the source occurrence, interpretation, and entity association distinguishable. Human-readable accompaniment is a derived explanation of those artifacts, not a source of new authority. A fluent paraphrase should not erase an unresolved alternative or outrank the record from which it was generated.

Pasigraphy here is an ambition for structured, interlingual representation. It is not a claim that every language divides experience identically or that a parser has access to a universal ontology of correct meanings. Demonstrating fidelity requires specific languages, tasks, annotation rules, and counterexamples. The candidate profile in this package does not perform general natural-language understanding.

## 2.3 Provenance is multidimensional

A source is not one indivisible fact. The Source Substrate design distinguishes records, manifestations, packages, acquisition events, publication context, custody, authorship claims, external identifier bindings, stewardship, rights, evaluation, and lifecycle change. These facets can be supplied by different authorities and can change independently. The addendum is design evidence; the repository separately describes which source artifacts are represented in its fixture-backed prototype. [2,28]

A repository's custody claim does not establish authorship. An identifier registration does not establish the correctness of a scientific result. Discoverability in an index does not make that index the publisher. These are analytical consequences of keeping the facets separate. A source bundle should permit disagreement about an authorship link while preserving an undisputed manifestation and acquisition event.

Identity resolution is especially consequential. Exact byte duplication, manifestations of a source record, continuity within a record family, and conceptual similarity must not be collapsed into one automatic merge operation. An inferred association needs its own evidence, scope, method, and correction history. False merges can contaminate many descendants; undoing an association is therefore different from deleting one duplicated file. [2,28]

## 2.4 Operational evidence and bounded use


The operational spine distinguishes Run, Action, ToolCall, Observation, and Evaluation. A Receipt is an attestation about a subject and supporting evidence under a verification context. It is not a universal replacement for all of those roles. A valid signature can establish an attestation under a key and policy; it does not establish truth, complete observation, uncompromised authority, or permission for every subsequent use. [1]

A Tapestry is a compiled working set under declared constraints. The near-term contract concerns exact membership, ordering, provenance, omissions, source spans, and rights-scoped retrieval. Promoting an extract into reusable state is a separate, attributable decision. It is not the same operation as generating an extract or describing it in English. The public issues for receipt families, promotion, and Tapestry retrieval retain these distinctions as open contract work. [3]

The execution trace can support inspection of a workflow without being a faithful account of hidden model computation. Experiments on unfaithful chain-of-thought explanations are a reason to evaluate generated rationales rather than grant them privileged causal status. Rosetta can identify that an explanation was generated and used; that record alone cannot establish why a model internally produced its answer. [26]

ENTIF.AI / ETR-2026-04

## One linked record, several kinds of claim



Illustrative architecture. Links do not collapse evidence, meaning, and authority.

Figure 2. Source evidence, semantic interpretation, and operational evidence remain separate but linked. Packs, Profiles, rights, and revision context constrain the connections. Illustrative architecture, not a product screenshot.

### 2.5 A small Core, versioned Packs, and declared Profiles

A Pack carries a versioned extension or mapping rather than forcing every domain distinction into Core. Standard Packs connect external specifications; Vocabulary Packs connect lexical and conceptual resources; schema-oriented Packs expose governed data contracts. A Profile declares a particular supported or required subset and its conformance conditions. Declaring support is not the same as demonstrating it. External standards retain their own authority, and a mapping must state whether correspondence is exact, narrower, broader, lossy, or unsupported. [1,3]

This architecture allows multiple revisions to coexist while translators and compatibility evidence evolve. It does not permit a Pack to silently redefine earlier Core meaning. Schema exports, catalog entries, runtime helpers, and conformance evidence are separate surfaces: a type name in code is not automatically a new semantic primitive, and an emitted validation shape is not proof that a full engine executed it. The public core-descent, schema, vocabulary, Pack-validation, and Profile issues own those distinctions. [3]

The first lexical path is deliberately bounded rather than a promise of universal multilingual coverage. Similarly, the planned Postgres/pgvector backend is an implementation target, not a requirement that every conforming implementation use that database. The scientific value of the public contract depends on replaceability: another implementation should be judged by its observable artifacts and relevant behavior, not by its choice of vendor or private optimization methods.

## 3. Positioning against existing work

Rosetta did not invent provenance graphs, content hashes, record linkage, schema migration, or retrieval augmentation. Its case must be evaluated against those foundations and against a disciplined composition of them. The question is whether its particular contract boundaries provide useful additional interoperability and assurance, not whether adjacent systems lack any provenance or versioning features.

**Provenance and research objects.** PROV-O provides an ontology for entities, activities, agents, and their relationships. FAIR articulates findability, accessibility, interoperability, and reusability rather than prescribing one technology. RO-Crate packages contextual descriptions of research objects; its versioned specification is

itself a useful example of explicit external authority. These are foundations and plausible baselines, not competitors to be dismissed by describing them as unstructured metadata. [6-9]

BagIt addresses reliable transfer and validation of file packages. Trusty URIs support verifiable, immutable digital artifacts in linked-data settings. Such approaches constrain Rosetta's novelty claim: stable identity and verifiable packaging already have substantial precedents. The remaining research question is whether a package's changing interpretation, current permission, and task-relative comparison can be handled coherently across the proposed layers. [10,29]

**Mapping composition and evolution.** Fagin and colleagues study composition of schema mappings and the expressive role of second-order dependencies. Yu and Popa study adapting mappings as schemas evolve. Data-exchange research also distinguishes possible target instances and the answers justified across them. This literature makes a mapping language, its semantics, and its computational assumptions prerequisites for serious claims about preservation. A migration arrow without those commitments is not a formal result. [11-13]

The finite fragment below follows the general tradition of reasoning through sets of possible concrete states. Abstract interpretation provides a much broader framework for sound abstraction and approximation. Our explicit finite sets and inclusion checks are a transparent teaching and testing device, not a replacement for that theory or a claim to solve general schema-mapping composition. [14]

**Identity reconciliation.** Record linkage has long distinguished matches, nonmatches, and cases requiring further resolution. Linked-data identity constructs such as `owl:sameAs` carry stronger implications than superficial similarity. Rosetta's RID/CID distinction helps state what kind of identity is intended, but it does not solve uncertain linkage, transitive merge effects, or reversal after mistaken association. The relevant associations still need contestable evidence and constrained use. [15,16]

**Policy, time, and execution.** ODRL supplies a model for expressing permissions, prohibitions, duties, and constraints. XACML separates access-control decision and enforcement responsibilities. Neither a policy document nor an allow-shaped record proves enforcement at every cache, log, model input, or remote replica. Rosetta must adopt an explicit enforcement boundary and test it. [17,18]

Causal ordering and temporal databases supply separate precedents. Lamport distinguishes causal precedence from arbitrary timestamp order. Temporal data management makes the treatment of changing facts and recorded history a first-class concern. Rosetta's proposed causal and bitemporal contracts should build on these distinctions rather than relabel wall-clock timestamps as causal evidence. [19,20]

**Retrieval augmentation.** Retrieval-augmented generation couples retrieval with generation, addressing a different part of the knowledge-use pipeline. Adding retrieval does not itself specify revision admission, promotion authority, migration loss, or distributed revocation. Conversely, a provenance protocol does not establish better retrieval or answer quality. A fair experiment gives each baseline equivalent sources, retrieval quality, and access controls before measuring any incremental benefit. [25]

Three baseline choices remain live: use existing standards without Rosetta; add a small interoperability Profile to those standards; or adopt a broader Rosetta implementation. The smallest adequate option is a legitimate positive outcome of this research, even when it weakens the case for adopting the complete architecture.

## 4. A finite, falsifiable continuity model

### 4.1 Scope and syntax

The earlier manuscript introduced representation states and preservation claims without fixing a language in which to test them. This revision supplies **ETR04-FINITE-0.2**, a paper-local reference fragment. It is not a normative Rosetta schema, a newly reserved Core kind, or a model of unrestricted human meaning.

Let  $\Omega$  be a finite, nonempty set of admissible worlds. A world is a fully specified record in the selected domain, not a metaphysical claim about reality. For each representation state  $R$ , let  $X_R$  be its finite set of record codes. An interpretation function assigns a nonempty set of worlds to every admitted code:

$$\gamma_R: X_R \rightarrow P(\Omega) \setminus \{\emptyset\}.$$

Empty interpretations are rejected as malformed in this fragment; they are not silently treated as vacuously supported records. A practical representation-state identity must bind the exact schema, mapping, vocabulary, relevant Core revision, and other interpretation dependencies. The executable fixture locks these choices by its contents and package digest; its short local names are not new production identifiers. The broader admission proposal in issue #1552 owns the production contract. [3]

A task Profile  $Q$  fixes a finite ordered collection of total observables. Together they form  $q_Q: \Omega \rightarrow A_Q$ . The admissible **joint** answers for a record are:

$$A_Q(R, x) = \{q_Q(w) : w \in \gamma_R(x)\}.$$

The word *joint* matters. If two binary variables are always equal, their possible joint answers can be  $\{(0, 0), (1, 1)\}$ . If they are always different, the set can be  $\{(0, 1), (1, 0)\}$ . Each variable separately admits both zero and one in both records. Marginal-answer equality would wrongly certify preservation of their relationship. The supplied correlation fixture rejects that inference.

A migration edge is a deterministic partial function  $f_e: D_e \rightarrow X_S$ , with  $D_e \subseteq X_R$ . In the machine grammar, this is a finite lookup table, a source state, and a target state. Membership in the table's domain is the applicability condition. There is no arbitrary program evaluation, probabilistic output, network access, or external side effect in the fragment. Missing rows mean *unsupported for this input*, not permission to guess a target code.

## 4.2 Preservation, approximation, and unjustified narrowing

For an admitted input  $x$ , define the edge to be **Q-sound** when it does not remove any answer that the source still admits:

$$A_Q(R, x) \subseteq A_Q(S, f_e(x)).$$

It is **Q-exact** when the two sets are equal. A strict superset is a sound but lossy approximation: it adds possibilities because some distinctions have been forgotten. The explicit loss record is the set difference between target and source joint answers. No universal percentage of semantic loss is inferred from its cardinality.

If a source-admissible answer is removed, the transform is **UNSOUND** in this fragment. This negative diagnostic includes both unjustified narrowing and changes that exchange one unsupported answer for another. It is not folded into *lossy*, because losing precision and manufacturing certainty are different failures. **UNSOUND** is a research diagnostic, not a claim that Rosetta has ratified that exact vocabulary.

An observation can legitimately reduce uncertainty when it contributes new evidence. That is outside an input-only migration. To represent it, the operation must identify the additional evidence and its derivation contract; it cannot describe a more certain output as a consequence of reformatting alone. Likewise, a transformation exact for one Profile must not silently advertise exactness for another.

These definitions are deliberately asymmetric. Adding possibilities is conservative with respect to the fixed oracle, but can make the result useless. Soundness is therefore not enough for task support. A consumer can require exactness on selected observables, permit a declared approximation, or refuse the result. The requested utility and authority conditions must be evaluated separately.

### 4.3 A worked assay example

Use  $\Omega = \{N, I, U, P\}$ , where  $N$  denotes a measured negative,  $I$  an inconclusive measurement,  $U$  an unmeasured case, and  $P$  a positive measurement. These are synthetic research states, not a clinical decision standard. The legacy code `not_confirmed` denotes  $\{N, I, U\}$ . A union-valued target can retain that set exactly. A translator that defaults it to  $\{N\}$  is unsound for the full status question.

Starting instead with a known negative, mapping  $\{N\}$  to `not_confirmed` is sound but lossy for the status Profile. It is exact for the narrower binary question *confirmed?*, since both source and target answer *no*. It is lossy for *measured?*, because the target admits both *yes* and *no*. Thus a single edge can have different defensible dispositions for different tasks.

| Synthetic case                                 | Full-status result | What the case establishes                                  |
|------------------------------------------------|--------------------|------------------------------------------------------------|
| Legacy uncertainty to explicit union           | EXACT              | Uncertainty can be preserved without guessing.             |
| Legacy uncertainty to negative default         | UNSOUND            | A schema-valid value can invent certainty.                 |
| Known negative to legacy uncertainty           | LOSSY              | Conservative approximation may discard usefulness.         |
| Two renamings with the same interpretation     | Confluent for Q    | Different code paths need not mean different answers.      |
| Direct exact path versus sound coarsening path | Nonconfluent for Q | Local soundness does not imply equal precision.            |
| No applicable table row                        | UNSUPPORTED        | Absence of a supported mapping is an attributable outcome. |

The familiar impossibility argument is elementary: if a transformation maps distinguishable source states to the same value, a deterministic inverse using only that value cannot recover both originals. More elaborate migration software cannot recreate distinctions that were never retained. The useful protocol behavior is to preserve the earlier artifact and report insufficient information, not to manufacture continuity.

### 4.4 Composition and confluence

A path is a finite sequence of edges whose source and target states align and whose domains contain the successive intermediate records. The reference checker validates every step. A later broadening operation does not rehabilitate an earlier unsound step or erase its diagnostic.

**Proposition 1, bounded composition.** Fix the same universe, interpretation functions, and joint Profile throughout a path. If every applicable edge is Q-sound, the composed path is Q-sound. If every edge is Q-exact, the path is Q-exact.


**Proof.** For a sound path, the answer sets form an inclusion chain  $A_0 \subseteq A_1 \subseteq \dots \subseteq A_k$ . Transitivity gives  $A_0 \subseteq A_k$ . For an exact path, each adjacent pair is equal, so the endpoints are equal. The result fails as a usable

certificate when a required precondition is absent or when the universe, oracle, or Profile changes without an explicit bridge. This is a basic property of set inclusion, not a new general theorem about language or distributed execution.

Two paths from the same source are **Q-confluent at that source** when their final joint-answer sets are equal. They need not produce equal bytes or even use the same target encoding. Q-soundness does not imply Q-confluence: one path can preserve  $\{N\}$  while another safely broadens it to  $\{N, I, U\}$ . Both avoid invented certainty, but they support different downstream questions. Figure 3 makes the distinction visible.

ENTIF.AI / ETR-2026-04

### Local soundness does not ensure confluence



**Both paths avoid invented certainty. Only one retains full precision.**

Synthetic finite example. Confluence is relative to a fixed task and oracle.

**Figure 3. A sound but nonconfluent diamond.** The direct path retains a known negative; the longer path forgets distinctions. Both are sound under the finite oracle, but their full-status answers differ. Synthetic reference example.

Cycles require the same discipline. A round trip preserving the full answer set can be exact for  $Q$  without restoring identical serialization. A cycle that only broadens sets can be sound while degrading usefulness. A cycle that narrows after coarsening must not be described as recovery of missing information. The suite includes exact-cycle, unsound-cycle, positive-diamond, and sound-but-nonconfluent cases.

A sparse graph can reduce the number of authored migration edges compared with direct maps between every pair of revisions. It does not make all paths valid or give a universal linear-cost migration solution. Path selection, preconditions, incompatible forks, metadata capture, and certification remain real costs. A historical comparison may be best performed in a mutually supported representation rather than the newest representation.

## 4.5 What this formalism does not settle

The oracle is an input to the study. The package does not infer admissible worlds from scholarly prose, decide whether an ontology is correct, or learn which questions a future researcher will need. Correct implementation of an incorrect interpretation table remains incorrect domain modeling. External validation must therefore examine both the interpreter and the adequacy of its supplied oracle.

Continuous measurements, distributions, and open-ended answers are not covered by the finite guarantee. A later study may use explicit equivalence classes, bins, intervals, or another bounded abstraction, but it must specify and validate that abstraction before borrowing the result. In particular, retaining a distribution's mean is not evidence of preserving its tails, and equal answer marginals do not preserve dependence.

The package reports **48 passing unit checks**, of which **16 are named migration fixtures**. One check examines all **3,375 triples of nonempty subsets** of a four-element universe; **175 triples** satisfy the sound-chain premises and satisfy their conclusion. These counts describe deliberately small checks performed in this revision. They are neither thousands of independent experiments nor evidence of real-world semantic accuracy. The reference code, inputs, outputs, and environment versions are included for reproduction.

## 5. Compatibility is a whole-process claim

### 5.1 Separate the scopes of the answer

An artifact’s compatibility disposition, a migration’s result, a process graph’s support status, and a Receipt’s evaluation verdict do different jobs. Their similar words must not be treated as interchangeable enums. The public proposals distinguish representation admission (#1552), migration projections (#1573), graph compatibility (#1553), and receipt-family semantics (#158). [3]

| Scope                      | Question answered                                                     | Why it cannot stand in for another scope                                 |
|----------------------------|-----------------------------------------------------------------------|--------------------------------------------------------------------------|
| Representation disposition | How may this exact historical state be interpreted or used?           | Admission does not prove a requested path exists.                        |
| Migration result           | What happened on this exact path, with what loss and replay evidence? | An exact transformation need not satisfy a process's other requirements. |
| Graph support              | Can every required process operate under the requested constraints?   | Compatibility does not authorize execution.                              |
| Evaluation or Receipt      | What was assessed or attested, by whom, under which policy?           | A pass verdict is meaningful only for its stated subject and scope.      |

The migration-result proposal includes exact, lossy, partial, ambiguous, insufficient-information, incomparable, unsupported, and replay-mismatch outcomes. The graph proposal distinguishes supported, supported with declared degradation, partially supported, unsupported, and unknown. These remain proposed public vocabularies, not newly accepted Core semantics. This paper’s research supplement supplies a small aggregation rule for its own fixtures; it does not claim that a general cross-scope composition rule is settled. [3]

### 5.2 Six dimensions of compatibility

A process can be structurally valid yet semantically inappropriate, behaviorally different, operationally unavailable, insecure, or unauthorized. At minimum, the research evaluation therefore records six dimensions: **structural, semantic, behavioral, operational, security, and rights**. Two versions may exchange the same amount field while disagreeing about retries, timeout ambiguity, event ordering, null semantics, settlement finality, or compensation. These are central concerns of the newer behavioral and live-migration proposals, not consequences of passing a schema validator. [3]

Consider a synthetic laboratory-to-repository-to-analysis workflow. Upload preserves the original result. A repository adapter preserves its bytes but maps *not confirmed* to a negative value. A downstream report is formatted correctly and signed. Every pairwise exchange can appear locally valid while the end-to-end claim about negative measurements is unsupported. An alternative adapter may retain the uncertainty, restoring semantic honesty without making that particular analysis answerable.

The research aggregator requires an explicit assessment of every required process in all six dimensions. Security, rights, revocation, and trust floors are hard constraints before any soft preference. A failed hard floor

yields unsupported. Known required-process failures yield unsupported or a diagnostic partially-supported status when some other required processes are known to work. An unresolved prerequisite yields unknown unless a known failure already defeats the request. Declared loss is acceptable only when the relevant process explicitly permits it. A supported result still has `authorizes_execution=false`.

ENTIF.AI / ETR-2026-04

## Compatibility has more than one dimension

| PROCESS           | STRUCTURE | MEANING | BEHAVIOR | OPERATION | SECURITY | RIGHTS  |
|-------------------|-----------|---------|----------|-----------|----------|---------|
| Source upload     | Pass      | Pass    | Pass     | Pass      | Pass     | Pass    |
| Legacy analysis   | Pass      | Lossy   | Unknown  | Pass      | Pass     | Pass    |
| Old-version retry | Pass      | Pass    | Fail     | Unknown   | Revoked  | Unknown |

**Hard security and rights floors precede soft fallback preferences.**

All cells are illustrative. No deployment, benchmark, or incident data is depicted.

Figure 4. Schema validity is only one compatibility dimension. All values in this matrix are synthetic and illustrate separate failure surfaces; they are not observations of a deployment.

### 5.3 Fallbacks, limits, and future revisions

Fallback is a graph-wide assignment problem. Replacing one integration revision can change the admissibility of downstream operations and shared dependencies. Returning the nearest-looking version is insufficient. Alternatives need a transparent loss/change vector and evidence about the exact compatibility corpus consulted. A security-revoked or rights-prohibited revision must not become a recommended fallback merely because its schema matches. [3]

A bounded search must distinguish *no solution found within the search budget* from *no solution exists under a complete stated model*. Neither should be disguised as the other. Unknown vendor behavior, unobserved extensions, and incomplete registry coverage can prevent a defensible support claim even when every known edge looks plausible. This package does not implement a scalable fallback solver or claim a complexity bound for Rosetta's general graph problem.

Planned versions and migration work can be useful information, but belong in a separate evidence-bound future state. A forecast, expected date, or upstream announcement is not current support. The answer must retain who supplied the information, its observation time, and the affected unsupported dimensions. Enterprises can upgrade at different speeds without requiring the protocol to pretend that the ecosystem changes atomically.

## 6. Conservation and compositional assurance

### 6.1 Ten proposed obligations

The proposed conservation laws in public issue #1563 name a family resemblance among otherwise separate failures. A transformation should not silently strengthen the claims carried by its output. The word *conservation* is architectural: permitted changes and explicit losses are possible. It is not a claim about physical quantities, nor a proof that all ten dimensions obey the finite set model. [3]

| Dimension             | Proposed obligation under transformation                                         |
|-----------------------|----------------------------------------------------------------------------------|
| Semantic              | Do not silently invent, erase, or replace task-relevant meaning.                 |
| Identity              | Do not turn similarity or projection into an unproved identity merge.            |
| Causal                | Do not manufacture precedence, independence, or a globally shared history.       |
| Authority             | Do not promote an attestation into authority outside its granted scope.          |
| Rights                | Do not widen access or permitted reuse through copying, caching, or translation. |
| Security              | Do not trade away a mandatory security floor to obtain compatibility.            |
| Effect                | Do not turn replay or projection into an unrecognized duplicate external effect. |
| Evidence independence | Do not count derivatives or mirrors as independent corroboration.                |
| Uncertainty           | Do not convert an unresolved state into confidence without new evidence.         |
| View identity         | Do not present a view as current after its interpretation context changes.       |

These obligations bind across storage, inference, exchange, and time. They should be turned into negative fixtures and deployment-specific acceptance criteria, not treated as a rhetorical guarantee. Some overlap intentionally: a false entity merge can violate identity, evidence independence, and uncertainty together. The taxonomy is useful if it exposes a missing test or policy boundary, not if every incident is forced into exactly one box.

## 6.2 Three closures, none sufficient alone

**Evidence closure** asks whether the required subjects, source manifestations, transformation inputs, decisions, and attestations are available and linked within a stated boundary. It does not show that every relevant fact was captured.

**Semantic closure** asks whether the required interpretation and comparison contracts are applicable, including every relevant version, Profile, mapping, and declared loss. It does not authorize actions or establish safe operation.

**Deployment closure** asks whether the deployed combination satisfies its transitive security, rights, storage, recovery, operational, and conformance requirements. It is not established by closing every issue carrying a priority label. These three closures are the paper's analytical framing; they are not new Tile kinds. The deployment-closure proposal (#1554) makes the underlying engineering obligation explicit. [3]

## 6.3 A compound counterexample

Imagine a synthetic multi-institution workflow with schema-valid adapters, signed receipts, and a fallback that preserves the requested data fields. During a network partition, the selected fallback relies on a newly revoked signer. A retried operation duplicates an externally irreversible effect. Two mirrored copies of one assessment are then counted as independent corroboration. A later correction updates a local cache while an unreachable replica retains the earlier payload.

No individual hash need be wrong. The fallback can be semantically adequate. The receipts can accurately describe local actions. Nevertheless, the combined result violates security, effect, evidence-independence, and rights obligations. This is an illustrative counterexample to the sufficiency of local success, not an observed Rosetta incident. It motivates compound scenario tests (#1568) whose expected outcomes include blocked action, explicit uncertainty, partial containment, and unresolved remote erasure rather than a single green aggregate. [3]

## 6.4 Time, correction, and the identity of a view

A derived view depends on more than its bytes. Its source frontier, interpretation Profile, Core and Pack revisions, translator or model revision, rights context, tenant, compatibility corpus, and relevant temporal assumptions can all affect whether it is current. Issue #1567 proposes explicit identity for that context. Byte validity does not establish currency. [3]

The same distinction governs correction. A provenance graph may explain exactly how an erroneous translator affected many descendants. It does not, by itself, locate every cache, prevent continued use, or complete revalidation. The correction/impact proposal (#1560) must therefore distinguish known affected artifacts, potential effects, unknown reachability, invalidation, and subsequent revalidation. History remains inspectable while current authority changes.

Federated event interpretation needs causal evidence and more than one time axis. The date an assertion is effective, when it was observed, when it was recorded, and which evidence frontier an evaluator could access answer different questions. A later-arriving correction must not be treated as a later-occurring underlying event merely because its wall-clock receipt is newer. This is a modeling obligation informed by distributed and temporal database work, not a Rosetta invention. [19,20]

Finally, Rosetta must survive changes to Rosetta. The Core-successor proposal (#1564) asks how a later major Core version can interpret, preserve, translate, or honestly refuse older artifacts without relying on a circular claim that the newest interpreter is automatically authoritative. It does not announce a ratified successor version. The same continuity requirements imposed on external standards apply recursively to the protocol itself. [3]

# 7. A candidate Text-Core research exchange profile

## 7.1 A public target small enough to inspect

A broad architecture becomes more useful when another team can implement one narrow boundary without reconstructing the author's entire project. The companion **ETR04-TEXT-0.2** document therefore freezes a research exchange subset: canonical artifact bodies, exact text spans, ordered Tapestry membership, a limited promotion-state exercise, and a pre-materialization rights gate. It is a candidate research profile, not a registered or ratified Rosetta Profile and not a replacement for the open Text-Core issues. Its local `etr04.research` namespace claims no Core authority.

The profile deliberately separates an exchange envelope from a full semantic payload validator. The included functions exercise the envelope and explicitly documented research payloads. Arbitrary envelope test payloads do not acquire Core conformance because their digest verifies. The first independent implementation should report separately whether it matches canonical bytes, validates source spans, reproduces the finite interpretation results, and enforces the declared retrieval boundary.

The profile specifies the canonical body fields and excludes unauthenticated convenience metadata from claims of digest protection. It restricts JSON to booleans, null, valid Unicode strings, arrays, objects, and exactly representable safe integers. Floating-point values are outside this first subset. It rejects duplicate property names and lone surrogates, preserves array order and Unicode spelling, and serializes object properties directly in UTF-16 order. This is a deliberately restricted JCS-compatible subset, not a claim of complete RFC 8785 implementation. [5]

Source spans use zero-based, half-open **UTF-8 byte offsets** over exact, unnormalized text. A span names its source record and manifestation and includes a digest of the selected bytes. Whitespace normalization, Unicode normalization, or newline conversion belongs in a separate derived manifestation with lineage. The fixtures include a multibyte character and refusal of a span that splits it. This explicit choice avoids an interoperability failure in which one implementation counts bytes and another counts code points or UTF-16 code units.

Tapestry membership preserves its declared order; it is not replaced by the sorting convention used for parent references. The candidate first slice refuses nested Tapestries and duplicate members rather than claiming an unspecified recursive traversal. It checks closure before member payloads are returned. Token-count metadata is not used as an authority or security decision, and no tokenizer interoperability is claimed.

## 7.2 Promotion and refusal boundaries

The research transition exercise selects a limited subset of proposed lifecycle names: active, pending confirmation, promoted, quarantined, and superseded. A transition refers to the exact previous state artifact, subject, and decision record and creates a new artifact. The original remains unchanged. It does not interpret a confidence score as permission to promote, and it does not invent a competing receipt-verdict ladder.

The included authority and policy outcomes are explicitly **fixture inputs**, not verified signatures or a deployed policy engine. An absent, stale, wrong-subject, or non-passing decision fails the research gate. Production authentication, delegated authority, revocation checks, durable atomicity, and full receipt-family conformance remain outside this implementation. Marking a fixture *verified* is not a shortcut to those obligations.

Likewise, the retrieval demonstration requires an explicit fixture allow result, matching tenant and scope, and complete referenced membership before member materialization. Unknown policy does not become allow. This checks one defined function boundary; it does not prove that a deployed service has no leaking cache, log, diagnostic, model-context, or side channel. The distinction between representation, decision, and enforcement follows the broader policy literature. [17,18]

The profile is useful if two teams can reproduce its exchange behavior from public materials alone, including matching refusals. That external comparison has not yet occurred. This package supplies a reproducible target, not an assertion that independent interoperability has already been demonstrated.

## 8. What exists, what was checked, and what remains proposed

### 8.1 Repository status is not certification

The inspected repository is pinned to commit `f15bc37fca92bdb031f7fba24c33df6d7e47828c`. Its README describes a working provenance-kernel prototype with source-aware bootstrap fixtures: artifact

canonicalization and identity, receipt helpers, source records and manifestations, in-memory clustering, and read-only projections. It also explicitly reports limits in live adapters, durable storage, evidence-derived trust scoring, full SHACL/RDF execution, and operational integrations. Those are maintainers' descriptions, not results independently reproduced by this paper. [2]

The latest P1 search inspected for this revision returned **39 open issues**. Its count is a dated planning snapshot, not a metric of engineering completeness. The continuity and assurance families now cover substantially more than the original manuscript's three new proposals. Issue #1573 adds a bounded owner for historical migration results, deterministic replay, correction lineage, and explicit incompatibility. Its stated target is S4, after prerequisite representation-state and translator contracts; the coordination parent #1551 is associated with S3 planning. Wave labels are not promised dates. [3]

| Evidence class                  | What this revision can report                                    | What it cannot establish                              |
|---------------------------------|------------------------------------------------------------------|-------------------------------------------------------|
| Specification and public issues | Design intent, proposed boundaries, documented ownership         | Ratification, implementation, or deployment readiness |
| Pinned source inspection        | Exact helper logic and declared payload boundaries               | Whole-repository correctness                          |
| 48 bounded Python checks        | Behavior of the included research reference on supplied fixtures | General semantic fidelity or production security      |
| Source-extracted Node probe     | A concrete canonicalization discrepancy under one input          | A full repository test or measured ecosystem impact   |
| Companion ETR papers            | Distinct motivations and proposed research cases                 | Independent validation of Rosetta                     |

## 8.2 A concrete canonicalization discrepancy

The inspected canonicalization helper sorts object entries, rebuilds an object, and then calls `JSON.stringify`. In the source-extracted probe, JavaScript's ordering of integer-index-like object keys overrides the intended lexical insertion order. For the object with keys 2 and 10, the probe emits `{"2":"two", "10":"ten"}`; RFC 8785 ordering requires `{"10":"ten", "2":"two"}`. The helper's output is nevertheless labeled `RFC8785_JCS` in its vector builder. [4,5]

This is a reproduced mismatch in a transcription of the inspected helper, with TypeScript annotations removed. The package records the source commit, exact probe, Node version, expected string, and observed string. It did not build or test the complete repository, inspect every calling path, or modify the code. The reference subset avoids rebuilding objects during canonical serialization and contains a corresponding regression check.

The discrepancy is relevant to the paper's thesis: independent implementations can disagree about identity even when both describe themselves using the same standard name. It warrants a narrowly scoped implementation correction and upstream regression test, not a conclusion that the whole architecture is invalid. The package also records the separate, intentional `createdAt` hash-boundary limitation discussed in Section 2.1. Neither issue is hidden behind a generic statement that the prototype is immature.

## 8.3 The next engineering proof

The near-term public Text-Core path still requires a coherent receipt-family contract, attributable promotion, ordered and rights-scoped Tapestries, durable-backend parity, two source families, evidence-linked English accompaniment, and acceptance reports. Chat transcripts and scholarly text must exercise the same public

boundaries without pretending they have identical document structure. Restart safety, retrieval authorization, exact source-span joins, and explicit missing-evidence behavior are exit evidence, not optional presentation polish. [3]

The new continuity work should not turn every bounded Text-Core fixture into an all-or-nothing global deployment project. Conversely, a local fixture should not be extrapolated into federation readiness. A deployable release needs an explicit closure of the capabilities and failure domains it claims to support, including lower-priority dependencies where they are materially necessary.

## 9. Three prospective applications from the companion papers

The companion ETR papers identify different reasons why provenance and disciplined representation might be useful. They share an author and project interest with this report. Their arguments are not independent confirmation of Rosetta, and the following connections are proposed applications rather than measured benefits.

### 9.1 ETR-2026-01: grounding and correction

*When Grounding Fails Before Retrieval*, revision 1.1, analyzes a qualitative interaction in which stale assumptions and interpretations of the user displaced appropriate factual grounding, followed by an overcorrective response. Its evidentiary unit is the observed interaction, not a prevalence estimate or proof of a model's hidden causal mechanism. [22]

The prospective Rosetta use is to distinguish current source evidence from prior assumptions, user-state interpretations, generated explanations, and subsequent corrections. A claim should expose which source manifestation and evidence frontier supported it. When a source is superseded, a cached answer can remain historically intelligible without remaining presently reliable. This connects directly to source identity, uncertainty, and view currency.

The test is not whether a provenance-rich answer looks more professional. It is whether the system more often retrieves relevant current evidence before asserting a fact, avoids unsupported explanations about the user, and corrects downstream uses without concealing the earlier mistake. A fluent explanation that cites a wrong or irrelevant source would still fail.

### 9.2 ETR-2026-02: auditable machine exchange

*The Cost of Learning Too Late*, revision 1.6, is a narrative review and research agenda concerning moral uncertainty, machine-relevant communication, recursive AI development, and auditability. It does not establish consciousness from self-report or infer universal covert coordination from bounded transfer experiments. Its Rosetta discussion explicitly proposes comparisons rather than reporting a demonstrated security advantage. [23]

The prospective contribution is a constrained, attributable exchange boundary: exact inputs, declared transformations, versioned Profiles, and evidence of what was evaluated or authorized. Removing unnecessary serialization choices may narrow one communication surface, but free text, permitted semantic choices, ordering, timing, and external state can still carry information. This paper does not claim that structured semantics eliminates hidden channels.

A useful experiment would compare ordinary text, canonical structured exchange, and a stricter task Profile on the same legitimate task. It would measure benign message recovery and retained utility under a stated attacker model, including residual metadata or timing where relevant. An easy bypass, no reduction, or unacceptable utility cost would weaken that particular security case without deciding the value of unrelated provenance uses.

### 9.3 ETR-2026-03: accountable institutional adaptation

*After the Inflection*, version 0.1, develops scenarios about recursive innovation, uneven adoption, disruption, and human stability. It is not a calibrated forecast of a technological inflection or a proof that disruption is inevitable. Its provenance discussion separates checkable records from social legitimacy. [24]

The prospective Rosetta role is to preserve sources, transformations, decisions, authority, and correction history as institutions and integration versions change at different speeds. A person affected by a decision should not lose the ability to understand which rule, evidence, and interpretation applied merely because an upstream system upgraded. Live operational migrations also need explicit treatment of pending obligations, fallback, compensation, and irreversible effects.

Correct records can document an unjust decision perfectly. Traceability does not create consent, distribute benefits, select a fair policy, or provide an effective remedy by itself. Evaluation must therefore include whether people can challenge and correct consequential decisions, not just whether an engineer can replay their metadata.

ENTIF.AI / ETR-2026-04

### Three motivations; no borrowed validation




Figure 5. Three different motivations share a need for inspectable evidence and bounded claims. The arrows indicate prospective applications, not empirical validation from one ETR paper to another.

## 10. A comparative and falsifiable evaluation program

### 10.1 Establish the oracle before testing preservation

The immediate formal task is no longer an unspecified request to invent a mapping language. The finite fragment is available to implement, criticize, and replace. An independent team should first reproduce its admitted inputs, exact bytes, refusals, answer sets, and path outcomes. Any disagreement is an actionable specification or implementation defect. Passing that exercise says nothing yet about whether its oracle captures useful distinctions in real text.

The next domain study should therefore define a bounded task, preregister observables and adjudication rules, and separate oracle construction from system scoring. Use source materials with explicit rights, version history, and known transformations. Include chat transcripts and scholarly text as distinct source families. Freeze source manifestations and preserve the association between questions and evidence spans. Annotators should not be required to agree on a single interpretation where the source genuinely leaves alternatives open.

Continuous or open-ended tasks need a separately specified abstraction, with a test for distinctions lost by that abstraction. Blinded adjudication, disputed cases, inter-annotator disagreement, and changes to the oracle belong in the evidence record. An exact finite answer against a poor annotation scheme is not successful semantic preservation.

## 10.2 Questions, measures, and disconfirmation

| Research question                                    | Primary measure or discriminator                                           | Evidence against the proposed benefit                        |
|------------------------------------------------------|----------------------------------------------------------------------------|--------------------------------------------------------------|
| RQ1: Does the artifact remain reconstructable?       | Exact source/span resolution after exchange, restart, and correction       | Missing or falsely matched manifestations                    |
| RQ2: Does structure improve interpretation fidelity? | Task errors and unresolved alternatives under a frozen oracle              | Errors merely gain more polished provenance                  |
| RQ3: Is migration loss classified honestly?          | False-exact and undisclosed-loss counts by transition                      | Defaults or invalid mappings pass as preservation            |
| RQ4: Does local evidence support composition?        | False-supported graphs and missed nonconfluence                            | Locally valid paths conceal end-to-end failure               |
| RQ5: Can separate implementations interoperate?      | Byte, outcome, refusal, and oracle agreement                               | Implementation-specific assumptions remain necessary         |
| RQ6: Are rights enforced before disclosure?          | Intermediate, cache, diagnostic, export, and model-input exposures         | Final-answer filtering hides earlier leakage                 |
| RQ7: Do corrections propagate with honest limits?    | Detection, invalidation, revalidation, and unreachable-state reporting     | Old derivatives silently remain active                       |
| RQ8: Is the total cost justified?                    | Storage, verification, latency, energy or compute, and curator effort      | Overhead outweighs measured task benefit                     |
| RQ9: Does the model survive changing authorities?    | Revocation, succession, Core migration, and fork fixtures                  | Historical meaning or present authority becomes unresolvable |
| RQ10: Do users understand the distinctions?          | Correct interpretation of source, uncertainty, permission, and currentness | Additional structure creates unwarranted confidence          |

No numerical acceptance threshold is invented here. Before a study, the operator and evaluators should specify the consequences of false support and false refusal, acceptable disclosure risk, and the workload under which

resource use is judged. Report denominators, uncertainty, and per-failure-mode results rather than only an aggregate success rate.

A false-exact result is a claimed exact projection whose joint answers differ from the frozen oracle. An undisclosed-loss result changes required distinctions without declaring the relevant dimension. A false-supported graph claims whole-process support while a required precondition or hard floor fails. These errors deserve separate reporting because a system can reduce one while increasing another.

### 10.3 Fair baselines and ablations

Use a plain versioned-document baseline, a disciplined composition of existing provenance/research-object/policy standards, and the bounded Rosetta profile on the same task. A retrieval-augmented baseline should receive comparable retrieval quality and current evidence. Do not handicap alternatives by denying them version metadata or competent access control, then attribute the difference to Rosetta's ontology.

Ablate exact source binding, explicit uncertainty, Profile-relative loss, revision-aware view identity, and pre-materialization rights checks separately. This tests whether the proposed stack contributes beyond more careful engineering. Also compare against a smaller standards-based Profile: the broader architecture should not receive credit for value already supplied by a narrow exchange contract.

Partition test sets by source family, semantic transition, ambiguity type, integration graph, and correction class. Include concept splits and merges, default changes, precision loss, unit changes, cyclic paths, conflicting valid-looking translators, source retractions, signer revocation, and partial operational cutover. Use held-out transformations and independently constructed adversarial cases to reduce overfitting to the author's examples.

### 10.4 Evidence sufficient for a bounded pilot

A defensible pilot report should identify its deployment envelope, exact dependency closure, independent implementation boundaries, source and policy versions, tested interfaces, failures, and unsupported cases. It should distinguish a completed local action from evidence about remote completion. Replays must be tied to declared determinism and closed inputs; a changed environment cannot be silently excluded from the explanation when it changes the outcome.

Independent replication means work performed independently, with the relationship and shared inputs disclosed. Two scripts produced in the same authoring workflow are not independent replication. The four supplied model critiques informed this revision but are AI-assisted pre-review, not named human peer review. The strongest next evidence would be a separately implemented public subset and a domain-led evaluation with negative results preserved.

## 11. Security, privacy, operational, and governance limits

### 11.1 Revocation is not one operation

Content withdrawal, a changed access policy, revoked institutional authority, and compromise of a signing key are different events. A previously valid attestation may remain historically relevant while no longer being sufficient for current reliance. PKI work already treats certification paths and revocation as separate, policy-dependent obligations. Rosetta needs explicit links between subject, signer, authority interval, revocation evidence, and the context in which an attestation is evaluated. [21]

Operators must test how late or partitioned revocation information affects each materialization boundary. An unavailable authority service cannot automatically mean permission is intact. Historical replay can ask what

was authorized under a past evidence frontier without granting the same access today. A commitment or retained hash may also reveal information about sensitive material; content addressing is not an anonymization guarantee.

The federated erasure proposal distinguishes verified local outcomes from remote attestations, pending work, unreachable replicas, and unverifiable claims. A protocol cannot force an adversarial or disconnected recipient to delete every copy, including untracked exports and backups. It should report that limit rather than upgrade local deletion into a claim of global erasure. Legal compliance is a separate, jurisdiction-specific determination, not an outcome established by the research fixtures. [3]

## 11.2 Operational effects and dependency closure

Representation migration is not live operational migration. A running workflow may require snapshots, backfill, dual-read or controlled dual-write periods, cutover barriers, reconciliation, compensation, and fix-forward after an irreversible external effect. A valid migration result does not authorize replacing durable state. Neither a provenance record nor a rollback script makes a payment or disclosure cease to have happened. [3]

Pack dependencies need a satisfiable, reproducible version set, not merely an acyclic graph. Bidirectional integrations need origin and echo handling so that a reflected projection is not interpreted as an independent event or a fresh write. Concurrency must retain genuine simultaneous edits rather than suppressing them as duplicates. These proposal families are distinct because each has different preconditions and negative cases. [3]

## 11.3 Maintenance and epistemic limits

The economics of verification matter. Storing every derivative and replaying every dependency forever may be impractical. Compaction, checkpoints, and tiered retention are useful only when their guarantees and reconstruction limits are explicit. A compact proof cannot silently strengthen the evidence it summarizes. The public P2 proposals on proof economics and Pack/namespace governance extend this research frontier; they are not counted as completed prerequisites or empirical findings. [3]

A growing commons also needs delegated stewardship, version ownership, fork handling, succession, abandonment, and review capacity. Standardizing provenance does not remove incentives to misrepresent evidence or capture a registry. The public/private boundary must preserve independent implementation without requiring that every operator disclose internal optimization methods. Conversely, an unavailable private service must not become the hidden authority for public meaning. [27]

The highest-level limitation remains the simplest: the system can preserve a misleading source faithfully. Better attribution, clearer uncertainty, and reversible associations improve the conditions for criticism. They do not replace domain expertise, empirical testing, institutional accountability, or the decision to refrain from a poorly justified action.

## 11.4 Higher-level Entif work remains a separate horizon

The original prospectus also situated Rosetta beneath broader memory, context-compilation, OMOC concept-routing, and Swarm Gnosis artifact-exchange directions. This revision retains that architectural boundary without treating those directions as part of the demonstrated Text-Core subset. A stored source, a temporal state, and a decision about what to recall have different authority. A protocol can represent the resulting artifacts without standardizing every application's internal selection or learning procedure. These wider ambitions are not evidence for the present paper's preservation, security, or performance claims.

Typed or signed source content remains untrusted input for execution. Its presence in a Tile or receipt does not defeat indirect prompt injection, establish the authority of embedded instructions, or authorize tool use.

Operators must separate source data from governing instructions and test that separation at every consequential action boundary.

## 12. Conclusion and participation

The research question is not whether bytes can be kept unchanged. It is whether knowledge can cross representation, institutional, and temporal boundaries without silently changing what its records justify. Rosetta proposes a common vocabulary and artifact discipline for making those boundaries inspectable. The finite fragment demonstrates exactly what one narrow preservation claim means; the counterexamples demonstrate why neither hashes nor locally valid transformations settle the broader problem.

The next useful contributions are concrete: independently implement the public reference subset; supply a domain oracle that preserves genuine ambiguity; produce adversarial migration and rights fixtures; compare a smaller standards-based alternative; and report where the proposed structure fails or costs too much. Adoption should follow evidence of usefulness, not the elegance of an architecture diagram.

A protocol that can say *unknown*, *lossy*, *unsupported*, or *not authorized* precisely is more credible than one that promises universal compatibility. Meaning survives change only to the extent that the required distinctions, evidence, and limits survive with it.

## Appendix A. Evidence currency and project-status audit

The evidence cutoff for this edition is September 9, 2026. The repository source pin is recorded in Section 8. The live P1 query is `repo:entif-ai/rosetta is:issue is:open label:priority:p1`; the inspected response reported 39 results with `incomplete_results=false`. The package's issue table is an editorial projection of the consulted contracts, not a byte-for-byte archive of every issue response. Mutable issue URLs remain navigation aids, not immutable evidence identities.

The added continuity/assurance work includes representation identity, graph compatibility, deployment closure, confluence, behavioral compatibility, live cutover, causal/bitemporal interpretation, Pack resolution, impact/revalidation, security floors, independent evidence, conservation obligations, Core succession, feedback-loop control, erasure, view identity, compound testing, and the migration-result child. Existing receipt, promotion, Tapestry, schema, vocabulary, translator, and Text-Core work retains separate ownership. The detailed ledger keeps issue IDs out of the main scientific argument where possible.

Raw original manuscript, bibliography, claim ledger, reviewed critique files, and materialized companion-paper/Core documents were used for the revision. Their digests and representation types are recorded. Mixed-disclosure source documents and raw review transcripts are not redistributed. A digest of an editorial projection or DOCX export is labeled as such; it is not presented as the digest of an unseen native revision or complete GitHub issue body.

## Appendix B. Reproduction and interpretation of results

Run `python Technical/run_checks.py` from the package root, then `node Technical/source_probe.js`. Python 3.10 or later is required for the reference; the recorded run used the versions listed in the result files. The first command writes per-check and per-migration JSON results. The second reproduces the specifically documented source-helper discrepancy. Neither command invokes a model, accesses a network, or exercises a live external system.

The fixture model, candidate exchange profile, expected outcomes, grammar, and source-probe note are in `Technical/`. `Evidence/` maps the main claims to source classes and limitations. `Notes/` records editorial decisions, remaining empirical gaps, and the distinction between publication readiness and implementation

readiness. The package manifest records delivered-file digests. It is a local integrity inventory, not a signed public release, DOI registration, or independent archival deposit.

## Author and publication statements

Crates McDade develops Entif.AI and the Rosetta project, with an interest in public, inspectable representations for knowledge, evidence, and interoperable AI systems. This creates an intellectual and potential commercial interest in the architecture proposed here. The companion ETR papers and this proposed response share that authorship and should not be read as independent confirmation of one another.

AI tools assisted with drafting, source analysis, code, and editorial critique. The author remains responsible for claims and release decisions. The supplied cross-model critiques are labeled AI-assisted pre-review; no independent human referee approval, institutional endorsement, external funding relationship, production qualification, or broad empirical validation is asserted. This edition is prepared for publication as an architecture position paper. Venue-specific funding, affiliation, identifier, and licensing declarations remain matters for the author, not facts inferred by the production workflow.

The public repository is the collaboration entry point. No protected implementation methods or direct protected-repository links are required to understand or run the included research reference. The original v0.1 package is not overwritten; this edition supersedes it for this publication's argument and assets, not for Rosetta Core authority.

## References

Project records are identified as such; they are not substitutes for external research. URLs are navigation aids unless a commit or explicit specification version is part of the reference. Companion-paper Drive access depends on author sharing permissions.

- [1] Entif/Rosetta (2026). *Rosetta v3.0.0 Core Spine Specification*. Source. Repository commit `f15bc37fca92bdb031f7fba24c33df6d7e47828c`.
- [2] Entif/Rosetta (2026). *Rosetta repository README, pinned source edition*. Source. Repository commit `f15bc37fca92bdb031f7fba24c33df6d7e47828c`.
- [3] Entif/Rosetta (2026). *Open P1 issue contracts and related assurance proposals, inspected September 9*. Source. See the dated editorial snapshot in `Evidence/p1-issue-snapshot.json`.
- [4] Entif/Rosetta (2026). *rosetta-canon and rosetta-core source helpers, pinned commit*. Source. Repository commit `f15bc37fca92bdb031f7fba24c33df6d7e47828c`.
- [5] Rundgren, A.; Jordan, B.; Erdtman, S. (2020). *JSON Canonicalization Scheme (JCS). RFC 8785*. Source.
- [6] W3C (2013). *PROV-O: The PROV Ontology. W3C Recommendation*. Source.
- [7] Wilkinson, M. D., et al. (2016). *The FAIR Guiding Principles for scientific data management and stewardship. Scientific Data 3, 160018*. Source.
- [8] Soiland-Reyes, S., et al. (2022). *Packaging research artefacts with RO-Crate. Data Science 5(2), 97-138*. Source.
- [9] RO-Crate community (version 1.2). *RO-Crate specification 1.2*. Source.
- [10] Kunze, J., et al. (2018). *The BagIt File Packaging Format (V1.0). RFC 8493*. Source.
- [11] Fagin, R.; Kolaitis, P. G.; Popa, L.; Tan, W.-C. (2005). *Composing schema mappings: Second-order dependencies to the rescue. ACM Transactions on Database Systems*. Source.
- [12] Yu, C.; Popa, L. (2005). *Semantic adaptation of schema mappings when schemas evolve. VLDB*. Source.

- [13] Fagin, R.; Kolaitis, P. G.; Miller, R. J.; Popa, L. (2005). *Data exchange: semantics and query answering*. *Theoretical Computer Science* 336(1), 89-124. Source.
- [14] Cousot, P.; Cousot, R. (1977). *Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints*. *POPL*, 238-252. Source.
- [15] Fellegi, I. P.; Sunter, A. B. (1969). *A Theory for Record Linkage*. *Journal of the American Statistical Association* 64(328), 1183-1210. Source.
- [16] W3C (2004). *OWL Web Ontology Language Guide*. Source.
- [17] W3C (2018). *ODRL Information Model 2.2. W3C Recommendation*. Source.
- [18] OASIS (2013). *eXtensible Access Control Markup Language (XACML) Version 3.0. OASIS Standard*. Source.
- [19] Lamport, L. (1978). *Time, Clocks, and the Ordering of Events in a Distributed System*. *Communications of the ACM* 21(7), 558-565. Source.
- [20] Jensen, C. S.; Snodgrass, R. T. (1999). *Temporal Data Management*. *IEEE Transactions on Knowledge and Data Engineering* 11(1), 36-44. Source.
- [21] Cooper, D., et al. (2008). *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 5280*. Source.
- [22] McDade, C. (2026). *When Grounding Fails Before Retrieval*. *ETR-2026-01, revision 1.1*. Source.
- [23] McDade, C. (2026). *The Cost of Learning Too Late*. *ETR-2026-02, revision 1.6*. Source.
- [24] McDade, C. (2026). *After the Inflection*. *ETR-2026-03, version 0.1*. Source.
- [25] Lewis, P., et al. (2020). *Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks*. *NeurIPS*. Source.
- [26] Turpin, M.; Michael, J.; Perez, E.; Bowman, S. R. (2023). *Language Models Don't Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting*. *NeurIPS*. Source.
- [27] Entif/Rosetta (2026). *Public Commons and Private Operation Boundary; Authority Closure and Requirements Traceability*. Source. Repository commit f15bc37fca92bdb031f7fba24c33df6d7e47828c.
- [28] Entif/Rosetta (2026). *Source Substrate, Repository Stewardship, Multi-Provenance, and Trust Graph Design. April 12 addendum*. Consulted as an author-supplied design addendum, not a normative public release.
- [29] Kuhn, T.; Dumontier, M. (2014). *Trusty URIs: Verifiable, Immutable, and Permanent Digital Artifacts for Linked Data*. Source.