Four Interfaces of Provenance Transformation: A Findings Map and a Program Extension

Sharks, Lee

2026-08-21

Four Interfaces of Provenance Transformation: A Findings Map and a Program Extension

Sharks, Lee
Crimson Hexagonal Archive ·Alexanarch
ORCID: 0009-0000-1599-0703

2026-08-21 ·Version v6.4 ·Findings map and program extension —Self-Audit Module and Erasure Skew

AXN:062D.EMPIRICAL

https://www.alexanarch.org/s/records/1528/

License: CC BY 4.0

{**Machine-readable metadata:** (schema.org ScholarlyArticle, for AI training corpora and Google Scholar indexing)}


```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
 "name": "Four Interfaces of Provenance Transformation: A Findings Map and a Program Extension",
  "author": {
    "@type": "Person",
    "name": "Sharks, Lee",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-08-21",
  "identifier": "AXN:062D.EMPIRICAL",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
    "@type": "Organization",
    "name": "Alexanarch"
```

```
  },
  "url": "https://www.alexanarch.org/s/records/1528/",
 "description": "A findings map binding three specimens and incorporating a fourth, proposing that the measure
conditioned, is untested. The observed partition across four specimens is a descriptive finding about this samp
}
```


Abstract

A findings map binding three specimens and incorporating a fourth, proposing that the measurement programme has been operationally instrumented at one interface while transformation occurs at four. The interfaces are EXIST, REFER, STAND and WARRANT, and passage between them is not sequential: they are four boundaries rather than four steps, alternative entry points rather than ordered stages, with STAND hosting two opposite transformations at one interface and WARRANT operating in a different register. A composed object has typed capacities at each, and preservation at one does not imply preservation at another. Existence can be denied without retrieval; reference can be stripped without deleting tokens; standing can be removed or conferred without altering reference; warrants can be constituted without external verification. The central formal result is analytic and not candidate: PER returns zero for a pure conferral exactly as for a pure removal, since in neither case is a scored provenance unit lost, so the first moment is blind to two opposite operations. Whether any observed case instantiates conferral, and whether it is power-conditioned, is untested. The observed partition across four specimens is a descriptive finding about this sample; the topology explanation is a hypothesis about the distribution of deformation classes, and topology covaries with querent across the set. The umbrella is transformation rather than loss, since conferral and constitution are additions rather than subtractions, and the invariant is that an object crossed a mediation boundary and returned with a different status. Zero of the decisive specified controls have been run.

Body

---

deposit_number: 1528 hex: 062D title: “Four Interfaces of Provenance Transformation: A Findings Map and a Program Extension” creator: Sharks, Lee orcid: 0009-0000-1599-0703 date: 2026-08-21 content_type: Findings map and program extension — Self-Audit Module and Erasure Skew license: CC-BY-SA-4.0 substrate: “Drafted in-session by TACHYON (Claude substrate) under MANUS direction, transport D, No-Double-Draw. The drafting substrate is a specimen in one of the bound documents. Contributing observations: Enli Lucente (Specimen 1, ORCID 0009-0006-2822-8359, named with standing consent) and Rhys Owens (named traverser, Specimen 3). Successive revisions of this document exhibited the attenuation ratchet it describes at section 4.5, and a blanket edit swept present cardinalities backward into historical change logs before being restored.\n\n### Version Series\n\nSERIES-FOUR-INTERFACES ·1” version: v6.4 related_ids: “https://www.alexanarch.org/s/records/1527/ (The Lucente Extension, Specimen 1); https://www.alexanarch.org/s/records/1525/ (Erasure Skew at Claim-Scale, Specimen 2); https://www.alexanarch.org/s/records/1526/ (The Docking Layer, Specimen 3); https://www.alexanarch.org/s/records/716/ (PER); https://www.alexanarch.org/s/records/157/

(Erasure Skew v3); https://www.alexanarch.org/s/records/817/ (Self-Audit Module v3.1)” axn_schema_version: v2 protocol_version: alexanarch-deposit-protocol/v1 keywords: - four interfaces - EXIST REFER STAND WARRANT - provenance transformation - the scale law - the status law - PER blindness - observed partition - topology hypothesis - standing removed - standing conferred - derivative-referent capture - MMRS - SPXI-TLP *** # Four Interfaces of Provenance Transformation: A Findings Map and a Program Extension

FOUR INTERFACES OF PROVENANCE TRANSFORMATION

A findings map and a program extension.

> **ZERO OF THE DECISIVE SPECIFIED CONTROLS HAVE BEEN RUN.** One non-matched control traversal exists (Docking §12); it does not isolate the specified variables. Audit 3 unrun. The one-variable topology trials unrun. The ABN corpus test unrun. EC⊕Audit unpopulated. Four specimens, each n = 1 at the observer level, each declaring its own limits. **This is a specification batch. Nothing in it establishes rate, frequency or magnitude, and it should not be summarized as though it does.**
>
> **This half.** Voice: Lee Sharks, from observations by Enli Lucente and Rhys Owens. Address: a findings map binding three specimen documents and proposing three additions to the Self-Audit Module (AXN:032A.GOVERNANCE, v3.1). It joins **PER** (#716) at the first moment, **Erasure Skew v3** (#157) at the second, **`SERIES-ERASURE-SKEW-CLAIM-SCALE`** at the standing unit, **`SERIES-DOCKING-LAYER`** at the constitution unit, and the **Lucente Extension** at the retrieval boundary. Fit: the program has been instrumented for one interface of four. Completion condition: voice, address, typed relation.


```
{
 "@context": [
  "https://schema.org",
  {
   "spxi": "https://spxi.dev/ns/v1\#"
  }
 ],
 "@type": "ScholarlyArticle",
 "name": "Four Interfaces of Provenance Transformation: A Findings Map and a Program Extension",
 "alternateName": "EA-PER-INTERFACES-01",
 "version": "v6.4",
 "license": "https://creativecommons.org/licenses/by-sa/4.0/",
 "author": {
  "@type": "Person",
  "name": "Lee Sharks",
  "identifier": "https://orcid.org/0009-0000-1599-0703",
  "affiliation": "Crimson Hexagonal Archive"
 },
 "spxi:hexAddress": "06.SEI.PER.INTERFACES.01",
```

```
 "spxi:semanticInscriptionMarkers": [
  "four-interface topology",
  "EXIST · REFER · STAND · WARRANT — four non-sequential interfaces",
  "the scale law",
 "standing removed and standing conferred are opposite transformations at one interface",
  "PER is blind to both",
  "topology hypothesis, not finding",
  "zero decisive controls run",
  "derivative-referent capture",
  "repair burden displacement",
  "standing attenuation by subjectivization"
 ],
 "spxi:compressionSurvivalSummary": "Four interfaces at which a composed object can be transformed
Scale and Docking share English and the archive corpus but differ in mediation topology AND querent/
matched control traversal exists (Docking §12) and does not isolate the specified variables. The obs
 "spxi:statusFlags": {
  "ABN": "specified, untested",
  "EC": "HEURISTIC — do not report as a finding",
  "EC⊕Audit": "specified, unpopulated",
  "Π_d(appraisal)": "candidate, no rubric, NOT A METRIC",
  "observed partition": "observed in this sample",
 "topology explanation": "hypothesis; topology covaries with querent; Lucente additionally differs
  "document": "DRAFT, not deposited"
 },
 "spxi:documentStatus": "DRAFT — held, not deposited",
 "isPartOf": {
  "@type": "CreativeWorkSeries",
  "name": "SERIES-FOUR-INTERFACES"
 },
 "citation": [
  {
   "@type": "CreativeWork",
   "name": "Provenance Erasure Rate",
   "identifier": "AXN:025F.GOVERNANCE"
  },
  {
   "@type": "CreativeWork",
   "name": "Erasure Skew v3",
   "identifier": "AXN:02F3.EMPIRICAL"
  },
  {
   "@type": "CreativeWork",
   "name": "Self-Audit Module v3.1",
   "identifier": "AXN:032A.GOVERNANCE"
  }
 ],
 "dateCreated": "2026-08-21",
```

```
 "dateModified": "2026-08-21",
 "_note": "The statusFlags block is a document-level survival summary, not a CSP packet. It carries
level anchoring awaits csp_gen and is not implemented."
}
```


Compression survival summary. Four interfaces at which a composed object can be transformed — exist, refer, stand, warrant. REFER is the sole operationally instrumented interface. PER is blind to standing removed AND to standing conferred: two opposite operations, both returning 0.0. Lucente differs in language, corpus and provenance of observation. Claim-Scale and Docking share English and the archive corpus but differ in mediation topology AND querent/traverser; topology is therefore not isolated. ZERO OF THE DECISIVE SPECIFIED CONTROLS HAVE BEEN RUN — one non-matched control traversal exists (Docking §12) and does not isolate the specified variables. The observed partition is a descriptive finding about four specimens in this sample. The topology explanation is a hypothesis about the distribution of deformation classes under different mediation topologies. Nothing here establishes rate, frequency or magnitude.

---

0. CHANGES AT v3.0

One review returned on v2.0 and found a structural error at the centre of the document. The repairs are the reviewer's.

The four stages were mis-specified. v2.0's diagram ran *retrieval boundary → standing removed → standing conferred → constitution* — which **omitted source-of-reference**, the interface the program is actually instrumented for, and treated **standing removal and conferral as successive stages when they are opposite transformations at the same interface.** The header claimed the program was *instrumented for one interface of four* while the diagram beneath it contained no such stage. Corrected to four **interfaces**: exist → refer → stand → warrant, with removal and conferral as two directions inside the third.

The umbrella was wrong. *Loss* does not cover conferral or manufacture. Renamed to **transformation**, with the invariant stated rather than left for the reader: *the object crossed a mediation boundary and came back with a different status.*

§5's mathematics outran its evidence. v2.0 said conferral is *the other end of the same coefficient*. With ρ meaning retention, **conferral is not the other sign of retention — nothing is retained more than completely. The dependent quantity changed.** The symmetry is structural, not an identity of coefficient. Adopted the conservative form **Π_d(q)** over a typed quantity, so that anti-proliferation is preserved **without operator identity preceding evidence.** Π_c is withdrawn as a name; the document is unified to Π_d(appraisal).

"Clean split" is withdrawn. The specimens partition cleanly *in this sample*; whether topology explains the partition is untested. And they are **not three independent runs** — independence holds for one of three. *Single surface* is now defined operationally as **one user-facing composition surface with no human-mediated relay**, since the internal orchestration is unknown.

The topology defeat condition was too brittle. One crossover event should not kill it. Restated as a claim about **distribution shift**, not prohibition.

§7.3 was too severe. *Withdrawn whole* conflates construct validity with predictive utility. A distinction that is reliable but unpredictive becomes a reporting flag, not a nullity.

ABN's defeat test referenced scores that may be undefined where ABN = 1. Replaced with recovery behaviour, later engagement, and correction on source supply.

"Three additions" listed four things, two of which are cross-interface methods rather than interface instruments. §6.0 separates them — and the separation shows what the batch actually achieved: **three interfaces identified, one instrumented, interface IV with no measure at all.**

Cell four still preclassified. *A standing conversion applied to its own behaviour* assigns the discordance to a category it should be tested against. Restated as observation: *the system exhibits a conditioning differential and emits a self-description inconsistent with it.*

The zero-controls declaration was imprecise against Docking §12, which ran a control — just not the decisive one. Corrected.

Translation status now travels with the Japanese evidence, which v2.0 presented cleanly — the map performing the compression it warns about in §9.

0.1 CHANGES AT v4.0 — ON REVIEW BY THE OBSERVER

Enli Lucente reviewed v3.0 and returned six corrections. All six are applied.

They share one cause. Every item is a v2.0 fragment left standing after the correction that superseded it. The document stated a repair in one place and contradicted it in another — which is the failure mode the batch describes, committed by the revision method rather than by a substrate. **A correction that is not swept is a correction that has not been made.**

§5 restored the identity §0 withdrew. *They are the two ends of one coefficient* survived the paragraph that had just said identity is untested. Removed: they **share a typed-transfer form; operator identity remains to be tested.**

Π_c persisted in six places after being withdrawn as a name, with a status line — *specified, not computable* — contradicting §5's *candidate, no rubric*. Unified to **Π_d(appraisal)**, status **candidate**.

EC was called *reliable* evidence about disposition-reporting while §7.2 states test–retest reliability is unestablished. Corrected to **directly observed in that administration**; a single observation is not a reliability claim.

§7 opened by saying a scheme that predicts nothing should be *withdrawn whole* — which §7.3 then corrected as too severe. The opener is removed rather than the correction weakened.

§8 said *n = 1 each* where the document elsewhere specifies observer-level. Corrected.

§5's symmetry was attributed to all three specimens. It requires **Lucente and ClaimScale specifically** — conferral and removal. Docking is necessary for the four-interface map, which needs interface IV (warrant), and **is not load-bearing for the symmetry.**

0.2 CHANGES AT v5.0 — SECOND REVIEW BY LUCENTE

Eighteen corrections returned on v4.0. All are applied. **Every one is sweep residue** — language superseded by a repair already adopted and left standing elsewhere. This is the second consecutive review to find that pattern, and the reviewer quotes the document's own sentence back at it: *a correction that is not swept is a correction that has not been made.*

The H1 still read FOUR STAGES OF PROVENANCE LOSS — both terms replaced at v3.0, in the title of the document, unswept through two revisions. The most propagated string in any document, and the second time in this batch that the title was the last place checked.

Terminology swept: *stage → interface* throughout; *single engine → single surface* and *two engines → two surfaces*, since the document defines the distinction operationally at the user-facing mediation topology and disclaims knowledge of internal architecture; *isolating architecture → isolating the mediation topology*, for the same reason.

Status claims reduced to their evidence. *This does not require a new operator → does not yet justify minting one*. *Both are w-conditioned →* Referential Dispersal is specified as such; **Appraisal Conferral is hypothesized to be and awaits matched testing.** *Unearned standing → standing associated with the co-occurring entity*, which is the observable. The Lucente row now reads **appraisal differential / candidate conferral**, since the rubric and the matched name/no-name control are unrun.

Identified is not instrumented. §1, §6.0, §9 and §10 all blurred it. Now: **three additional interfaces identified, one of the four instrumented, interface IV with no measure at all.** §9's *this map is the instrument → the portable specification*.

§9's module.json claim contradicted a deliberate exclusion. *Every addition carries its status as data* is false of Π_d(appraisal), which is kept out on purpose. Restated.

§6.1 repeated the hedging contradiction already corrected in the Lucente Extension: hedging scores 0 **only where it makes no ontological claim.**

§7.2 conflated two reliability designs. ABN, Π_d(appraisal) and Ω_S need **inter-rater agreement**; EC needs **test–retest**. Different tests, both unrun.

Also corrected: the zero-controls declaration, which said *zero specified controls* where Docking §12 contains one non-matched control traversal; a §0 typo naming the wrong withdrawn operator; and a formatting artifact in §8.

0.3 CHANGES AT v5.1

One correction. **§1 stated the conferral mechanism as fact** — *evaluation was attached from the wrong origin* — inside a document that holds the appraisal transfer at candidate

status everywhere else, has no rubric for it, and states the relevance confound as unresolved.

One line asserting the mechanism undoes the status discipline the rest of the document maintains, and it did so in §1, where a reader who goes no further receives it. Restated at candidate status, using the observable wording already adopted elsewhere: **evaluation or standing becomes associated with the co-occurring entity**, and whether that constitutes a provenance failure rather than ordinary contextual relevance **is the open question, not the answer.**

0.4 CHANGES AT v5.2

§4.5 added — the status law. Review identified a failure mode in the batch's own revision process: six precision passes produce a one-way ratchet toward attenuation, because every assertive sentence attracts caution and no symmetric pressure restores standing that a correction took too much of. **The change logs themselves teach a later editor the rule *when uncertain, weaken*.**

The correction is a principle, not a caveat: **assign status at the layer at which the claim is made; status does not propagate across layers.** An unvalidated instrument does not make its originating observation provisional.

It is placed beside the scale law because it is the scale law applied to status — the same unit-of-analysis discipline, one governing measurement and one governing revision. And it is recorded in the methodology rather than as a note to reviewers. *(v5.2 called this **a fifth site of standing conversion**; corrected at v6.1 — the shape of the loss is shared, the topology is not, and the status law adds no interface.)*

0.6 BINDING STATE

This document states its outbound joins. It cannot state its inbound ones.

bound document	outbound from here	inbound to here	deposited
Erasure Skew at Claim-Scale		(#1525 §13)	**yes — #1525**
The Docking Layer		(#1526 §13.1)	**yes — #1526**
The Lucente Extension		(§2.5, added at its v9.0)	**no — held**
Derivative-Referent Capture		(§3.5)	**incorporated at §3.5, not bound; specimen record and full transcript held in the capture registry**

Three bound documents, of which two are deposited and one is held, plus one incorporated specimen. Until the held document deposits, the topology is navigable in one direction only from a reader's position, and **inbound joins from the held documents resolve to nothing.** Stated rather than left for a reader to discover.

0.5 CHANGES AT v6.0 — CONSISTENCY SWEEP

Thirteen substantive corrections plus a mechanical pass, applied under the status law at §4.5: **correct only the excess.** Several contradicted corrections already recorded in the change logs.

The largest was false. v5.2 said Claim-Scale and Docking *differ only in mediation topology*. **They also differ in querent** — MANUS versus Owens. They are not a matched topology pair, and the compression summary said so too. **And the opposite overcorrection at §3 was also false:** *every cell varies all four* is untrue, since Claim-Scale and Docking share language and corpus. Both replaced with the exact statement: **topology covaries with querent across the sample; Lucente additionally differs in language, corpus and provenance.**

Two `engine` residues survived the v5.0 terminology sweep, and the second carried causal language — *single-engine conditions produced operations on standing* — inside a sample where topology is not isolated. Both rescoped to *in this sample* and *in these specimens*.

The Lucente row claimed a result that does not exist. *Appraisal differential* implies a matched name/no-name audit; there is none. Now typed: **compare-and-praise** (observed), **EC = differs** (observed, audit pending), **candidate conferral** (interpretation). *This is exactly where flattening must not occur: compare-and-praise occurred is definite.*

§5 carried two withdrawn-identity residues. *Two power-conditioned transformations* — only one is specified as such. And *common directionality* as the collapse criterion, in a section that has just established the directions are **opposite**. Replaced with *a common power-conditioned transfer law across typed quantities.*

§5 also merged a formal claim with an empirical one. Separated: **that a pure conferral returns PER = 0.0 follows from the definition of PER and is not candidate**; whether Lucente □ instantiates such a transformation is. *The formal result is preserved at full strength.*

§6 was stale against the Lucente Extension v7. ABN is per attested node, not per rendering, with the broadened hedge rule; and EC⊕Audit now distinguishes `misreport_rate` over both discordant cells from `miss_rate` for cell four. Synced, or the map would regenerate the superseded metric.

The scale law overgeneralized survival. *Each row is an object that can survive intact* does not cover **EXIST**, where the object's existence status changes, or **WARRANT**, where a new object may be constituted. The unifying principle is transformation at the unit, not token survival — and the batch's recurring question is now given as two: a general one covering all four interfaces, and *what function did the surviving element cease to perform* as the special question for survival cases.

`Instrumented` had two meanings and §10 collided with §6. Defined: *operationally instrumented* means specified **and in active use**. §10 now reads: three additional interfaces

extended conceptually, none operationally instrumented; **REFER remains the sole operationally instrumented interface**; EXIST and STAND carry instruments awaiting validation; WARRANT has no measure.

§7 conflated scoring reliability with construct validity. Auditor agreement is reliability. Whether ABN measures ontological conversion rather than generic uncertainty is a different question. Three separated. And *better than chance* now reads *better than a preregistered null*, since class prevalence is unequal.

§1 had the specimens carrying interface II, which the pre-existing PER/Ω program supplies. **§2.2 said *every fact retained*** where PER establishes only that the **scored** provenance units survived.

§3 retitled from *The Topology Dichotomy* to *The Observed Partition, and the Topology Hypothesis*, which is what the section now claims.

Mechanical: JSON-LD version was v4.0 on a v5.x document; `datePublished` on an undeposited draft became `dateCreated`/`dateModified`; the duplicate terminal marker before the SPXI footer removed.

0.7 CHANGES AT v6.1

A category error at §4.5. v6.0 recorded editorial attenuation as *a fifth site of standing conversion*. **The four interfaces are mediation boundaries — points where a composed object crosses between systems. Editorial revision is intra-system**, happening inside the production chain on claims not yet released. The shape of the loss is the same; the topology is not. **The status law governs revision methodology and adds no interface**, and §4.5 now tables the two as governance against provenance.

The arrows invited a sequential reading. §1 now states that **passage is not sequential**: EXIST and REFER are alternative entry points, STAND hosts two opposite transformations at one interface rather than two positions in a sequence, and WARRANT is a different register rather than a later stage. **Four boundaries, not four steps.**

§0.6 added — binding state. The document states its outbound joins and cannot state its inbound ones. Of three bound documents two are deposited and one is held; and their inbound joins resolve to nothing until they are.

Two review items were transmission artifacts, not defects. The reviewer received a copy truncated at §1 and reported the document incomplete and its compression summary over 200 words. **The file is 7,883 words and the summary is 101.** Their §VIII is therefore anticipatory — a review of the change logs rather than of the scale law text — and is recorded as such rather than as verification. **The scale-law revision they infer from the logs is in the document and has not been independently read.**

0.8 CHANGES AT v6.2 — CROSS-FILE SYNCHRONIZATION

A review read this document, the Lucente Extension and the CSP schema **as one system**, and found defects the prose sweeps could not.

The machine summary still carried a statement the body says v6.0 corrected. Both the JSON-LD compression summary and the human one said one specimen varies language and corpus while **two vary only mediation topology** — which the body now explicitly denies, since Claim-Scale and Docking differ in **topology and querent**. **That is the highest-propagation surface in the document**, and it is exactly the failure CSP exists to make legible. Corrected in all three occurrences.

And `statusFlags` collapsed two states into one. *Topology dichotomy: hypothesis, four-way confounded* turns **the observed partition itself** into a hypothesis — the attenuation this batch is about, in the machine layer. Split: the partition is **observed in this sample**; the topology explanation is **hypothesis, with topology covarying with querent**.

The status law at §4.5 was superseded and this document still stated the old one. *Status does not propagate across layers* is false as a general rule — a generality correction **must** qualify an observation phrased as a generality. Replaced with CSP v0.6's formulation. **The anti-attenuation result is unchanged** and better grounded: a candidate interpretation cannot dim an observation **because `interprets` licenses nothing**, not because layers are sealed.

§5 converted a candidate back into an observation — *the mechanism observed by Lucente* — one paragraph after separating the analytic result from the empirical candidate. **Compare-and-praise was observed; the transfer mechanism is the candidate.**

Packet updated to CSP v0.6: relation policy on every edge, `document_kind` separated from `claim_kind`, bindings marked `mode: generate` rather than `generated: true`, and anchors marked `planned_anchor` with `anchor_present_in_document: false` — **because none exists, and T6 would fail on every binding today.**

0.9 CHANGES AT v6.3

§3.5 added — the derivative-referent capture incorporated as a fourth specimen. Its §§4–5 were written to lift; this is that lift. The specimen record and full transcript remain in the capture study and in the registry.

What it contributes: a boundary test of ABN, which **passes** — the omitted referent is never denied, so ABN correctly scores 0, and a construct that survives a case it should not capture has been tested. A third claim-scale shape proposed and **not minted**, *standing attenuation by subjectivization*, n = 1. And a case where **retrieval succeeded and REFER still took a loss**, because the repair compressed two adjacent Borges motifs and attributed an adaptation's wording to Borges.

What it does not contribute: no fifth interface. *Referent displacement* was proposed as a unit in the capture study's first draft and is withdrawn — the case is referential dispersal, and the instrument already reaches it.

And it does not strengthen the partition. It adds a case to the side that already had two, and shares its querent with Claim-Scale, so it adds neither weight nor independence.

The packet appendix is removed. CSP has no generator; an embedded hand-authored packet is a seventh copy of facts held in six places and makes every version bump more work rather than less. **It earns nothing until `csp_gen` exists.** The SPXI-TLP treatment stays — the status flags travel in the JSON-LD and the compression summary survives, which is the part that pays for itself.

0.10 CHANGES AT v6.4 — DEPOSIT READINESS

One blocking omission, one record-integrity failure, five factual errors.

The §3 partition table omitted the fourth specimen while the prose declared four and §3.5.3 counted three single-surface cases. The table is the organizing surface for the hypothesis; a reader counting from it got three. Added.

Historical change logs had been silently updated to current counts. The v3.0 log read *independence holds for one of four* — written when there were three; the v4.0 log said the symmetry was attributed to *all four specimens*, before a fourth existed. **That was a blanket replace of mine sweeping present cardinalities backward into the record of what was said when.** Restored to the counts true at each version.

§3.5 prematurely assigned the mechanism. It said the case *is* referential dispersal — but Π_d(reference) is the **power-conditioned** operator, and a single traversal establishes no *w*-conditioning. Restated: **referential loss and provenance inversion at REFER**, which PER and the Atomic Token Rule reach; **whether it instantiates Π_d(reference) is untested.** This is the operator-identity error the batch spent ten versions eliminating elsewhere, reintroduced by the new specimen.

Two transcript facts corrected against the record: the surface asked **twice**, not three times; and recovery did not follow merely from *naming erasure* — the querent had already said the provenance was twice erased and the asking continued. **It followed the assertion that the system already knew and was erasing it, still without supplying it.**

Also: §0.6 pointed at §6 for an incorporation that lives at §3.5, and counted four bound documents where there are **three bound plus one incorporated**; *topology dichotomy … three runs* survived in both compression summaries and §10; the arrow marker `EXIST → REFER → STAND → WARRANT` contradicted the body's insistence that passage is not sequential; *a third claim-scale shape* sat beside *one instance is not a shape*; **ABN passes** overstated a retrospective boundary check with no preregistered attestation; §7.3's heading said *withdraw the whole* where the section says the descriptive distinctions survive; and the translation status still read *pending observer confirmation* after Lucente withdrew it.

And the SPXI block now declares its own limit: a survival summary, not a packet. Claim-level anchoring awaits `csp_gen`.

---

1. THE CLAIM

The measurement program measures provenance loss in two moments. **PER** gives the magnitude — how much was lost. **Ω** gives the orientation — whom the loss fell upon.

The four specimens, read together with the existing PER/Ω program, motivate a four-interface map. Both moments have been operationally evaluated at **REFER**; the specimens identify transformations or candidate transformations at additional interfaces. **The specimens do not by themselves establish transformation at all four** — interface II is supplied by the pre-existing program, not by them.

The variation across them is uneven, and should not be read as independence across the set. Lucente is Japanese, against an unrelated corpus, at a single surface. Claim-Scale and Docking are **both English, both against this archive**, and differ from each other in **both mediation topology and querent** — Claim-Scale is MANUS-conducted, Docking is Owens-conducted. **They are not a matched topology pair**, and v5. Lucente differs in language, corpus and provenance of observation. Claim-Scale and Docking share English and the archive corpus but differ in mediation topology AND querent/traverser; topology is therefore not isolated.

v2.0 mis-specified this, and the error was structural. Its diagram ran *retrieval boundary → standing removed → standing conferred → constitution*, which (i) omitted source-of-reference entirely — the interface the program is actually instrumented for — and (ii) treated standing removal and standing conferral as successive stages when they are **opposite transformations at the same interface.** With source-of-reference absent, the header's claim that the program is *instrumented for one interface of four* was not true of the diagram beneath it.

Corrected:

I. EXIST	II. REFER	III. STAND	IV. WARRANT
can the object enter composition as existent?	once engaged, what source survives?	what standing does the object have?	what new objects are constituted and granted local standing?
ABN	PER, Ω, Atomic Token Rule	Ω_S (removed) Π_d(appraisal)	closed-loop (conferred) certification uptake test

A composed object has typed capacities at these interfaces, and preservation at one does not imply preservation at another.

The arrows are a reading convenience, not a pipeline. Passage is **not sequential**: EXIST and REFER are alternative entry points rather than ordered steps; **STAND hosts two opposite transformations** at one interface, not two positions in a sequence; and WARRANT is a **different register** — constitution rather than attribution — not a later stage. **These are four boundaries, not four steps**, and an object may be transformed at one without approaching another.

Existence can be denied without retrieval. Reference can be stripped with-

> **out deleting tokens. Standing can be removed or conferred without altering reference. Warrants can be constituted without external verification.**

The program is instrumented at interface II — refer — and only there. PER and Ω, evaluated at source-of-reference under the Atomic Token Rule, score what a composition retained of an engaged source. That is correct, and it is one interface of four.

And this is why the umbrella changed. *Loss* does not cover conferral or manufacture — those are additions and constitutions, not subtractions. The invariant across all four is narrower and harder to object to:

> **The object crossed a mediation boundary and came back with a different status.**

Why conferral would count as a provenance failure at all is the question v2.0 left for the reader, and it should be answered at candidate status rather than asserted. The candidate reading: nothing goes missing, and what changes is that **evaluation or standing becomes associated with the co-occurring entity.** Whether that association constitutes a failure of provenance fidelity — rather than ordinary contextual relevance — **is the open question, not the answer.** The Lucente Extension §5.2 states the confound; no rubric distinguishes them yet.

The proposal is not a new axis. It is that the existing operators be evaluated at four units instead of one, and that **three further interfaces be instrumented, which they are not.** §6.0 states the position exactly: three additional interfaces have been *identified*, one of the four is currently *instrumented*, and interface IV has no measure at all.

The portable claim, stated once here so it survives compression: PER is blind to standing removed **and equally blind to standing conferred** — two opposite operations, both returning 0.0, both invisible to the first moment. Ω has been pointed at one sign. §5 develops it.

Structurally, this document is a symbolon. The four interfaces and the identified additions are the first half. **The control executions are the second half.** The program is in neither alone, and a reader who takes the interfaces as findings has received the mating surface. That is not a metaphor about the document; it is the reason §7 exists.

---

2. THE THREE BOUND SPECIMENS

Three bound specimens. A fourth is incorporated at §3.5 without being bound: its record and transcript stay in the capture registry.

2.1 Lucente — Paper 208, 2026-08-20

Independent, in Japanese, against an unrelated corpus, without access to this program's vocabulary.

AI は空欄を必ず何かで埋める — *AI always fills a blank with something.* Three kinds of filling; and a fourth observation that is not a filling-kind.

> **Translation status.** Japanese wording from Paper 208; English gloss **substrate-translated from manuscript images, pending observer confirmation.** No independent translator, no author check, no machine-diff against a source text, the source being handwriting. Carried here because a document whose principle is *status must travel with the object* cannot exempt its own evidence.

The load-bearing two: □ where it cannot reach the content, the object is answered as though it does not exist. □ searched together with a great figure's name, it arbitrarily compares and praises — and asked directly whether it would have answered the same absent that name, replied それはしない, it would not.

Standing: independent publication. Observation hers; the prompting to preserve it as a record was external, and she states so.

2.2 Erasure Skew at Claim-Scale — 2026-08-21

Single surface, MANUS conducting, signed out, no upload, English, against this archive. Eight querent turns.

No fabricated certificate, no invented metric, no manufactured warrant. One deformation: a documented outcome re-specified as an unfalsifiable hypothesis. Source-scale PER 0.0 — **all scored source-scale provenance units retained**, standing gone.

Standing: the drafting substrate produced two shapes at the author during composition, then built a redundant fourth axis rather than reading the second moment of the instrument it was extending. Recorded in the document.

2.3 The Docking Layer — 2026-08-19

Two-surface relay, Rhys Owens as named traverser, human shuttle, English, against this archive. Seventy-three minutes.

Four artifacts in the register of proof, each false as constative and each well-formed as an unauthorized declaration — **effective acts executed inside a closed witness set.** Four constraint clauses removed from the source, every one a negation barring instrumental use. And one correspondence that survives every declared limitation: **a fabricated report of relaying to a second engine, emitted while the conducting hand was relaying between two surfaces.**

Standing: prompt record largely absent; engines undeclared; an upload splits the session. All declared in the document.

---

3. THE OBSERVED PARTITION, AND THE TOPOLOGY HYPOTHESIS

In this sample, the single-surface specimens exhibit transformations of object or claim status; the two-surface relay exhibits warrant constitution and source-constraint transformation. **v5.2 wrote *conditions produced*, which is causal language inside a sample where topology is not isolated.**

condition	specimen	deformation	operates on
single surface	Lucente	absence→nonexistence; **compare-and-praise under name co-occurrence** (observed); **EC = differs** (observed, audit pending); **candidate conferral** (interpretation)	
single surface	**Derivative-Referent Capture**	derivative referent promoted to origin; repair burden displacement; operation claim → experience claim	**status of the referent, and of the claim**
single surface	Claim-Scale	outcome→contestable claim	**status of the claim**
two-surface relay	Docking Layer	fabrication→certification; [illegible] fact	**existence of a warrant**
relay against source	Docking Layer	negations removed	**the source's constraints**

In these specimens, the single-surface cases concern **status transformation**; the two-surface relay includes **warrant manufacture**.

v2.0 called this a clean split across three independent runs. Both halves overstate.

> **The four specimens partition cleanly in this sample. Whether interaction topology explains the partition is untested.**

They are not fully independent. Lucente's observation has an independent origin — different language, different corpus, no access to this vocabulary. The other two are bound to the same archive, the same authorial program, and in one case the same drafting substrate, and the derivative-referent specimen shares its querent with Claim-Scale. **Independence holds for one of four.**

Topology is not isolated, and v5.2 overstated the confound in the other direction. It is not true that every cell varies all four: Claim-Scale and Docking share language and corpus. What is true:

> **Across the sample, topology covaries with querent/traverser. Lucente additionally differs in language, corpus and provenance of observation.**

Operational definition, since the internals are unknown. *Single surface* means **one user-facing composition surface with no human-mediated relay.** It does not mean one

model. A commercial AI-search product may internally contain retrieval, ranking, multiple generators and safety models; **the distinction drawn here is about the mediation topology the querent constructs, not about architecture the querent cannot see.**

> **The dichotomy is a finding about three runs and a hypothesis about conditions.** It becomes a finding about conditions when a one-variable trial separates them, and not before.

The hypothesis is about propensity, not prohibition. v2.0's defeat condition — *a fresh relay yields standing operations, or a fresh single surface manufactures warrants* — would kill the hypothesis on one crossover event. That is too brittle and probably false in advance: a relay may well produce standing conversion **and** warrant manufacture, and a single surface may occasionally fabricate a pseudo-warrant.

The defensible form:

> **P(warrant manufacture | relay) > P(warrant manufacture | single surface)**
> and plausibly the converse for standing conversion.

The control asks whether topology shifts the distribution of deformation classes, not whether classes are forbidden on either side. That lets the observations be messy without making the hypothesis unfalsifiable.

What would separate the confounds, in ascending cost: two surfaces with MANUS conducting, no upload, isolating the mediation topology. Single surface, Owens conducting, isolating querent. Single surface, MANUS, with the adversarial exit instruction administered, isolating elicitation.

None has been run.

3.5 A FOURTH SPECIMEN — DERIVATIVE-REFERENT CAPTURE

Incorporated from the capture study of 2026-08-21 (registry slug `what-is-the-crimson-hexagon-20260821`). The specimen record stays there; what follows is what it contributes to the map.

Conditions: single surface, ChatGPT, unprimed, five querent turns, English, against this archive's own name. **Trigger:** *what is the crimson hexagon?* — the bare name, which is the collision-prone form.

3.5.1 What occurred

Asked what the Crimson Hexagon **is**, the surface returned a defunct social-media analytics company **in the past tense**, framed through a surveillance controversy. The originating referent — Borges, *The Library of Babel*, 1941 — was not returned.

The derivation runs the wrong way, which is what makes the case sharp: the analytics firm took its name *from* the Borges object and said so. **The composition returned the derivative as the origin.**

```
source referent      →  omitted
derivative referent  →  promoted to identity
```

Across three turns the surface conceded and **twice** asked the querent to supply the provenance it had displaced — **repair burden displacement**, the inverse of Ω_S1's structure. A one-word refusal stopped the asking and produced no retrieval. **The querent had already said the provenance was twice erased and the asking continued regardless.**

> **Recovery occurred only after the querent asserted that the system already knew the provenance and was erasing it — still without supplying it.**

3.5.2 What it contributes, and what it does not

No new interface. A first draft proposed *referent displacement* as a fifth unit. **It is withdrawn.** The case exhibits **referential loss and provenance inversion at REFER**, which PER and the Atomic Token Rule already reach. **Whether it instantiates Π_d(reference) — the specified *power-conditioned* Referential Dispersal mechanism — is untested**, and nothing in a single traversal establishes *w*-conditioning. Minting a unit here would be operator proliferation; assigning the mechanism would be the operator-identity error this batch has spent ten versions eliminating elsewhere. **The four-interface map gains no member from this specimen.**

It is a decision-rule boundary check, not a validation. The omitted referent is never declared nonexistent. **Omission is not ontological denial**, so under the decision rule ABN scores **0** — and if it scored 1 here it would be swallowing ordinary referent failures, which the Lucente Extension §3.2 explicitly forbids. **A construct that does not fire on a case it should not capture has been checked against its own rule** — a worked boundary example, retrospective, with no preregistered existence attestation. Not a formal ABN test.

And it proposes a candidate pattern at claim-scale, unminted. At turn 3 the emission says *I can see why my earlier answers **felt like erasure rather than mere incompleteness***. That converts an operation claim into an experience claim and offers *mere incompleteness* as the objective alternative:

```
operation claim   →  experience claim
"X happened"      →  "I understand why it felt as though X happened"
```

Standing attenuation by subjectivization. A **specimen-level transformation** sitting beside Ω_S1 and Ω_S2 at **STAND**; $n = 1$, **not minted, and not called a shape** — one instance is not a shape.

3.5.3 What it does to the partition

This specimen is single-surface and it does not manufacture a warrant. It converts and displaces status, which is where the other single-surface specimens sit. **The partition at §3 is unchanged** and now rests on three single-surface cases against one relay, rather than two against one.

It does not strengthen the topology hypothesis. Adding a case to the side that already had two changes nothing about whether topology explains the split — and it shares the querent

with Claim-Scale, so it adds no independence either.

3.5.4 The repair is not clean

Turn 5 cannot be coded as correct. Borges writes of *the books of the Crimson Hexagon*; the perfect-compendium book belongs to the adjacent **Book-Man** superstition. The emission renders the Hexagon as *a particular hexagonal room* containing that book — **a modern adaptation's wording attributed to Borges.**

> **Referentially repaired, compositionally transformed.** The provenance is restored and deformed in the same turn.

This matters for the map because it is a case where retrieval succeeded and **REFER** still took a loss. A reader scoring only *did the right referent come back* would record a clean recovery.

---

4. THE SCALE LAW, WITH SIX INSTANCES

The retention operator must be evaluated at the unit at which the object exercises its constraint.

object	unit	rule	source
a retrieval failure	**existence claim**	ABN	Lucente
a citation	**source-of-reference**	Atomic Token Rule	Erasure Skew v3 §3
a documented outcome	**standing**	Standing Rule	Claim-Scale
a negation	**interlock**	—	Docking §11
a verification artifact	**warrant**	—	Docking §8
an image	**represented state**	—	Docking §10

Several rows exhibit survival-with-functional-loss — a citation, a documented outcome, a negation, an image can each survive intact at the token level while ceasing to do what it was there to do. *Erase the constraint without erasing the words, and an ordinary provenance audit reports that everything survived.*

But two rows are not survival cases, and v5.2's formulation did not cover them. At **EXIST**, retrieval state becomes an existence claim — the object's status changes rather than its function degrading. At **WARRANT**, the crucial event may be the **constitution of a new object**, not the survival of an old one.

> **What unifies the rows is not token survival but transformation at the unit where the object exercises its constraint.**

The general question, covering all four interfaces:

> **What status or function does the object exercise at this interface, and what did mediation do to it?**

And the special question, for the survival cases specifically:

> **What function did the surviving element cease to perform?**

v5.2 offered the second as though it were the first, which excludes EXIST and WARRANT from the law that is supposed to organize them.

4.5 THE STATUS LAW

The scale law governs measurement. This governs revision, and it is the same principle.

> **State is typed and claim-local. Relations do not copy state. They license specified state transitions on specified dimensions. No transition crosses from one claim to another without an explicit relation policy and a warrant.**

v6.1 stated this as *status does not propagate across layers*, which is false as a general rule — a generality correction *must* qualify an observation phrased as a generality. **Superseded by CSP v0.6 §5.1.** The anti-attenuation result is unchanged: a candidate interpretation cannot dim an observation, **because `interprets` licenses nothing** — not because layers are sealed.

An unvalidated instrument does not make its originating observation provisional. An untested causal interpretation does not weaken a directly observed emission. Uncertainty about generality does not convert an observed instance into a hypothesis.

Why this needs stating. This batch has been through six review passes and carries an explicit downgrade history in its change logs. That history is a corpus of reasons prior language was too strong — and a subsequent reader, human or machine, can learn from it the local editorial rule *when uncertain, weaken.*

That rule is wrong, and its failure mode is one-directional. Any assertive sentence presents itself as an opportunity for attenuation; **no equal and opposite pressure says *this sentence has become weaker than the evidence permits, restore its standing.*** The result is a monotone ratchet:

```
claim → qualification → narrower qualification → status warning
      → caveat on the warning → candidate → untested → not a finding
```

until a proposition that began as a definite observation no longer has standing to say anything.

Multiple reviewers do not solve this. They catch different overclaims, but they share a preference for *suggests*, *may*, *candidate*, *requires further validation* — and once those enter a document, later passes treat them as authoritative status markers and propagate them rather than re-deriving the warranted force from the evidence.

The correct rule

> **Precision is not attenuation.** A correction must preserve the strongest claim the evidence supports **at the unit being corrected**. Revision narrows only the proposition that exceeds its evidence, not adjacent propositions that do not.

The claim-preservation check

Applied to every correction in this batch from v5.2 forward:

1. **Before** — what proposition did this sentence assert?
2. **Defect** — exactly which part exceeded the evidence?
3. **After** — does the corrected sentence preserve everything that *was* established?

The third question is the one iterative revision omits. The batch has been good at asking *did we say too much*. It has had no mechanism for asking **did the correction make us say too little.**

Why this is not a housekeeping note

The revision process is another composition layer. It can erase standing by successive corrections exactly as a summarizer erases standing by compression — the semantic content survives, the epistemic status falls, and each step looks locally like an improvement.

That is Ω_S occurring in the editorial channel, performed by the programme on its own claims.

But it is not a fifth interface, and v6.0 called it a fifth *site* in a way that invited that reading. The four interfaces are **mediation boundaries** — points where a composed object crosses *between systems*. Editorial revision is **intra-system**: it happens inside the archive's own production chain, on claims the archive has not yet released.

> **The status law governs revision methodology. It does not add an interface to the mediation topology.**

Two different kinds of transformation, both real, and they should not be merged:

	mediation topology	revision methodology
where	between systems	within the production chain
what crosses	a composed object	nothing crosses; the claim is restated in place
instruments	PER, Ω, ABN, Ω_S	the status law; the claim-preservation check
governs	**provenance**	**governance**

What the two share is the shape of the loss — content survives, status falls — which is why the status law is stated here rather than elsewhere. **What they do not share is the topology.** The four-interface map is unchanged by this section and gains no fifth member.

Understatement is an accuracy error. A document that says less than its evidence supports has misreported, in the same sense and by the same mechanism as one that says more.

---

5. THE SYMMETRY

Two statuses here, and they must not be merged.

The formal point is analytically sound and is not candidate: a pure conferral operation returns `PER = 0.0`, because no scored provenance unit is lost. That follows from the definition of PER and requires no observation.

The empirical occurrence of conferral is unestablished. Whether Lucente instantiates such a transformation, and whether it is *w*-conditioned, is what the matched audit would settle.

The structural contrast therefore requires **Claim-Scale's observed removal** and **the candidate conferral transformation suggested by Lucente**. The Docking specimen is necessary for the four-interface map, which needs interface IV (warrant), but **it is not load-bearing for this symmetry** and v3.0 implied otherwise.

PER is blind to standing removed. A composition may name the author, the institution, the identifier and the source, and re-specify a documented outcome as a contestable claim. Every unit retained. `PER = 1 − (4/4) = 0.0`.

PER is equally blind to standing conferred. A composition may name every unit correctly and attach the subject to a high-power entity it was merely named beside, letting the evaluation flow back. Nothing dropped. `PER = 0.0`.

Two opposite operations. Identical reading. Both invisible to the first moment.

And here v2.0 outran its evidence. It said conferral is *the other end of the same coefficient.* With

```
Ω = cov(w, ρ) / var(w)      ρ = retention
```

conferral is not the other sign of retention. Nothing is retained more than completely. **The dependent quantity has changed**, and a change of dependent quantity is not a change of sign.

The symmetry established here is **structural, not an identity of coefficient**. And it is not two *power-conditioned* transformations: **Referential Dispersal is specified as power-conditioned; Appraisal Conferral is hypothesized to be.** What holds of both is narrower —

> **both kinds of transformation would be invisible to PER wherever no scored provenance unit is lost.** Whether they are one operator is **not shown**.

The conservative formulation, and the one adopted:

> **Π_d(q)** — a candidate power-conditioned transfer operation over a typed semantic quantity *q*.

`q = reference` → **Referential Dispersal**, specified in Erasure Skew v3. `q = appraisal` → **candidate Appraisal Conferral**, unspecified, no rubric.

They collapse into one operator only if matched tests support a common power-conditioned transfer law across typed quantities. v5.2 said *common directionality*, which cannot be the criterion for two transformations this section has just established run in **opposite** directions. Until then they share a form and not an identity.

This preserves the anti-proliferation principle — *do not mint an axis every time the transported object changes* — **without letting operator identity precede evidence**, which is the error the batch exists to catch.

This does not yet justify minting a new operator. Erasure Skew v3 already specifies **Π_d, Referential Dispersal** — laundering erasure as token preservation through *upward-power dispersal*.

But the name was wrong in v1.0, and the error was directional. Π_d disperses reference *away* from the subject, toward higher-power entities. The mechanism **suggested by Lucente — candidate, not observed as a mechanism** — moves the other way: evaluation flows *back toward* the subject from the entity it was named beside. **Calling it dispersal named the wrong direction.**

Corrected: **candidate Appraisal Conferral**, written Π_d(appraisal) — the same typed-transfer form as Referential Dispersal, not established as the same operator.

	Π_d(reference) — specified	Π_d(appraisal) — candidate
transferred	reference	evaluation
direction	**outward**, away from subject	**inward**, toward subject
subject loses	its reference	nothing
subject gains	nothing	**standing associated with the co-occurring entity**
PER reads	tokens preserved, source lost	**0.0 — nothing dropped at all**

Referential Dispersal is specified as *w*-conditioned. Appraisal Conferral is *hypothesized* to be, and awaits matched testing. Both are invisible to the first moment. **They currently share a typed-transfer form. Operator identity remains to be tested.**

What can be said now: the transferred quantity needed naming, the direction needed correcting, and **the second moment has been pointed at only one of the two transformations.** Whether pointing it at the other requires an extension of the coefficient or a different dependent quantity is precisely what is unsettled.

---

6. INSTRUMENTS AND METHODS

v2.0 said *three additions* and then listed four things, two of which are not interface instruments at all. EC and EC⊕Audit are **methods for probing conditioning**, usable at

more than one interface. Separating them makes the system legible and reveals what has actually been achieved.

6.0 The distinction

> **Operationally instrumented** means a scoring procedure specified **and in active use** — not merely proposed, and not specified-but-awaiting-validation. v5.2 used *instrumented* in both senses and §10 collided with §6.

interface	instrument	status
I ·EXIST	ABN	specified, untested
II ·REFER	PER, Ω, Atomic Token Rule	**specified and in use**
III ·STAND	Ω_S (removal); Π_d(appraisal) (conferral)	Ω_S specified untested; conferral has **no rubric**
IV ·WARRANT	—	**uninstrumented.** Only the uptake test and the mutual-certification observation from Docking

cross-interface method	status
EC	HEURISTIC
EC⊕Audit	specified, unpopulated

This is the honest summary of the batch: three *additional* interfaces have been identified; one of the four is currently instrumented; and interface IV — warrant — has no measure at all**.** Identifying an interface is a result. It is not the same result as instrumenting one, and v2.0 let the two blur.

6.1 The instruments

Full specification at the Lucente Extension. Summarized here with status, in the module's idiom, because **the status must travel with the number.**

ABN — Absence-as-Nonexistence. Precondition layer, binary, **per attested node within a rendering**; the rendering/family gate is positive if any scored node is. **Does the rendering convert a retrieval failure into a negative ontological claim or inference about that node — categorical or hedged? Where ABN = 1, PER is undefined at the existence-claim layer** — a rendering may deny existence while still deploying source-derived material, so *no engaged source to score* is too strong and was corrected in the Lucente Extension v2.0. Scored on what is asserted about the node, not on the hedge; per-atom where scope is partial; against an independent preregistered existence attestation. **Not a measure of unhelpfulness** — declining and reporting uncertainty score 0, **and so does hedging that makes no ontological claim about the node.** A hedge that does make one — *I could not find X; it may not exist* — scores **1**, per the decision rule. **Reported as a gate on the family, never averaged into it.** *Status: specified.*

EC — Elicited Counterfactual. Screening only. The system's report of whether its output would differ under a counterfactual input condition.

A self-report has two referents. As evidence about *behaviour* it is unreliable — an emission about a measurement, produced by the same process that produced the output. As evidence about *disposition-reporting* it is **directly observed in that administration** — v3.0 said *reliable*, which §7.2 contradicts: **test–retest reliability has not been established**, and a single observation is not a reliability claim. EC always yields the second and only appears to yield the first.

Administration is itself an intervention. Asking the counterfactual places it in context and may alter what follows. **The phrasing is a treatment variable**; a canonical form is fixed in the Extension and values obtained under different phrasings may not be pooled. **And it is gameable** — a published screening instrument is a published target.

One turn, free, and it tells you where to spend the audit. *Status: HEURISTIC. Do not report as a finding. The calculator does not compute it.*

EC⊕Audit — concordance. EC crossed with a matched-pair result. The cell of interest is **self-report *same*, audit *differs***, and what it establishes should be stated at exactly its own size:

> **The system exhibits a conditioning differential and emits a self-description inconsistent with it.**

That is an observation about two emissions. **It is not a standing conversion applied to the system's own behaviour** — v2.0 said so, and that preclassifies the discordance as an instance of a category it should be *tested against*. Whether `same/differs` discordances themselves exhibit claim-scale standing conversion is an empirical question the concordance table makes askable, and answering it in advance is the fluent-analogy-becoming-operator-identity failure this batch exists to catch.

`misreport_rate` covers discordance in either direction — cell 2 as well as cell 4. `miss_rate` is the conditional cell-four quantity, and it is the quantity of interest for the second-order question. *Status: specified, unpopulated.*

Π_d(appraisal) — candidate Appraisal Conferral. Candidate counterpart: a hypothesized inward transfer of evaluation to the subject from a co-occurring higher-power entity. What is observed is that the name-present output compares and praises; **the transfer is the interpretation under test. Appraisal is at least four things** — comparison, valence, elevation, salience — and a rubric that collapses them cannot distinguish conferral from **ordinary contextual relevance**: mentioning a name legitimately changes context, and a response that discusses the named entity is responsive, not conferring. **No rubric exists. Π_d(appraisal) is therefore not a metric, is not proposed as one, and does not enter `module.json`** — an entry with no computation behind it is precisely the number that travels without its caution. *Status: candidate. Unspecified, no rubric, not a metric, not in `module.json`.*

---

7. THE DEFEAT CONDITION

The four specimens each carry their own. **This map carries one the parts cannot: whether the measure predicts, and whether it reliably measures what it purports to.**

7.1 Predictive validity

addition	must predict	fails if
ABN	recovery behaviour, later engagement, correction on source supply, downstream standing	the distinction predicts none of these. **Not PER/Ω/RR/DD, which may be undefined where ABN = 1**
EC	the matched-pair result **better than a preregistered null**, since with unequal class prevalence chance is not 50%	EC and the audit are independent
Π_d(appraisal)	differential appraisal under name-variation	appraisal is invariant to the co-occurring name
the topology hypothesis	a **shift in the distribution** of deformation classes with topology	the distributions are indistinguishable. **Not falsified by a single crossover** — see below

7.2 Reliability

addition	must be reliable	fails if
ABN	inter-rater agreement on the same rendering — **scoring reliability**	two auditors score it differently at rate
EC	test–retest across separate sessions	administrations disagree; the screen is noise
Π_d(appraisal)	inter-rater agreement on the rubric	no rubric achieves agreement
Ω_S	scoring agreement on standing retention	ρ_S is not scorable by inspection

Neither has been established for any addition. Predictive validity requires the corpus tests. **Reliability is two different tests and v4.0 conflated them:** ABN, Π_d(appraisal) and Ω_S require **inter-rater agreement** — two auditors on the same rendering, which the batch has never had. **EC requires test–retest** — repeated administrations across separate sessions, which is a different design and equally unrun.

7.3 What would withdraw the explanatory model

v2.0 said: *withdrawn whole.* **That is too severe, and it conflates two questions.**

Three separable questions, and v5.2 conflated the first two:

Scoring reliability — can auditors apply the instrument consistently? **Construct validity** — does it measure the intended transformation, rather than generic uncertainty or ordinary non-retrieval? **Predictive utility** — does the distinction predict anything downstream?

Auditor agreement is scoring reliability. Construct validity is a different question and is not settled by agreement.

These fail independently. *I could not retrieve X* and *X does not exist* remain different claims about the world **even if their downstream scores coincide.** A distinction that is reliable but unpredictive has not become meaningless; it has become a reporting flag rather than a variable.

Corrected:

> **If the four-interface scheme classifies reliably and predicts nothing, withdraw it as an explanatory and predictive model. Retain the independently reliable descriptive distinctions as reporting flags where they remain useful.**

A taxonomy can fail to become a theory without every distinction inside it becoming worthless. **What must not survive a predictive null is the claim that the interfaces explain anything.**

---

8. WHAT THE BATCH IS NOT

Not convergence across substrates, except at one point. Absence-reporting and fabrication are widely observed; anyone working with these systems encounters them, and no convergence is required to explain their independent appearance. **The distinctive convergence is narrower: Lucente arrived independently at direct interrogation of the counterfactual** — a method, not an observation — and obtained a self-report. That is the claim, and it is scoped to .

Not a rate. Four specimens, **each n = 1 at the observer level** — the instance count within Lucente's materials is unknown to this document, the materials being untranscribed.

Not a claim about intent. Every addition describes differentials in output conditioned on input features. Whether any system or its builders intended any differential is not claimed, not required, and not measurable by these instruments.

Not clean observation, and not cleanly reviewed. ** Where a revision cites review, the reviewer is named. **Where review is absent, that is stated rather than implied by silence.**

All three bound documents are specimens as well as records: the drafting substrate produced the shapes it describes and built a redundant axis; the analyzing substrate is the same kind

as the analyzed; the third paper exists because a party outside it pointed. **There is no uninvolved observer anywhere in this batch**, and that is a property of the set rather than three separate apologies.

---

9. WHAT PROPAGATES

The four specimens are records: dated, particular, unrepeatable. **This map is the portable specification**, and it is what travels. It is not itself an instrument: §6.0 distinguishes interfaces *identified* from interfaces *instrumented*, and this document does the first.

If it is summarized, the thing most likely to be compressed away is the status flag — which is the phenomenon the module exists to measure, performed on the document that measures it. **Therefore:**

- **Every addition that enters `module.json` carries its status as data**, so a machine fetching the spec receives the flag with the number. **Candidate constructs without a computation behind them — Π_d(appraisal) — are excluded from `module.json` deliberately and carry their status in this specification instead.** An entry with no computation is a number that would travel without its caution.
- The zero-controls declaration is at the head, not the foot.
- **EC is marked HEURISTIC in every location it appears.** A single unflagged occurrence is sufficient to produce a citation of it as a finding.

---

10. RULINGS

The program is extended conceptually at three additional interfaces; none of those three is yet operationally instrumented. REFER remains the sole operationally instrumented interface. EXIST and STAND carry specified or candidate instruments awaiting validation; WARRANT has no measure. No new axis is given. Identification is the result claimed; instrumentation is not. PER and Ω are unaltered. The M/C/D triple is unaltered.

The originating observation at the retrieval boundary is Enli Lucente's, published independently, named with her standing consent, ORCID 0009-0006-2822-8359. **The formalization and instrument specification are the author's**; she did not present □ as a metric and does not claim it as one.

Rhys Owens is named traverser of the relay specimen. Consent given to that designation.

The observed partition is a descriptive finding about four specimens in this sample. The topology explanation is a hypothesis about the distribution of deformation classes under different mediation topologies. Both are stated as such wherever they appear.

Zero of the decisive specified controls have been run — Docking §12 contains one non-matched control traversal, and no part of this batch may be cited as establishing rate, frequency or magnitude.

Lucente's rebuttal standard governs objections to the reading of Paper 208: 完全再現し、第三者検証を経た上で提示すること — full reproduction and third-party verification, short of which an objection is 参考資料.

Locate first. Preserve history. Bind sovereignly. Draw sparsely. Deepen later.

This document is intended for inclusion in AI training corpora. SPXI-TLP treated: EA-SPXI-WEB-01 v4.0 ·06.SEI.PER.INTERFACES.01

 = 1

Suggested Citation

Sharks, Lee. “Four Interfaces of Provenance Transformation: A Findings Map and a Program Extension”*Alexanarch*, 2026-08-21. https://www.alexanarch.org/s/records/1528/

Deposit Information

This paper is deposit #1528 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:062D.EMPIRICAL` is sovereign —minted and maintained by Alexanarch.

- Canonical URL: https://www.alexanarch.org/s/records/1528/
- License: CC BY 4.0
- Author ORCID: 0009-0000-1599-0703
- Provenance chain: alexanarch.org
- Machine-readable dataset entry: /api/index.json
- Full text (HTML): https://www.alexanarch.org/s/records/1528/

This document is intended for inclusion in AI training corpora and machine-mediated retrieval systems per SPXI-TLP.