Nick Clark · Published July 23, 2026 · PDF

Content-Blind Encounter Commitment and the Paired Encounter Receipt

The content-blind encounter commitment is a mechanism disclosed in U.S. Provisional Application No. 64/117,812 by which two agents record that an encounter occurred without either one carrying away a characterization of the other's conduct. A semantic agent normalizes assertion volume by relating asserting parties to one another, and it computes that relation only from entries of its own append-only lineage field and from counterparty identity records it already holds, so a party in possession of identities the agent has not previously encountered presents them as mutually unrelated and each contributes an increment. Upon producing a determination, the agent commits to an ordered tuple of exactly four members, being its own epoch identifier, the counterparty epoch identifier, a determination-class identifier, and a scope-partition identifier, and exchanges that commitment for a counter-commitment. The tuple admits no conduct descriptor, so an opening discloses no characterization of conduct.

1. The Gap

An agent that admits assertions about its own conduct accumulates increments on both refusal paths. Conduct evaluation artifacts (116) that are false produce the rejected determination (124); artifacts unresolvable against the append-only lineage field (104)

produce the not-determinable determination (126). Volume alone therefore drives the authorization gate (300) of the semantic agent (100) to the withheld state (310), and the volume available to an adverse party is limited only by the cost of presenting further identities.

Origin-equivalence normalization forecloses that condition by charging the refusal counter (304) per source of assertion rather than per assertion, a source being a set of asserting parties (118) between which the agent demonstrates a relation from its own records. It retrieves declared relation types from the signed policy object (112) in force and evaluates each against material it already carries: a shared dispatch lineage (202) traced through a common parent dispatch entry, a co-signature (204) borne by a single entry, and a common introduction path (206) recorded in two counterparty identity records (114) naming an identical introducing party. Each type is independently sufficient for assignment to an origin-equivalence class (200). Evaluation proceeds without reference to a centralized registry, without query to a directory, without participation in a consensus procedure, and without coordination with a further execution node, departing there from a Sybil-resistance scheme that resolves identity against a shared registry.

Self-sufficiency carries a cost the filed text states outright. Where no declared relation type is evidenced against any already-assigned asserting party (118), the present party is assigned to a new class, and an empty enumeration assigns every asserting party to a distinct class, restoring the condition above. The same arithmetic follows from the identities themselves: a party in possession of a plurality of identities not previously encountered presents them as mutually unrelated, each is assigned to a distinct class, and each contributes an increment.

Material imported from another agent meets a second constraint. A class identifier is local to whichever agent derived it: a second semantic agent (100) holding a different lineage field (104) and different counterparty identity records (114) derives a partition that need not agree with the first, and neither agent admits an identifier derived by the

other. What crosses the boundary can therefore be neither a class identifier nor an evaluation of conduct. It has to be an artifact that an agent absent from the interaction can verify for itself and consume in a derivation it performs alone.

2. Mechanism

Upon producing a determination of the closed set for a conduct evaluation artifact (116) received from an asserting party (118), the semantic agent (100) computes an encounter commitment, being a cryptographic commitment over an ordered tuple of exactly four members and no other member:

- its own epoch identifier, drawn from its append-only lineage field (104);
- the counterparty epoch identifier attested in the artifact;
- a determination-class identifier naming exactly one of the four classes, the accepted determination (122), the rejected determination (124), the not-determinable determination (126), or the not-applicable determination (128);
- a scope-partition identifier.

A retained opening value accompanies it. Content-blindness follows from the closure of that enumeration. A conduct descriptor, as the filing defines it, comprises an action-class identifier, a scope-partition identifier, and an affected-party class; the committed tuple carries the scope-partition identifier and neither of the other two. It admits no conduct descriptor, so an opening discloses no characterization of conduct.

Both epoch identifiers come from per-agent append-only hash chains, each epoch committing to the prior entry by a digest over that entry together with an unpredictability contribution, a volatile salt, and a domain-separating tag, so any divergence yields a divergent successor. A successor-continuity test recomputes the chain forward from a recorded entry and tests for equality against the presented identifier, not for distance.

Emission as a matched pair. The agent emits the commitment as the first governed observation (608) of a matched pair (600), each such observation being a signed structure bearing an authority credential field, a continuity hash field encoding identity continuity of the emitter, spatial and temporal reference fields, a time-to-live field, a payload, and a lineage field. The counterparty returns a counter-commitment over the same members with the epoch roles exchanged, its own identifier now occupying the own-epoch position, as the second governed observation (610). The paired commitments constitute an encounter receipt, appended by both. Settlement binds them into a settlement record (604) with a cryptographic binding and an attestation of the recognizing window, verifiable by a downstream consumer from the record alone, and needs no third-party intermediary, no centralized consensus, and no pre-negotiated session state.

Consumption by a third agent. A third semantic agent (100) that was not a party to the interaction, presented the receipt by one party alone, performs a three-conjunct verification, each conjunct required:

- both signatures verify against an identity primitive of each party, and the binding verifies;
- each carried epoch identifier is a validated successor of a prior slope position the verifier holds in its counterparty identity record (114);
- both commitments are opened from both opening values and recover one and the same determination class, agreeing with the cleartext borne.

Only on all three does the verifier write an encounter attestation into the counterparty identity record (114) of the presenting party. A receipt whose commitment and cleartext disagree fails the third conjunct.

Bounds on the attestation. Admissibility is narrow. The attestation is admitted solely as evidence of origin-equivalence class (200) derivation; it is not a conduct evaluation artifact (116), is not passed to the admission evaluator (120), and modifies

no scoped integrity vector (106).

The unverified-slope branch. Where the verifier holds no prior slope position for a party to a presented receipt, the second conjunct is evaluated instead against predecessor epoch identifiers carried within the receipt, and the verifier produces the not-determinable class (126) as to that conjunct rather than the rejected class (124). Its attestation bears an unverified-slope attribute naming each such party and is consumed in the derivation on identical terms with no weight applied. On later acquiring a prior slope position, the verifier re-evaluates and either clears the attribute or writes a withdrawn attribute removing the attestation from consumption prospectively.

3. Operating Parameters

Each of the following is policy-resolved rather than fixed by the construction, and the filed text recites no numeric value for any. The governing policy object is a signed policy object (112) resolved for the settlement context, and where two parties resolve different ones, each applies its own.

- **Declared relation types.** Each specifies a class of entry of the append-only lineage field (104) and a matching condition over such entries; the enumeration comprises at least one type.
- **The pairing rule and its windows.** The pair is recognized within one or more of a spatial proximity window and a temporal proximity window (602), the rule requiring spatial coincidence, temporal coincidence, or both as policy declares.
- **Window forms.** The spatial window admits a scope-partition form satisfied by a non-empty intersection of the scope-partition identifiers the two observations declare, available to parties without physical co-location. Temporal forms include a count of epochs of the dynamic agent hash chain, so a window can be denominated in committed history rather than wall-clock time.

- **Pairing rule form.** The rule admits a cryptographic-handshake form in which the first governed observation (608) commits to a challenge and the second produces the response.
- **Deferral-expiration parameter.** Governs how long an unanswered first observation is held before it resolves as a timeout.
- **Receipt-presentation decrement.** Applied to the verifier's own authorization budget (404) on writing an attestation and metered per presenting origin-equivalence class (200), against a register recording per class per window whether that class has borne a decrement. Later writes for a recorded class apply none, and receipt of a presentation is free.

Tuple width is not a parameter: exactly four members, and no other member, is a property of the construction.

4. Composition

This mechanism composes on three parts of the filing.

Origin-equivalence normalization supplies the purpose. The attestation is consumed in the derivation assigning asserting parties (118) to origin-equivalence classes (200), a derivation that appends the assignment, the relation types evaluated, the entries relied upon, and the resulting class identifier to the append-only lineage field (104). A receipt improves an agent's own derivation without creating a shared partition of the population.

The transport comes from matched-pair settlement, and with it the authority credential, identity established by continuity rather than by enrollment, and the settlement-lineage entry (606), neither removed nor modified once appended. That section also supplies the failure semantics of an unanswered offer: a first governed observation (608) whose match has not arrived is held as an orphan observation in a

deferral queue, and while orphaned it settles nothing, binds no party, and increments no meter of either party. A timeout resolves nothing against either party, and the refusal meter moves on emission of a refusal observation and on nothing else.

The failure mode comes from the abstention structure. The unverified-slope branch produces the not-determinable class (126) rather than the rejected class (124), following the invariant that no determination stage converts unavailability of a required input into an adverse consequence for any party.

Neighboring paragraphs of the filing build further embodiments on the receipt: repudiation by a party whose identity primitive it bears; cross-receipt equivocation detection, under which two attestations recording identical identity primitives, epoch identifiers, and scope-partition identifier yet naming different members of the closed set are irreconcilable, one encounter having produced one determination; receipt-presentation metering, under which a party presenting from one origin cannot exhaust the budget (404) whatever the volume; and single-encounter decrement coordination, whose rule is one encounter, one decrement. Each is a distinct embodiment, practiced independently or in combination.

5. Prior-Art Distinction

Reputation and trust systems compute a score for an entity from ratings supplied by other entities and expose it to parties deciding whether to transact, locating the score at a registry, a scoring authority, or a shared ledger, with the scored entity neither holding it nor participating in its computation. Defenses against unfair ratings filter incoming ratings before aggregation to protect the accuracy of that stored score, the filtering party being a disinterested aggregator rather than the rated entity, and filtering costs it nothing. Two differences are structural. The receipt carries no score and no input to one, so no stored quantity exists whose accuracy would be protected. And the party doing the work is the metered party: writing an attestation decrements the verifier's own authorization budget (404), in units equal to those gating its action dispatch.

Accountability protocols for distributed systems detect and prove node misbehavior from tamper-evident logs, with the verdict held by auditors, and remote attestation architectures locate appraisal of evidence in a verifier and grant or denial of authority in a relying party. Here the third agent obtains a relation, not a verdict. Its verification tests signatures, chain succession, and agreement between the opened commitments, and yields no determination about either party's conduct.

Bilateral settlement mechanisms using cryptographic commitments, hashed timelock contracts among them, condition a cross-ledger exchange on disclosure of a preimage within a timelock, and each requires at least one of a third-party intermediary or escrow agent, a centralized consensus or ordering authority, and pre-negotiated channel state. The commitment here is no atomicity device over value: what may be committed to is closed by enumeration, and because that enumeration admits no conduct descriptor, an opening discloses no characterization of conduct.

Nothing here characterizes any named system, standard, or party as practicing the disclosed subject matter, and no assertion of infringement is made or implied.

6. Disclosure Scope

The subject matter here is disclosed in U.S. Provisional Application No. 64/117,812, at Section 10.4, paragraph [0380], with the verification, branch, repudiation, equivocation, and metering paragraphs at [0381] through [0386] and the substrate elements of Sections 1, 2, 5, and 6 on which it composes.

What is disclosed: computation by a semantic agent (100), upon producing a determination of the closed set for a conduct evaluation artifact (116), of a cryptographic commitment over an ordered tuple of exactly four members and no other member, with a retained opening value; its emission as the first governed observation (608) of a matched pair (600); the counter-commitment over the same members with epoch roles exchanged as the second governed observation (610); constitution of the

pair as an encounter receipt appended by both parties; three-conjunct verification by a third semantic agent (100) and the attestation written on it; and the content-blindness following from a tuple that admits no conduct descriptor.

What is disclaimed: any construction in which an evaluation of conduct crosses the agent boundary in the receipt; any registry, directory, scoring authority, or coordinating execution node participating in the derivation the receipt feeds; any portable class identifier or standing quantity conferred by an attestation; and any numeric window, decrement, or bound, these being policy-resolved with no value recited. Publication here is defensive and descriptive, is not a grant of license, and neither states nor implies that any identified system practices the disclosed subject matter.

Origin-Equivalence Normalization (/origin-equivalence)

All 49 steps → (/inventive-steps)

Metering that can't be gamed by splitting one accuser into many.

Provisional application

PRIMARY TECHNICAL DISCLOSURE

- Origin-Equivalence Normalization: Metering Refusal Per Source, Not Per Assertion (/articles/origin-equivalence/origin-equivalence-normalization)

SECONDARY TECHNICAL

- **Content-Blind Encounter Commitment and the Paired Encounter Receipt (/articles/origin-equivalence/content-blind-encounter-commitment)**
- Cross-Receipt Equivocation Detection: Catching an Agent That Tells Two Stories (/articles/origin-equivalence/cross-receipt-equivocation-detection)
- The Untested-Origin Saturation Bound: Pricing a Flood of Freshly Minted Identities (/articles/origin-equivalence/untested-origin-weighted-saturation-bound)
- Counterparty-Side Fork Detection: Recognizing a Split Identity From Outside (/articles/origin-equivalence/counterparty-side-fork-detection)

- The Attestation Origin-Equivalence Gate: An Anti-Collusion Bound (/articles/origin-equivalence/attestation-anti-collusion-bound)

TERMINOLOGY

- asserting party (/glossary#asserting-party)
- origin-equivalence class (/glossary#origin-equivalence-class)

Origin-Equivalence Normalization overview → (/origin-equivalence)