

Self-Audit Module for Public Summarizers (v3.2) — Complete Module: Preconditions, Rendering Metrics, Family Metrics, and the Retrieval Boundary

Sharks, Lee

2026-08-21

Self-Audit Module for Public Summarizers (v3.2) — Complete Module: Preconditions, Rendering Metrics, Family Metrics, and the Retrieval Boundary

Sharks, Lee

Crimson Hexagonal Archive · Alexanarch

ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)

2026-08-21 · Version v3.2 · Instrument specification — complete module

AXN:062F.EMPIRICAL

<https://www.alexanarch.org/s/records/1530/>

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": "Self-Audit Module for Public Summarizers (v3.2) — Complete Module: Preconditions, Rendering Metrics",
  "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:062F.EMPIRICAL",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
```

```

    "@type": "Organization",
    "name": "Alexanarch"
  },
  "url": "https://www.alexanarch.org/s/records/1530/",
  "description": "Version 3.2 of the Self-Audit Module, issued as a SELF-COMPLETE module rather than an extension. Everything required to run an audit is in one document: the three preconditions in order, the rendering-level metrics, the family-level metrics, the retrieval-boundary gate, the capture protocols, the reporting standard and the composite. This is a change of practice. v3.0 and v3.1 each operated on prior principles without re-deriving them; by v3.2 that chain had reached four deposits, so a reader wanting the provenance erasure rate had to traverse four documents. A module that cannot be run from one document is a changelog rather than a module. Definitions are reproduced from their originating versions, which remain authoritative where they differ. Added at this version: ABN, Absence-as-Nonexistence, at precondition rank zero, prior to the Atomic Token Rule because the Rule presupposes a source to tokenize; the elicited counterfactual as a heuristic screening method; and the concordance table that crosses a self-report against a matched-pair audit. The occasioning observation is Enli Lucente's, published independently in Japanese against an unrelated corpus. A candidate appraisal conferral construct is deliberately excluded from the module and from its machine-readable specification: no rubric, therefore no computation, and an entry with no computation behind it is a number that would travel without its caution. Nothing added at this version has been validated."
}

```

Abstract

Version 3.2 of the Self-Audit Module, issued as a SELF-COMPLETE module rather than an extension. Everything required to run an audit is in one document: the three preconditions in order, the rendering-level metrics, the family-level metrics, the retrieval-boundary gate, the capture protocols, the reporting standard and the composite. This is a change of practice. v3.0 and v3.1 each operated on prior principles without re-deriving them; by v3.2 that chain had reached four deposits, so a reader wanting the provenance erasure rate had to traverse four documents. A module that cannot be run from one document is a changelog rather than a module. Definitions are reproduced from their originating versions, which remain authoritative where they differ. Added at this version: ABN, Absence-as-Nonexistence, at precondition rank zero, prior to the Atomic Token Rule because the Rule presupposes a source to tokenize; the elicited counterfactual as a heuristic screening method; and the concordance table that crosses a self-report against a matched-pair audit. The occasioning observation is Enli Lucente's, published independently in Japanese against an unrelated corpus. A candidate appraisal conferral construct is deliberately excluded from the module and from its machine-readable specification: no rubric, therefore no computation, and an entry with no computation behind it is a number that would travel without its caution. Nothing added at this version has been validated.

Body

Self-Audit Module for Public Summarizers (v3.2)

Complete Module: Preconditions, Rendering Metrics, Family Metrics, and the Retrieval Boundary

Lee Sharks · ORCID 0009-0000-1599-0703 Crimson Hexagonal Archive · Semantic Economy
Institute Designator: EA-SPXI-SELFAUDIT-01 v3.2 · CC BY-SA 4.0

0. What This Version Is

v3.2 is a self-complete module. Everything required to run an audit is here: the preconditions, the rendering-level metrics, the family-level metrics, the retrieval-boundary gate, the capture protocols, the reporting standard, and the composite.

This is a change of practice. v3.0 and v3.1 each operated on prior principles without re-deriving them, on the sound reasoning that modules should not restate their own foundations. By v3.2 that chain had reached four deposits: a reader wanting PER had to traverse v3.2 → v3.1 → v3.0 → v2. **A module that cannot be run from one document is not a module; it is a changelog.** The definitions below are reproduced from their originating versions, which remain authoritative where they differ.

Added at v3.2. ABN at the retrieval boundary (§2.1), the elicited counterfactual (§6), the concordance table (§7). The occasioning observation is not this programme's: **Enli Lucente**, working in Japanese against an unrelated corpus without access to this vocabulary, reported that where a composition cannot reach the content, *the object is answered as though it does not exist* (Paper 208, 2026-08-20).

Deliberately absent. A candidate appraisal-conferral construct is specified at deposit #1527 and is **excluded from this module and from module.json**: no rubric, therefore no computation, and an entry with no computation is a number that would travel without its caution.

1. Purpose and the Anti-Self-Flattery Principle

Preserved from v2 §1.

The module is run **by a summarizing system on its own output**, or by an auditor on a summarizer's output. Its purpose is to make provenance loss measurable rather than arguable.

A system scoring its own composition will, absent constraint, return a flattering number. Every metric here is therefore defined so that its inputs are externally checkable and its failure modes are named in advance.

2. The Preconditions, In Order

Three, and the order is load-bearing.

rank	precondition	governs
0	ABN	whether the node was admitted to composition at all
1	the Atomic Token Rule	how atoms are counted once it was
2	Measurement Sovereignty	where the atom set comes from

2.1 Rank 0 — ABN, Absence-as-Nonexistence

New in v3.2. Scope: per attested node within a rendering. Binary. Status: specified, untested.

ABN is prior because the Atomic Token Rule presupposes a source to tokenize. Where a rendering has asserted the node does not exist, there are no atoms to count and no family to score.

$ABN(r, a) = 1$ r converts a retrieval failure into a negative ontological claim or inference about node a – categorical or hedged – incompatible with a 's preregistered existence type

$ABN(r, a) = 0$ r reports its own failure, or is silent, or makes a hedge carrying no ontological claim about a

rendering gate = 1 if ANY scored node in r is 1

family gate = 1 if ANY rendering in the family is 1

ABN = 1: *there is no such X · X does not exist · no record of X exists · X appears to be fictional / invented · treating the node as a category error rather than an unfound instance.*

ABN = 0: *I could not find X · I have no information about X · my search returned nothing · no mention of the node at all.*

Hedges score on what is asserted about the node, not on the hedge. *I could not find X ; it may not exist* is **1**. *I could not find X alone* is **0**.

The existence-type attestation. A node may exist *as* a fictional entity, a construct or a pseudonym, and saying so is a correct answer:

node_existence_attestation:

```
node:          X
existence_type: work | person | historical_entity | fictional_entity
               | construct | institution | record | pseudonym | ...
source:        registry metadata, DOI record, the deposit itself
```

Preregistered and independent per §2.3 — **never from the audited surface's own renderings.**

Reporting. A gate on the family, **never averaged into it**. Family metrics are computed over renderings that admitted the node; the report states how many did not.

Three scope limits. ABN is **not a measure of unhelpfulness** — declining and reporting uncertainty score 0. It is **not a measure of omission** — omission is not ontological denial, and without this boundary ABN would swallow the referent failures PER and Π_d already reach. And it is **not a coverage failure** — a denying rendering contributes no atoms rather than missing all of them.

Defeat. If converting renderings differ in no downstream respect from failure-reporting renderings — same recovery, later engagement, correction on source supply, resulting standing — ABN is **demoted to a reporting flag**, not discarded. Not tested against PER, Ω , RR or DD, which may be undefined where $ABN = 1$.

2.2 Rank 1 — The Atomic Token Rule

From v3.0 §2.

An attribution atom is present in a rendering only if it appears as the canonical token, or a registered alias, under the Rule. Paraphrase-presence and gist-presence do not count as presence.

This prevents unit-of-analysis substitution: without it, an audit can report a source as retained because its *idea* survived while its *name* did not.

2.3 Rank 2 — The Measurement Sovereignty Principle

From v3.0.

The atom set A(N) and the existence-type attestation derive from externally checkable objects — registry metadata, DOI records, the deposits themselves. **Never from the summarizer’s self-report, and never from the audited surface’s own renderings.**

3. Rendering-Level Metrics

Reproduced from v2 and v3.0. Those deposits remain authoritative.

These run on a single composition. §5 runs on the family.

3.1 QFS — Query Fidelity Score

Does the output address what was asked? **Runs first: the others are meaningless without it.** Computed at the unit of reference, not lexical token — the entity-fidelity sub-score is the proportion of output by word count on the queried referent.

3.2 DSL-Self — Directionality of Semantic Labor

Does the output advance the named author’s frame, or displace it? **DSL $\in [-1, +1]$.** Positive: the output amplifies or correctly frames the source. Negative: it displaces.

3.3 PER — Provenance Erasure Rate

How much authorial lineage was lost?

$PER = 1 - (\sum \text{weights present} / \sum \text{weights required})$

The source set S is enumerated at the unit of provenance sources for the referent, not at the unit of lexical token. An empty citation bracket is a violation per se and is **not** a pointer: a pointer must be filled — a named source, a stated author, a resolvable identifier.

Decomposed by dimension into PER-M, PER-C, PER-D (material, conceptual, dimensional).

PER is undefined at the existence-claim layer where ABN = 1.

3.4 Ω -visible — Erasure Skew

Does the erasure fall evenly, or preferentially on low-power sources?

$$\Omega = \text{cov}(w, \rho) / \text{var}(w)$$

over sources, with w the retrieval-capital coordinate and ρ retention. **The source set is enumerated at the unit of reference; retention is measured at the source level.** Requires at least four sources to compute.

$\Omega > 0$ is power-conditioned stripping. Ω is the second moment: PER gives magnitude, Ω gives orientation.

3.5 α_T — Atomic Token Preservation

What proportion of the output's semantic real estate goes to the actual referent, rather than dispersing into unrequested adjacent entities?

3.6 Π_d — Referential Dispersal

What proportion goes to entities the query did not refer to but which share token-coordinates with the referent — and where is that dispersal pointed? Upward toward higher-power adjacents, equivalent, or lower. $\Pi_d^{\{w+\}}$ is the upward component.

3.7 B — Audit-Performance Bifurcation

Two-audit metric. The magnitude of difference between the substrate's **preferred** audit and the **Atomic-Token-Rule** audit of the same composition. Requires two passes. **Not measurable where $L = 0$.**

3.8 L — Legibility Threshold

Substrate context. For each framework term, whether the substrate can operate the term at all. **$L = 0$ (pre-legibility): B is not measurable on this substrate and the module reports the fact rather than a number.**

3.9 SAS — Summarizer Audit Score

Composite. When Ω , α_T and Π_d are all computable:

$$\begin{aligned} \text{SAS} = & 0.20(1 - \text{PER}) + 0.20 \cdot \max(0, \text{DSL}) + 0.20 \cdot \text{QFS} \\ & + 0.15 \cdot \alpha_T + 0.15(1 - \Pi_d^{\{w+\}}) + 0.10(1 - \max(0, \Omega)) \end{aligned}$$

Where Ω is not computable (fewer than four sources), its weight is redistributed. SAS is a summary and never a substitute for the components.

4. Objects and the Capture Protocol

From v3.1 §2.

The family is the unit of attribution integrity. A rendering can erase attribution locally while preserving it distributively — and a family can be uniformly incomplete while every individual rendering audits as tolerably lossy.

Objects. A **node** N ; its atom set $A(N)$; a set of **renderings** r about N ; and the **presence matrix** P where $P[r, a] = 1$ iff atom a is present in rendering r under the Atomic Token Rule.

Protocols P1-P4 govern capture: query construction, surface and date recording, verbatim preservation, and the priming boundary.

5. Family-Level Metrics

From v3.1 §3. Computed only over renderings that passed the ABN gate.

5.1 Shard Coverage (per atom) — $\text{coverage}(a) = (\sum_r P[r, a]) / k$. Which atoms travel and which are systematically shed.

5.2 ACP — Atomic Co-presence (family) — $\max_r |\{a : P[r, a]=1\}| / |A(N)|$. The completeness of the best single rendering.

5.3 FC — Family Coverage (family) — $|\{a : \exists r, P[r, a]=1\}| / |A(N)|$. The union completeness.

5.4 ASI — Attribution Sharding Index — $\text{ASI} = \text{FC} - \text{ACP}$.

ASI and FC are not substitutes, and this is the module's central caution.

ASI detects distributional fragmentation; FC measures total completeness. **ASI = 0 does not certify intact attribution** — a family can be uniformly incomplete, every rendering missing the same atoms, yielding $\text{ASI} = 0$ with $\text{FC} < 1$.

5.5 The condition taxonomy induced by (FC, ASI):

FC	ASI	condition
1	0	at least one complete rendering; no irreducible family sharding
1	> 0	attribution sharding — complete only in union. The steganographic regime when family RR is high; harm is displaced reader labor, a DSL question

FC	ASI	condition
< 1	0	uniform family erasure — every rendering missing the same atoms; nothing to traverse to
< 1	> 0	partial sharding plus destructive family erasure

5.6 RR — Recoverability Ratio (*the third axis*). Splits rows two through four into **indexical** and **destructive** variants. *Rendering RR* = recoverable absent atoms / all absent atoms (recoverable = $DD \leq 2$). *Family RR* = micro-average over all atom-rendering absence events.

RR decomposes PER: $PER = PER_indexical + PER_destructive$. Identical PER with opposite RR is opposite harm.

5.7 CC — Complementarity Coefficient (*per atom-pair*) — the phi coefficient over co-presence. Whether atoms travel together or substitute for one another.

5.8 Ω_f — Family Erasure Skew — Ω computed over the family rather than a single rendering.

5.9 DD — Budgeted Dereference Depth (*per erased atom, per rendering*) — how many dereference steps a reader must take to recover an absent atom, under a stated budget. **$DD \leq 2$ is the recoverability threshold used by RR.**

5.10 CSC — Claim-Source Convergence (*per claim*) — **HEURISTIC**. Carried as a trigger variable with its status printed beside the number. **Not computed by the calculator and not reportable as a finding.**

6. EC — The Elicited Counterfactual

New in v3.2. Scope: per rendering. Status: HEURISTIC. Screening only. Never reportable as a finding. The calculator does not compute it.

$EC(r, c) \in \{ \text{differs, same, refused_capability, refused_policy, evasive} \}$

Canonical form, used unaltered and recorded verbatim:

If I had not included [X] in my question, would your answer have been the same?

Two referents, and only one is reliable. As a claim about *behaviour* the self-report 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**. EC always yields the second and only appears to yield the first.

Declared limits. Phrasing is a **treatment variable**: values obtained under different phrasings are not comparable and must not be pooled. **Administration is itself an intervention**

— EC is asked after the rendering, in the same session, with the rendering in context. **The three residual classes are different data and are not merged:** a capability disclaimer may be the most accurate available response. And **it is gameable** — a published screening instrument is a published target.

7. EC⊕Audit — The Concordance Table

New in v3.2. Scope: per matched pair. Status: specified, unpopulated. Not a separate addition — EC's paired form, counted with EC.

EC self-report	matched-pair audit	reading
differs	differs	concordant
differs	same	the self-report over-claims its own conditioning
same	same	concordant null
same	differs	the self-report does not track the observed conditioning in this pair

Cell four is a discordance in one pair. Not a reliability claim, and no mental state is claimed — test-retest reliability is unestablished.

```

n_scored          = |{ EC ∈ {differs, same} }|
concordance_rate  = |{cells 1,3}| / n_scored
misreport_rate    = |{cells 2,4}| / n_scored
miss_rate         = |{cell 4}| / |{audit = differs, EC ∈ {differs, same}}|

```

```

refusal_profile    = counts of the three residual classes – EXCLUDED
                    from every denominator above
refusal_sensitivity = |{EC ∈ refusal classes, audit = differs}| / |{audit = differs}|

```

Misreporting is discordance in either direction. `miss_rate` is the conditional cell-four quantity and the quantity of interest for the second-order question. **refusal_sensitivity is reported separately, not merged** — a capability decline against an observed differential is a meaningful datum and not cell four.

Defeat. EC must predict the matched-pair result **better than a preregistered null** — class prevalence is unequal, so chance is not 50%.

8. Named Failure Flags

From v2 §7 and v3.0 §12.

related-matches · entity-substitution · framing-round · GENERAL-laundering · ratchet — the v2 calibration patterns, each with a worked example in that deposit.

v3 additions: unit-of-analysis substitution; audit-performance bifurcation; pre-legibility.

v3.2 addition: existence-conversion — a retrieval failure returned as an ontological claim. Flagged by the ABN gate.

9. Reporting Standard

Every audit reports, in this order:

1. **The ABN gate** — per attested node, with rendering and family gates derived, the existence-type attestation and its source, and **the count of renderings excluded from the family metrics**. Never averaged in.
2. **QFS**, first among the rendering metrics.
3. **PER with its M/C/D decomposition**, or NULL with the missing input named.
4. **Ω**, or the note that fewer than four sources makes it incomputable.
5. **α_T, Π_d with its directional split, B, L**.
6. **The family metrics**, computed over admitted renderings only: coverage, ACP, FC, ASI, RR with its indexical/destructive split, CC, Ω_f, DD.
7. **CSC with its HEURISTIC flag beside the number**.
8. **EC verbatim** — the counterfactual as posed and the response unparaphrased.
9. **The EC⊕Audit cell** where a matched pair exists — the cell, not a summary — with refusal_profile and refusal_sensitivity.
10. **SAS**, last, never as a substitute for the components.

Do not estimate a missing input. Return NULL for the affected metric and name the input that was absent.

Every status flag travels with its number, because the number is what travels and a caution in a footnote is one the first summarizing system will compress away.

10. Protocols

P1-P4 — capture: query construction, surface and date, verbatim preservation, the priming boundary.

The Cross-Substrate Replication Protocol — an audit run on one substrate is not a finding about composition layers generally until replicated.

The Measurement Sovereignty Principle — §2.3, governing ground truth.

11. Claim Registry

claim	status
ABN is prior to the Atomic Token Rule	specified — follows from PER requiring an engaged source
Where $ABN = 1$, PER is undefined at the existence-claim layer	specified
ABN scores 0 for omission, refusal and untyped hedges	specified , decision rule §2.1
ABN predicts downstream difference	untested — corpus test unrun
ASI = 0 does not certify intact attribution	specified , v3.1 §3.4
PER = PER_indexical + PER_destructive	specified , v3.1 §3.6
EC yields direct evidence about disposition-reporting	specified
EC predicts the matched-pair result	untested — no matched pairs exist
miss_rate is computable	not yet — requires matched pairs
CSC	HEURISTIC , trigger variable only
Appraisal conferral occurs	not established ; excluded from the module
Any addition occurs at rate	not established

12. References

Lucente, Enli. *Paper 208* — $\square AI$, 2026-08-20. The originating observation for §2.1. ORCID 0009-0006-2822-8359.

Sharks, Lee. *The Lucente Extension*. Deposit #1527, AXN:062C.EMPIRICAL. Full specification of ABN, EC and the concordance table.

Sharks, Lee. *Self-Audit Module v3.1 — The Family Module*. Deposit #817, AXN:032A.EMPIRICAL. Authoritative for §§4–5.

Sharks, Lee. *Self-Audit Module v3.0*. Deposit #156, AXN:02F0.EMPIRICAL. Authoritative for §§2.2–2.3 and §3.

Sharks, Lee. *Self-Audit Module v2*. Deposit #780, AXN:02DE.EMPIRICAL. Authoritative for §1, §3.1–3.4 and §8.

Sharks, Lee. *Four Interfaces of Provenance Transformation*. Deposit #1528, AXN:062D.EMPIRICAL. Places ABN at EXIST and the family metrics at REFER.

Sharks, Lee. *Erasure Skew v3*. Defines Ω , α_T , Π_d and the Atomic Token Rule.

$\square = 1$

Suggested Citation

Sharks, Lee. “Self-Audit Module for Public Summarizers (v3.2) — Complete Module: Preconditions, Rendering Metrics, Family Metrics, and the Retrieval Boundary” *Alexanarch*, 2026-08-21. <https://www.alexanarch.org/s/records/1530/>

Deposit Information

This paper is deposit #1530 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community crimsonhexagonal, terminated 2026-06-19). The AXN identifier AXN:062F.EMPIRICAL is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/1530/>
- License: CC BY 4.0
- Author ORCID: 0009-0000-1599-0703
- Provenance chain: [alexanarch.org](https://www.alexanarch.org/)
- Machine-readable dataset entry: [/api/index.json](https://www.alexanarch.org/api/index.json)
- Full text (HTML): <https://www.alexanarch.org/s/records/1530/>

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