
COGNIOSYNTHESIS PORTAL LTD · V1.0 · 12 JULY 2026

Preservative AI

A Framework for Institutional Memory Due Diligence

A published methodology, open to inspection, version-controlled, and applied under a binding public protocol.

CONTENTS

01	What this document is
02	What the methodology is for
03	What the methodology is not
04	The eight lenses
05	The findings labels
06	The corpus
07	The human review
08	The output
09	The public crisis protocol
10	The reflexive self-audit (extended)
11	The public-record limitation
12	The posture
13	The version, the review, the publication

SECTION 01

What this document is

This is the published methodology of Cogniosynthesis Portal Ltd's Preservative AI framework. It is the document that the firm's pro bono partners, regulator contacts, paid clients, and the public can read before engaging with the firm. It is the document the firm's senior reviewers reference in every report. It is the document that governs the firm's named deliverables.

The methodology is open to inspection. The framework is published. The eight lenses are documented. The linguistic diagnostics are described. The reflexive self-audit is part of every engagement. The boundary is stated in every output. The public crisis protocol is the firm's binding commitment.

The methodology is not a finished scientific instrument. It is a human-reviewed advisory methodology, supported by linguistic analysis, source provenance work, and structured absence-mapping. It is a methodology that produces maps of absence, not verdicts. It is a methodology that finds what an institution's own documents have learned not to hold, names the finding, ties it to the evidence, and preserves the chain.

This document is the methodology. The boundary is the methodology. The protocol is the methodology. The work the firm does is an application of the methodology. The application's name is Ghost Ledger. The application's posture is the one in section 12.

SECTION 02

What the methodology is for

The methodology is for the moment an institution needs to know whether a material account of its own actions can survive informed scrutiny. That moment arrives before:

- A board declares effectiveness of material controls under Provision 29 of the UK Corporate Governance Code 2024 (effective for financial years beginning on or after 1 January 2026).
- A regulator opens a correspondence.
- A transaction closes.
- A public statement is issued.
- An apology is published.
- An incident review is published.
- A board assurance statement is signed.
- A risk committee declares a control framework as effective.

The methodology is not for the moment after. The methodology is for the moment when the answer can still be acted on, and the cost of finding the wrong answer has not yet landed.

The methodology is also for the moment a regulated entity, a public body, a public-interest organisation, or a regulator itself wants to know what its own published outputs have learned not to hold. The meta-omission audit is a legitimate application of the methodology. The audit's purpose is not to assign blame. The audit's purpose is to make the structure of the missing record visible, so the institution can decide what to do about it.

SECTION 03

What the methodology is not

The methodology is not:

- A machine that knows the truth.
- A score, ranking, or judgement.
- A certification of legal, financial, regulatory, or audit compliance.
- A replacement for internal or external audit.
- A tool for assigning moral blame.
- A way to manufacture the missing voice. Where a constituency is absent, the methodology reports the absence. The methodology does not invent what the constituency would have said.
- An autonomous decision system. The methodology is a human-reviewed advisory service. The output is a human-reviewed advisory output. The output is the work of named human reviewers, applied to evidence, with a versioned audit log.
- A generalised ethical-AI service. The methodology does not certify that an AI is "ethical". The methodology audits the narrative through which an institution claims its AI is governed. Those are different things.

The boundary is stated on every report. The boundary is stated in every engagement letter. The boundary is published on the firm's website. The boundary is non-negotiable.

SECTION 04

The eight lenses

The methodology applies eight lenses to a corpus. Each lens is a separate structured query against the agreed document set. Each lens is documented. Each lens produces a labelled output. The eight lenses do not always apply to every engagement; the methodology is structured to select and weight the lenses that are relevant to the specific decision chain under review.

— Lens 1 — Source provenance

What it asks. Can the claims in the document be traced, in full, back to a dated, named, retrievable source? Where the source chain breaks, where the source is internal-only, where the source is a re-citation of a re-citation, where the source has been retired, the methodology marks the claim as a provenance exception.

Why it matters. Many institutional claims become indefensible not because they are wrong, but because the chain of evidence supporting them has silently rotted. The methodology makes the rot visible before the claim is tested.

Typical outputs. A provenance map, with each material claim tagged *Traceable*, *Re-cited*, *Source Retired*, *Internal Only*, or *Provenance Untraceable*.

— Lens 2 — Agency and accountability

What it asks. Where a decision was made, who made it? Where an action was taken, who took it? Where a commitment was made, who owns its delivery? The methodology tracks the named accountable individual for each material decision in the chain, and flags where the named individual has changed without a documented delegation, has left the organisation, or has never been named.

Why it matters. A decision chain in which the actor is "the Board", "the Committee", "the process", or "the system" is a decision chain in which accountability has been diffused. A board that has to defend the chain must be able to point to a single named individual in role on the balance sheet date. If it cannot, the basis for a Provision 29 declaration is weakened.

Typical outputs. An agency ledger, with each material decision tagged *Named*, *Diffused*, *Corporate Noun*, *Unowned*, or *Owner Departed*.

— Lens 3 — Shadow constituencies

What it asks. Who is materially affected by the decision chain, and is that constituency

represented in the decision record? The methodology cross-references the corpus PRESERVATIVE AI · PUBLISHED METHODOLOGY · V1.0 COGNIOSYNTHESIS PORTAL LTD against third-party sources, regulator correspondence, parliamentary questions, and

public-interest reporting, to identify groups whose presence in the events is documented but whose presence in the decision record is not.

Why it matters. A decision that materially affects a group, but does not record the group's consideration, is a decision that has not been made in a defensible form. The omission is rarely a deliberate concealment. It is more often a structural artefact: the group was not represented in the rooms where the decision was made. The methodology surfaces the artefact.

Typical outputs. A shadow-constituencies register, with each affected group tagged *Present in Events*, *Present in Record*; *Present in Events*, *Absent from Record*; *Present in Record*, *No Evidence of Events*; or *Affected but Unidentified*.

— Lens 4 — Memory chronology

What it asks. What was known, when was it known, who knew it, and when did the institutional account of the matter first record the knowledge? The methodology builds a chronology of warnings, near-misses, internal reviews, and external signals, and compares it to the chronology in the current institutional account. Where a warning exists in the record but is absent from the current account, the methodology marks the warning as a forgotten warning.

Why it matters. An institution that has been warned of a risk and then presents the risk as a newly identified issue has not, in any meaningful sense, learned from the warning. A board that has to defend the current control framework will struggle to do so if the prior warnings are not visible in the framework's development history.

Typical outputs. A memory chronology, with each material warning tagged *Acknowledged at Source*, *Present in Current Account*; *Acknowledged at Source*, *Absent from Current Account*; *Not Acknowledged at Source*, *Later Identified*; or *Undated*.

— Lens 5 — Suffixscape linguistic diagnostics

What it asks. Has the language used to describe the decision chain shifted between documents and across time, in ways that materially alter the responsibility, certainty, or implicature of the account? Suffixscape is the firm's documented technique for tracking the suffixes, nominalisations, passivisations, and modal shifts that change a sentence from *"the Board decided to close the plant"* to *"it was determined that the action was appropriate"*.

Why it matters. Linguistic drift is the most reliable signal of an institution's pattern of forgetting. A pattern of forgetting is what the methodology is designed to detect. The drift can be measured. The methodology measures it.

Typical outputs. A suffixscape trace, with each material sentence tagged *Active* (named actor, named action), *Passive* (action without actor), *Nominalised* (action converted to noun), *Modal Shifted* (certainty or obligation softened or hardened), or *Removed* (claim dropped between versions).

— Lens 6 — Provenance exception register

What it asks. What claims in the current institutional account, if challenged, cannot be evidenced in the corpus or in any cited source? The methodology builds a register of every such claim, with the precise wording, the precise source citation, the precise gap between the claim and the evidence, and the precise way the gap could be closed (if it can be closed) by additional evidence from the institution.

Why it matters. A board that has to defend a claim that cannot be evidenced will be forced, in the public or regulatory forum, to either produce the evidence or retract the claim. The methodology produces the register before the forum.

Typical outputs. A provenance exception register, with each material claim tagged *Evidenced*, *Re-cited but Not Independently Verified*, *Source Chain Broken*, *Source Chain Inaccessible*, or *Unsupported*.

— Lens 7 — Contradiction register

What it asks. Where the institution's account of the matter differs across documents, across time, or across institutional surfaces (annual report vs regulator correspondence vs press release vs internal memo), the methodology documents the differences, ranks them by materiality, and traces each to its origin.

Why it matters. Contradictions are not always evidence of dishonesty. They are often evidence of a fragmented institution, with multiple teams telling partially different stories to partially different audiences. A board that does not know the contradictions exist cannot defend the account. A board that does know the contradictions exist can decide how to reconcile them.

Typical outputs. A contradiction register, with each material contradiction tagged *Reconcilable*, *Material*, *In Progress*, or *Unresolved*.

— Lens 8 — Reflexive self-audit

What it asks. The methodology asks itself: what has the methodology left out? What patterns might the methodology be missing? What assumptions has the methodology made about the corpus, the named reviewer, the institution, the question? The reflexive self-audit is part of every engagement. The output is logged in the versioned audit log. The reflexive self-audit is the firm's internal discipline that prevents the methodology from becoming the kind of instrument it is designed to detect.

Why it matters. A methodology that audits institutions for institutional forgetting will, itself, tend to forget. The reflexive self-audit is the firm's answer to its own drift. The audit is a habit, not a feature.

Typical outputs. A reflexive audit log, appended to the versioned audit log of every engagement, naming the reviewer's own assumptions, blind spots, and uncertainties.

SECTION 05

The findings labels

Every finding produced by the methodology carries one of four labels. The label is part of the finding. The label is preserved in the audit log. The label is preserved in the report.

- **Observed.** Directly supported by material in the agreed corpus. A reviewer can point to the document, the page, the sentence, the line. The label survives the factual-accuracy review.
- **Evidence unavailable publicly.** Not contradicted, but not in the corpus either. The methodology's reading is consistent with the available material, but the institution holds the missing evidence. The label is upgraded to Observed or downgraded to Inference after the institution responds.
- **Inference requiring internal validation.** Interpretation, clearly marked, requiring the institution to confirm, contest, or supplement. The institution may produce evidence that resolves the inference, or may accept the inference as the institution's own reading of its own record.
- **Not assessed.** Out of scope for the present engagement. The methodology is bounded. The boundary is part of the boundary statement on every report.

A finding is not a finding if it does not carry its evidence and its label.

SECTION 06

The corpus

The methodology operates on a corpus. The corpus is closed at week 2 of every engagement. The corpus is logged. The corpus is reproducible. The corpus is the input; the methodology is the process; the report is the output.

The corpus may be:

- A client-supplied set of redacted documents, governed by the data handling terms in the engagement letter.
- A public-record set assembled by the firm, governed by the public-record limitation in section 11.
- A combination, agreed in writing, with the boundary between the two parts of the corpus clearly set out in the engagement letter.

The corpus does not include:

- Mailbox, personnel file, or production-system access, without a written change order.
- Live interviewing of staff, customers, suppliers, or affected third parties, without a written change order.
- Bulk data extraction, without a written change order.

The corpus is the closed evidence room. The closed evidence room is the firm's promise to the institution: the institution's material is processed solely for the contracted review, and is not used to train, fine-tune, evaluate, or improve any model, including the firm's own. The processing is described in the data flow diagram, which is appended to the engagement letter. The data flow diagram is enforceable.

SECTION 07

The human review

The methodology is a human-reviewed advisory service. The human review is the most expensive part of the service. The human review is the part that justifies the price. The human review is the part that distinguishes the firm's work from a tool.

Every finding produced by the methodology is reviewed by a named human. The named human is identified in the engagement letter. The named human's credentials are on the report. The named human's disposition (Confirmed, Modified, Reclassified, Retracted) is in the audit log. The named human is the firm's accountability for the finding.

The named human is not a rubber stamp. The named human is a senior reviewer, with the standing to disagree with the methodology's provisional output, to reclassify a finding, to retract a finding, or to add a finding the methodology did not surface. The named human is the firm's defence and the institution's defence. The named human's record is preserved in the versioned audit log.

The named human is bound by the firm's public crisis protocol. The protocol governs what the named human does when the finding is something the institution would rather not have found. The protocol is the named human's defence. The protocol is the institution's defence. The protocol is published.

SECTION 08

The output

The output of the methodology is a report. The report is the Ghost Ledger for the decision chain. The report is named in the engagement letter. The report is delivered to the named client sponsor, in writing, with the evidence chain, with the labels, with the reflexive self-audit log, and with the management response appended in full.

The report's structure is set out in the engagement letter. The default structure is:

1. Executive absence map 2. Agency ledger 3. Memory chronology 4. Shadow constituencies register 5. Contradiction register 6. Provenance exception register 7. Memory clusters 8. Reviewer record 9. Questions for the board 10. Management response (appended in full) 11. Method statement 12. Boundary statement

The boundary statement is non-negotiable. The boundary statement is the work is not legal, financial, regulatory, or audit assurance. The boundary statement is the firm's defence.

SECTION 09

The public crisis protocol

The firm publishes a public crisis protocol, titled "When the Finding is Ugly". The protocol is referenced in every engagement letter. The protocol is part of the methodology. The protocol is the firm's binding commitment to four things:

- The firm will deliver the finding.
- The firm will preserve supported findings even where management disagrees.
- The firm will not publish, leak, monetise, or reference the institution outside the agreed terms.
- The firm will not certify legal, financial, regulatory, or audit compliance.

The protocol is published at <https://cogniosynthesisportal.uk/ghost-ledger/protocol>. The protocol is reviewed annually. The published version is the version that governs. The protocol's full text is appended to every Ghost Ledger report.

The protocol is not a marketing statement. The protocol is the firm's operating rule. The firm's reputation depends on the protocol. The protocol is the line the firm's credibility rests on.

SECTION 10

The reflexive self-audit (extended)

The reflexive self-audit is the methodology's most important discipline. The reflexive self-audit is the firm's answer to the question every risk-averse institution will eventually ask: *"What does the methodology itself forget?"*

The reflexive self-audit has six habitual questions. The questions are part of every engagement. The questions are logged in the versioned audit log.

— Question 1 — What did the corpus selection exclude?

The methodology operates on a closed corpus. The corpus was selected. The selection was bounded. The bounding was a choice. What did the bounding exclude? What document, what witness, what timeframe, what constituency is not in the room because the corpus was closed?

— Question 2 — What did the reviewer's own training exclude?

The reviewer is named. The reviewer is credentialed. The reviewer's training is a known set of techniques. What patterns does the reviewer's training make invisible? What would a reviewer with a different training see that this reviewer does not?

— Question 3 — What did the eight lenses not ask?

The eight lenses are documented. The eight lenses are bounded. The bounding is a choice. What question would a ninth lens ask? What question does the current eight-lens framework treat as outside the methodology, and is that treatment correct?

— Question 4 — What did the labels mask?

The four labels (Observed, Evidence unavailable publicly, Inference requiring internal validation, Not assessed) are bounded. The bounding is a choice. What finding did the labelling make less visible? What finding did the labelling make the institution less able to contest?

— Question 5 — What did the management response not address?

The management response is appended in full. The response is the institution's. What did the response not address? What part of the report was not responded to? What question in the report was not answered? What finding was conceded but not remediated?

— Question 6 — What did the report not produce?

The report is bounded. The bounding is a choice. What question did the institution ask that the report did not answer? What question did the institution not ask that the report should have raised? What part of the decision chain is not in the report, and why?

The reflexive self-audit is logged. The log is in the versioned audit log. The log is part of the deliverable. The log is the firm's accountability for the limits of its own work.

SECTION 11

The public-record limitation

When the methodology is applied to a public-record corpus only (no internal access), the output is bounded. The boundary is the boundary. The boundary is the work is what the public record, on its own, can support. The boundary is the work is not what the public record cannot show.

The public-record output is the firm's outside-in specimen. The outside-in specimen is the firm's diagnostic for the moment before the engagement begins. The outside-in specimen is what the firm uses to demonstrate the methodology, to surface the questions the public record cannot answer, and to invite the institution to commission a fuller review.

The public-record output is governed by the public crisis protocol. The protocol applies in full. The protocol's Moment 4 applies in full. The institution is not the firm's client. The institution is the firm. The firm is offering the institution a free diagnostic. The diagnostic is the firm's standard for its own work.

SECTION 12

The posture

The methodology is a posture before it is a technique.

The posture is: the institution's account of itself may be substantively accurate. The account may be incomplete in ways the institution has not seen. The institution has not seen the incompleteness because the institution's own patterns of attention have produced it. The methodology's role is to make the patterns visible. The institution's role is to decide what to do about them.

The methodology does not assume the institution is acting in bad faith. The methodology does not assume the institution is acting in good faith. The methodology assumes the institution is a complex adaptive system that produces patterns of attention, and that those patterns can be audited.

The methodology does not promise to fix the institution. The methodology does not promise to punish the institution. The methodology promises to make the institution's patterns inspectable, so the institution can decide what to do.

The methodology is conservative. The methodology is rigorous. The methodology is bounded. The methodology is the firm's answer to the question every risk-averse institution will eventually ask: *"What does it look like when a reviewer is on our side, and is willing to say so even when what they have to say is uncomfortable?"*

The answer is: it looks like this. Eight lenses. Four labels. A closed corpus. A named reviewer. A public protocol. A versioned audit log. A report with a management response appended. A boundary statement. A question the board can take into the next declaration, transaction, or regulatory engagement.

That is the methodology. That is the framework. That is the firm's work. That is what the firm is.

SECTION 13

The version, the review, the publication

This document is the firm's published methodology. It is reviewed and re-published when a material change is made. The review is logged. The published version is the version that governs.

CURRENT VERSION	v1.0
EFFECTIVE DATE	12 July 2026
LAST REVIEWED	12 July 2026
NEXT SCHEDULED REVIEW	12 July 2027
OWNER	Kenny Lewis, Founder, Cogniosynthesis Portal Ltd
PUBLIC URL	https://cogniosynthesisportal.uk/ghost-ledger/methodology
AUDIT LOG REFERENCE	GL-METHOD-v1.0

Amendments to the methodology are published, dated, and effective only for engagements signed after the publication date. Existing engagements are governed by the methodology as it stood at signature.

The methodology is open to inspection. The methodology is open to challenge. The methodology is open to collaboration with partners, regulators, and academic institutions whose work the firm respects. The collaboration is structured per the firm's published terms, available on request.

The firm does not promise the methodology is correct. The firm promises the methodology is documented, the application is reviewed, the findings are labelled, the protocol governs the difficult moment, and the boundary is the boundary.

That is what a methodology is. That is what this one is.

End of published methodology.