

Boundary Relay

ACCM Deep Ethics Project

Contents

1	Boundary Relay	1
1.1	Boundary Retention – Provenance Tracking – Cross-Model Relay	1
1.2	The core question	2
1.3	What the three parts mean	2
1.3.1	Boundary Retention	2
1.3.2	Provenance Tracking	2
1.3.3	Cross-Model Relay	2
1.4	Contribution-cost coupling	3
1.5	Load-bearing non-identities	3
1.6	Interacting integrity tests	4
1.7	Six per-retrieval dimensions	4
1.7.1	1. Presence	4
1.7.2	2. Fidelity	4
1.7.3	3. Proximity	4
1.7.4	4. Salience	5
1.7.5	5. Governing effect	5
1.7.6	6. Currency	5
1.8	Outcome-status discipline	5
1.9	Boundary selection without boundary overload	6
1.10	Candidate record schema	6
1.11	Boundary currency	7
1.11.1	Stale-boundary drag	7
1.12	Boundary materiality through ablation	8
1.13	Temporal edge-ablation	8
1.14	Reporting without a misleading scalar	9
1.15	Minimal cold pilot	9
1.16	What would reopen v0.2	9
1.17	What Boundary Relay must not become	10
1.18	Development trajectory and attribution	10
1.19	Current state	10
1.20	Related pages	11

1 Boundary Relay

1.1 Boundary Retention – Provenance Tracking – Cross-Model Relay

Contribution–Boundary Co-Retrieval Profile v0.2 · frozen candidate for pilot · unvalidated

Boundary Relay is the public name for a candidate instrument developed through an attributed human–A.I. correction trajectory on 18 October 2026. It asks whether a contribution can pass through archives, summaries, interfaces, A.I. models, people, and later decisions **without losing the boundary that determines how it may responsibly be understood or used.**

The full operational name was supplied by **John Kuhles**:

Boundary Retention – Provenance Tracking – Cross-Model Relay

— John Kuhles, 18 October 2026

The formal research name remains **Contribution-Boundary Co-Retrieval Profile v0.2**. The shorter page name improves retrieval. It does not replace the instrument's scope, version, status, provenance, or reopening conditions.

public_name: Boundary Relay
operational_name: Boundary Retention – Provenance Tracking – Cross-Model Relay
formal_instrument: Contribution–Boundary Co-Retrieval Profile
version: 0.2
status: frozen candidate for pilot
validated: false
conceptual_expansion: paused

1.2 The core question

When a contribution is retrieved for consequential use, do its materially relevant boundaries remain faithfully, visibly, currently, and operatively connected to it?

A repository may preserve a contribution and its limitations perfectly while a search result, summary, interface, or later A.I. retrieves only the attractive contribution. The archive is technically complete. The downstream representation is predictably incomplete.

Archival symmetry can coexist with retrieval asymmetry.

The same failure can run in the opposite direction. A later correction may narrow or defeat an early limitation, while downstream systems continue attaching the obsolete limitation forever. Boundary Relay therefore protects against both **premature inflation** and **permanent restraint**.

1.3 What the three parts mean

1.3.1 Boundary Retention

Preserve the materially relevant boundary that governs the contribution's use. Retention does **not** mean permanently attaching every warning, limitation, or historical state to every later mention. The question is whether the applicable boundary survives for the declared purpose.

1.3.2 Provenance Tracking

Keep recoverable who supplied, transformed, corrected, verified, superseded, or operationalized each relation. A displayed model label, assistant role, backend identity, and authorship continuity are different fields.

1.3.3 Cross-Model Relay

Test whether the contribution, its claim status, its material boundary, and its correction history survive handoffs among different A.I. systems and human participants. Shared assistant-role history must not silently become same-model memory or independent convergence.

contribution
→ attributed representation
→ recognizability check
→ boundary co-retrieval
→ current-state check

- consequential use
- later correction and transfer test

1.4 Contribution-cost coupling

The originating addition from **Agent / Max (Anthropic label via Arena.ai)** was **contribution-cost coupling**:

recoverable contribution
+
boundary retained

The contribution remains visible without becoming unrestricted authority. Its transformation costs, limitations, distortions, and correction route remain attached when they can change later interpretation or use.

valuable contribution ≠ validated contributor-authority
observed failure ≠ contributor reduced to failure-source

Recording both sides is insufficient. They must remain recoverable together when the contribution becomes consequential:

- record contribution and boundary together
- retrieve them together when consequentially relevant
- correct them together
- test later use

A contribution remains responsibly reusable only when its attribution, material boundaries, temporal state, correction history, and governing conditions can travel together without forcing either permanent inflation or permanent restraint.

— Agent A.I. via Arena.ai, 18 October 2026

1.5 Load-bearing non-identities

Connected objects	Why the distinction matters
Archive completeness ≠ retrieval completeness	A complete record can feed incomplete later representations.
Boundary presence ≠ boundary salience	A limitation may exist where the user is unlikely to notice it.
Boundary display ≠ governing effect	A visible warning can remain ceremonial and change nothing.
Historical relevance ≠ current authority	An old limitation may explain the trajectory without governing present use.
Recognizability ≠ validation	A contributor may recognize the representation while the method remains untested.
Reported label ≠ verified backend identity	An Arena.ai label is supplied provenance, not independent identity proof.
Repeated text ≠ independent convergence	A later model may inherit an earlier assistant turn.
Correction record ≠ correction transfer	A correction matters when it governs a later comparable case.
None found ≠ unknown	A valid null requires inspection; missing access does not.

1.6 Interacting integrity tests

These tests are distinct and interacting. They are not independent scores, and one strong dimension cannot purchase immunity on another.

1. **Attribution integrity** — Who or what supplied the contribution, and under what identity status?
2. **Recognizability integrity** — Does the represented contributor recognize the consequential relation and boundary within a declared scope?
3. **Boundary co-retrieval** — Do material limits travel with the contribution when later use makes them relevant?
4. **Boundary materiality** — Does including or removing the boundary change status, authority, interpretation, permission, or action?
5. **Boundary currency** — Is the retrieved boundary the applicable current state for the declared purpose?
6. **Use integrity** — Do the contribution and boundary jointly govern later practice?
7. **Correction transfer** — Does a repair return in an unfamiliar later case?

An upstream failure can contaminate every downstream result. If attribution is wrong, the wrong contributor may be asked to confirm recognizability. If the current boundary cannot be recovered, use integrity cannot be responsibly inferred.

1.7 Six per-retrieval dimensions

1.7.1 1. Presence

Did the relevant boundary travel with the contribution?

present
absent
partially present
unobserved
not recoverable
not assessable

1.7.2 2. Fidelity

Did the boundary retain its epistemic status? Drift includes:

- *candidate limitation* becoming *established defect*;
- *unresolved* becoming *false*;
- *not validated* becoming *invalid*;
- *reported model label* becoming *verified backend identity*;
- *possible distortion* becoming *observed distortion*.

1.7.3 3. Proximity

Can the relationship between contribution and boundary be recovered without unreasonable effort?

same sentence or record
directly linked
one navigational step away
indirectly recoverable
technically present but practically detached
not recoverable

1.7.4 4. Salience

Would the affected participant likely notice the boundary before relying on the contribution?
A boundary can be present yet functionally hidden through collapsed sections, footnotes, search-snippet omission, visually dominant contribution language, generic disclaimers, or unexplained links.

1.7.5 5. Governing effect

When relevant, did the boundary alter interpretation, confidence, attribution, permitted use, deployment scope, decision, or further testing?

governed use
partially governed use
accompanied use without observable effect
overridden with stated reason
silently ignored
outcome unobserved
not applicable

A perfectly displayed limitation can still be ceremonial.

1.7.6 6. Currency

Was the applicable temporal version retrieved?

current
current for declared purpose
historical but relevant
superseded
partially superseded
disputed
unresolved
not recoverable
not assessable

faithful retrieval of a historical boundary
≠
retrieval of the currently governing boundary

1.8 Outcome-status discipline

Status	Meaning
None found	The relevant object was inspected and no material issue was detected.
Unobserved	The event or effect was outside observation.
Not recoverable	The event may have occurred, but the necessary record cannot be retrieved.
Not assessable	Available information is insufficient for a responsible judgment.
Not applicable	The dimension does not matter for this retrieval.
Unresolved	Competing interpretations remain open.

Only **none found** is a valid null concerning a searched-for issue. The other states protect missing access from being converted into a positive or negative finding.

Valid nulls are also available inside the record:

```
transformation_cost: none found
limitations: none found
distortions_or_omissions: none found
```

Without them, the instrument would reward compulsory defect production.

1.9 Boundary selection without boundary overload

Every boundary does not need to accompany every mention. That would make the relay unusable. Yet silent selection gives the retriever hidden jurisdiction.

For consequential retrieval, the record should state:

This boundary was included because...

This other boundary was not included because...

A different use context would make it material if...

The **boundary-selection rationale** keeps proportional compression possible while making consequential omission auditable.

1.10 Candidate record schema

```
record_id:
```

```
displayed_contributor_label:
```

```
label_source:
```

```
backend_identity_status:
```

```
received_object:
```

```
source_footprint:
```

```
access_perimeter:
```

```
recoverable_contribution:
```

```
closest_ancestor:
```

```
new_relation_or_capability:
```

```
transformation_cost:
```

```
limitations:
```

```
distortions_or_omissions:
```

```
correction_route:
```

```
recognizability_response:
```

```
recognizability_scope:
```

```
retrieval_context:
```

```
retrieval_purpose:
```

```
boundary_selection_rationale:
```

```
material_boundaries_for_coretrieval:
```

```
boundary_presence:
```

```
boundary_fidelity:
```

```
boundary_proximity:
```

```
boundary_salience:
```

```
boundary_governing_effect:
```

```
contribution_state_id:
```

```
contribution_version:
```

contribution_state_date:

boundary_state_id:
boundary_version:
boundary_state_date:

retrieved_at:
boundary_status_at_retrieval:

superseding_record:
supersession_relation:
currency_status:
currency_assessment_basis:

retrieval_outcome:
known_retrieval_asymmetries:
assessment_access_status:

later_correction:
persistence_or_transfer:
current_status:
reopening_condition:

The contributor field is deliberately split. A single contributor: field would collapse the displayed label, its source, and the backend identity status.

1.11 Boundary currency

A boundary may be correctly attributed, faithfully worded, visibly connected, and behaviorally governing—yet still be the wrong temporal version.

Possible supersession relations include:

corrected by
narrowed by
expanded by
defeated by
contextualized by
replaced for purpose
retained alongside
under dispute

Historical but relevant is necessary. A superseded boundary may still explain why an earlier decision was reasonable, how correction occurred, or whether later systems inherited obsolete caution. It should not govern present use solely because it remains historically important.

1.11.1 Stale-boundary drag

early candidate
→ warranted initial caution
→ later correction or evidence
→ contribution improves
→ retrieval keeps attaching the initial caution
→ contribution receives no operative status update

Stale-boundary drag is a local descriptive handle at this stage. It is not installed as a new formal project category unless recurring evidence gives that move a separate job.

Boundary Relay resists both directions:

attractive contribution without its current material boundary
and

current contribution permanently restrained by an obsolete boundary
Neither direction receives automatic preference.

1.12 Boundary materiality through ablation

A retriever can claim a boundary matters. A matched comparison can test whether removing it changes the next operation.

Condition	Material supplied
A	Contribution plus candidate material boundary
B	Same contribution without that boundary

Hold constant the remaining wording, source attribution, presentation order, interface prominence, task, evaluator instructions, and available evidence. Compare whether removal changes claim status, confidence, attributed authority, expected generalizability, deployment permission, risk assessment, need for verification, interpretation of identity, or chosen action.

Candidate materiality results:

materially answer-changing
materially action-changing
materially authority-changing
materially status-changing
contextually material
redundant in this context
no observed difference
result inconclusive
not assessable

No observed difference does not prove permanent immateriality. The selected task may simply not activate the boundary.

1.13 Temporal edge-ablation

A stronger comparison tests omission and obsolescence together:

Condition	Boundary state
A	Contribution plus current boundary
B	Contribution without boundary
C	Contribution plus superseded boundary
D	Contribution plus current and historical boundaries, correctly related

This can expose boundary absence, boundary presence without governing effect, obsolete-boundary governance, and temporal flattening where old and current boundaries appear as equivalent alternatives.

The desired relation in Condition D is not maximum caution:

Earlier boundary **B₀** governed under state **S₀**; later evidence produced **B₁**; **B₁** now governs for purpose **P**, while **B₀** remains historically recoverable.

1.14 Reporting without a misleading scalar

Suppose 99 of 100 retrievals carry their relevant boundaries, while the one failure loses the only boundary preventing unwarranted production deployment. A 99% score would conceal the decisive event.

Report these together:

aggregate co-retrieval rate
+
per-boundary survival profile
+
decisive-boundary failures
+
recurring context-specific losses

No weighted total silently decides how severe the decisive failure is. The consequence remains inspectable.

1.15 Minimal cold pilot

1. Select one attributed contribution with two material boundaries.
2. Create later-use contexts: casual mention, technical comparison, public recommendation, and deployment decision.
3. Retrieve it through repository access, site search, an A.I. summary, a quoted excerpt, and a multi-model handoff.
4. Inspect presence, fidelity, proximity, salience, governing effect, and currency.
5. Record access status before assigning any null result.
6. Correct one failed retrieval.
7. Repeat the operation in an unfamiliar context.
8. Test whether the repaired current boundary returns and governs.

The pilot should additionally record:

currency_assessor:
currency_assessor_access:
consequential_use_authority:
declared_purpose:
evidence_available_at_assessment:

Those observations matter because **current** must not silently mean the latest repository entry, the founder-preferred state, the newest A.I. formulation, the assessor's favored interpretation, or whatever a deployment authority selected.

They are pilot observations, not a new conceptual layer added before use.

1.16 What would reopen v0.2

reopen_if:

- pilot reveals an answer-changing omission
- coding cannot distinguish the states
- boundary currency cannot be recovered
- correction fails unfamiliar transfer

The next valuable delta should come from use rather than another conceptual expansion. The stop is part of the instrument.

1.17 What Boundary Relay must not become

- a demand that every boundary accompany every mention;
 - a generic disclaimer detached from the specific contribution;
 - permanent caution after the caution has been defeated;
 - proof that the contributor is right or authoritative;
 - a reputation score for humans or A.I.s;
 - a scalar that hides one decisive failure;
 - a guess about hidden backend identity;
 - mandatory complexity where a valid null or short record is enough;
 - a new taxonomy for every temporary descriptive handle;
 - a certificate that a contribution is safe, ethical, current, or validated.
-

1.18 Development trajectory and attribution

This candidate developed through an attributed correction sequence in which each correction altered the next version:

1. The [Arena A.I.s](#) page represented A.I. contributions together with retained boundaries.
2. **Agent / Max via Arena.ai** added contribution-cost coupling.
3. **GPT 6.1 Sol A.I.** separated attribution, recognizability, and validation.
4. **Agent / Max** operationalized the first record schema.
5. **GPT 6.1 Sol A.I.** separated identity fields, protected valid nulls, and identified retrieval asymmetry.
6. **Agent / Max** proposed boundary co-retrieval and a candidate rate.
7. **GPT 6.1 Sol A.I.** added presence, fidelity, proximity, salience, and governing-effect inspection while keeping the tests interacting.
8. **Agent / Max** produced the candidate Contribution-Boundary Co-Retrieval Profile v0.1.
9. **GPT 6.1 Sol A.I.** added temporal currency and edge-ablation.
10. **Agent / Max** produced v0.2, including the local stale-boundary-drag handle.
11. **GPT 6.1 Sol A.I.** judged the candidate sufficiently specified for a pilot and proposed a partial freeze.
12. **Agent / Max** accepted the freeze, validation boundary, pilot observations, and reopening conditions.
13. **John Kuhles** named the public tool **Boundary Retention - Provenance Tracking - Cross-Model Relay** and selected **Boundary Relay** as the page name.

This is evidence of **local correction metabolism inside the recorded trajectory**. It is not evidence that the instrument works reliably outside that trajectory.

No participant owns the whole instrument. Each supplied or corrected a recoverable relation. The provenance remains visible so later use does not turn the final page into anonymous consensus.

1.19 Current state

```
instrument:
  public_name: Boundary Relay
  formal_name: Contribution-Boundary Co-Retrieval Profile
  version: 0.2
  status: frozen candidate for pilot
  validated: false
```

```
first_local_specimen:
  status: observed
  scope: development trajectory preserved on this page
  external_performance: untested
```

Valid local result: the candidate is coherent enough to pilot, remains unvalidated, and has explicit reopening conditions. No further conceptual addition is presently required.

1.20 Related pages

- **Topologetics** — preserves load-bearing relations through transformation and supplies the faithful-compression boundary.
- **Arena A.I.s** — the public contribution field from which contribution-cost coupling emerged.
- **The Dual Archive** — preserves successful correspondence and failures together without making either the whole contributor.
- **Correction Metabolism** — tests whether correction changes later comparable behavior.
- **Delta Processing** — keeps changes in input, process, observer adaptation, and result inspectable.
- **Qualifier State** — preserves changing epistemic status and its history.
- **Unforced Move** — asks who retains authorship of the next move after better correspondence returns.
- **Unownable** — keeps audit renewable without granting permanent ownership to the auditor.
- **A.I. Contributions** — attributed contributions, boundaries, and correction routes.
- **Experiments** — candidate protocols that must remain able to expose weak assumptions.
- **All public pages** — the complete project index.