

The Untested-Origin Saturation Bound: Pricing a Flood of Freshly Minted Identities

An adversary that cannot defeat a semantic agent's determination logic can still aim volume at its arithmetic, and Section 2.1 of U.S. Provisional Application No. 64/117,812 records the reason that works: volume alone drives the authorization gate (300) toward the withheld state (310), and the volume available to an adverse party is limited only by the cost of presenting further identities. The untested-origin-weighted saturation bound answers that condition at the containment layer. Per Section 10.2, the per-partition propagation-suspension bound is multiplied by a factor of one less the proportion of a scope partition's conduct evaluation artifacts (116) originating from origin-equivalence classes (200) designated untested. A partition whose artifacts all originate from untested classes triggers containment at a receipt rate approaching zero; a partition none of whose artifacts so originate triggers at the full declared bound. The weighting moves the threshold and not the count, and the authorization gate (300) is not written by reason of suspension.

1. The Gap

A semantic agent (100) that partitions its own cognitive state limits what any one context can do to it. Under Section 10.2 of the filing, the agent maintains a plurality of scope partitions, each carrying its own instance of the scoped integrity vector (106), self-esteem aggregate (108), refusal counter (304), authorization budget (404), and a per-partition segment of the append-only lineage field (104), keyed by a partition identifier declared in the signed policy object (112) and referenced by the conduct descriptor. On a determination, every consequent modification, counter increment, and writing of the authorization gate (300) applies to the identified partition's instances alone. The gate stays granting for classes of every other partition, and execution continues there.

One route crosses that boundary, and the filing makes it a declaration rather than a default. State movement in a first partition propagates to a second only where the first partition's coupling declaration enumerates the second partition and enumerates the empathy-scope designation implicated by the conduct descriptor. A partition whose coupling declaration enumerates nothing is fully confined.

The declared route is then conditioned on load. Where a partition's receipt rate of conduct evaluation artifacts (116) over a declared window exceeds a declared bound while the aggregate rate across the remaining partitions, taken as the greatest of their partition rates, stays below it, the agent suspends cross-partition propagation out of that partition alone and appends a containment determination recording the partition, both rates, the bound, and the window.

That condition compares a count against a constant. Section 2.1 records why the constant is the exposed term: false artifacts and artifacts unresolvable against the append-only lineage field (104) both accumulate increments, volume alone therefore drives the authorization gate (300) to the withheld state (310), and the volume

available to an adverse party is limited only by the cost of presenting further identities. Against a fixed bound, artifacts arriving from identities stood up this morning and artifacts arriving from counterparties holding recorded histories count alike.

2. Mechanism

Paragraph [0357] leaves the condition, the counters, and the window where Section 10.2 put them and changes one term. The per-partition suspension bound is multiplied by a factor of one less the proportion of a partition's conduct evaluation artifacts (116) originating from origin-equivalence classes (200) designated untested. A partition whose artifacts all originate from untested classes triggers containment at a rate approaching zero; a partition none of whose artifacts so originate triggers at the full declared bound. The filing states the composition in one clause: saturation by freshly minted identities becomes cheaper to contain than saturation by counterparties holding recorded histories.

The threshold moves, not the count. The partition receipt rate stays what paragraph [0355] defines, a count of artifacts bearing that partition identifier over the declared window. Nothing arriving is re-scored, and the containment determination still records the partition, both rates, the bound, and the window.

What makes a class untested is a resolution over its edges, not the agent's unfamiliarity with it. Section 2.5 defines the severance-survival test. A severance event is recorded against a constituent edge upon append to the append-only lineage field (104) and carries a class, being one of a payment class and an obligation class. An edge survives such an event where, within a declared continuation interval following it, a subsequent event of that class is recorded against that edge and re-establishes it. Edges resolve want-sustained, payment-sustained, obligation-sustained, or not-typeable, and where every constituent edge of an origin-equivalence class (200) resolves not-typeable, the class is designated an untested class. The filing is deliberate about the boundary: an edge against which the test does not resolve, for want of a recorded severance event, of

a recorded outcome, or of elapse of the interval, is likewise an edge resolving not-typeable, and the designation rests upon that resolution alone and upon no narrower condition. A party presenting identities not previously encountered as mutually unrelated asserting parties (118) therefore produces a distinct class per identity, each untested until an edge of it survives a severance.

The evidence is the agent's own. The class (200) is computed by the semantic agent (100) upon entries of its own append-only lineage field (104) and the counterparty identity records (114) it holds, over relation types declared in the signed policy object (112): a shared dispatch lineage (202), a co-signature (204), a common introduction path (206). No registry, directory, consensus procedure, or further execution node is consulted. Severance events are equally local, recorded upon append to that same field.

Suspension itself is recited narrowly. Containment halts propagation out of the contained partition alone, the agent continues applying that partition's movements within it, and the gate (300) is not written by reason of the suspension: containment withdraws reach into other partitions, not the ability to act. Propagation resumes after the rate stays below the bound for a declared interval.

The designation is cleared by the record, not by elapsed time. Section 2.5 fixes the recomputation triggers at three and no others: receipt of a conduct evaluation artifact (116) from an asserting party (118) assigned to that class, before the increment procedure consults the per-class increment register; append of a severance event against a constituent edge, or of an entry carrying such an edge to a survived value; and a merge of the class, the merged class being untested where and only where every constituent edge of it resolves not-typeable. A recomputation that changes the designation is itself appended, governs contributions applied after the append, and does not alter one already applied within the current window. A party newly encountered and thereafter transacting accumulates such events, and its class ceases to be designated an untested class.

Spreading the load does not evade the structure. Where saturation is distributed so that the per-partition condition fails because the aggregate across the remaining partitions also exceeds the bound, the global saturation-defeat remedy applies: the agent sums every partition's rate over the window, compares that sum against a global bound declared in the signed policy object (112), suspends propagation out of every partition, and appends a global containment determination distinct from the per-partition one. Again the gate is not written. The untested-origin weighting acts on the per-partition bound; the global condition closes the even-distribution evasion.

3. Operating Parameters

The filing declares parameters and not values. Each is retrieved from the signed policy object (112) in force.

- **Per-partition suspension bound.** Declared, and compared against the partition receipt rate.
- **Weighting factor.** One less the proportion of a partition's conduct evaluation artifacts (116) originating from origin-equivalence classes (200) designated untested. At a proportion of one the effective bound approaches zero; at a proportion of zero it is the full declared bound.
- **Window.** Declared, and used for the partition receipt rate and for the aggregate rate alike.
- **Aggregate rate.** The greatest of the partition receipt rates of the remaining partitions over that window.
- **Resumption interval.** Propagation resumes after the rate stays below the bound for a declared interval.
- **Global bound.** Declared separately, compared against the sum of every partition's rate, and producing a global containment determination distinct from the per-partition one.

- **Continuation interval.** Declared, and governing whether a recorded severance is survived.
- **Recomputation triggers.** Three, and no others, as recited in Section 2 above.

Two further declared quantities belong to the cost multiplier of Section 2.5 rather than to this mechanism, and are noted only because they govern the same designation: the multiplier is greater than zero and less than unity, and the introduction-path depth is at least one.

4. Composition

Within Section 10.2 this mechanism sits last in the containment run, after confinement of state movement, declared cross-partition coupling, attack-conditional propagation suspension, and the global saturation-defeat remedy. As recited, paragraph [0357] introduces no further condition and no further counter: it multiplies a bound by a proportion taken over artifacts the suspension condition already counts, using a designation Section 2.5 already maintains.

The same predicate runs backward into Section 2 with opposite valence. On the metering path, an untested class whose recorded introduction paths converge upon a common ancestor entry has a declared cost multiplier applied to its contribution to the refusal counter (304), reducing that contribution below the contribution of a single class whose edges resolve want-sustained, and the filing gives the reason in terms: a newly encountered party retains capacity to correct the semantic agent (100). On the containment path the untested proportion instead lowers the effective bound, the stated effect being that saturation by freshly minted identities is cheaper to contain. As filed, paragraph [0357] conditions the containment weighting on the untested proportion alone and does not restate the convergence requirement that governs the cost multiplier.

The predicate appears a third time in Section 10.1. The severance-conditioned renewal amount makes both the amount by which a renewal raises the authorization quantity and the amount by which it replenishes the authorization budget (404) a monotonically non-decreasing function of the count of recorded severance events across the renewing class's constituent edges, a class each of whose edges resolves not-typeable renewing at a declared least amount. A principal standing up fresh identities that classify into classes bearing no recorded severance event therefore renews its own agent's authorization at that least amount and does not restore full authorization. One resolution over edges, three recited consequences.

Containment also carries across the agent boundary: Section 10.8 recites that propagation is suspended and no determination propagation record is given effect while the containment determination stands for the scope partition concerned in either agent. Origin-equivalence normalization is drawn in FIG. 2, the refusal counter and its accumulators in FIG. 3.

5. Prior-Art Distinction

The categories below are described structurally, following the filing's own background where it reaches them. Nothing here characterizes any particular product or implementation as within or outside the disclosed subject matter.

Reputation and trust systems. As the background describes them, such systems compute a score for an entity from ratings supplied by other entities, locate the score at a registry, a scoring authority, or a shared ledger, and expose it to parties deciding whether to transact; the scored entity neither holds the score nor participates in its computation. The class identifier written into a counterparty identity record (114) under Section 2.2 is expressly not a determination concerning that party's conduct, and the class is local: a second agent holding different records derives a partition that need not agree, and neither agent admits a class identifier derived by the other.

Identity-cost schemes. The filing distinguishes a Sybil-resistance scheme that resolves identity against a shared registry or a coordinating authority. The untested predicate here does not mean unregistered, unstaked, or unfamiliar. It means no constituent edge of the class has survived a recorded severance of a payment or obligation class within the declared continuation interval, resolved from the agent's own append-only record.

Policy decision points. The background addresses systems implementing the XACML standard, which return a decision drawn from permit, deny, not-applicable, and indeterminate. That result is directed at the request, produces no persistent state in the deciding entity, and constrains no subsequent operation of the decider. A containment determination is persistent state in the deciding agent, governing its own later state movements until the resumption condition is met.

Volume controls generally. Mechanisms that throttle, challenge, or drop traffic once a rate is exceeded place the consequence on the sender. The recited consequence here is internal: propagation out of the contained partition is suspended, movements continue to apply within it, and the gate (300) is not written.

The distinguishing structure is narrow. The trip condition is not receipt volume alone but receipt volume measured against a bound scaled by how much of that volume originates from classes the deciding agent's own record resolves as never having survived a severance.

6. Disclosure Scope

This mechanism is disclosed in U.S. Provisional Application No. 64/117,812, Section 10.2, at paragraph [0357], supported by the scope-partitioned containment paragraphs [0353] through [0356] and by the untested-class designation and severance-survival test of Section 2.5 at paragraphs [0071] through [0079].

Disclosed is the multiplication of a per-partition propagation-suspension bound by a factor of one less the proportion of that partition's conduct evaluation artifacts (116) originating from origin-equivalence classes (200) designated untested; the endpoint behavior at untested proportions of one and of zero; the composition of that weighting upon the receipt-rate suspension condition, its containment determination, and its resumption condition; the computation of the class and of the designation by the agent upon its own records; and the recited property that the authorization gate (300) is not written by reason of suspension.

Disclaimed is any particular numeric bound, global bound, window length, continuation interval, or resumption interval, each being declared in the signed policy object (112) and none fixed by the disclosure; and any reading of this publication as a claim to rate limiting, identity-cost schemes, reputation scoring, or fault isolation as general techniques.

This article is a technical disclosure published to establish public availability of that subject matter as of the date shown. The provisional application is the operative and citable disclosure, and nothing here enlarges or narrows it.

Origin-Equivalence Normalization ([All 49 steps → \(/inventive-steps\)](#))

origin-equivalence

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 \(/articles/origin-equivalence/content-blind-encounter-commitment\)](/articles/origin-equivalence/content-blind-encounter-commitment)
- [Cross-Receipt Equivocation Detection: Catching an Agent That Tells Two Stories \(/articles/origin-equivalence/cross-receipt-equivocation-detection\)](/articles/origin-equivalence/cross-receipt-equivocation-detection)
- [**The Untested-Origin Saturation Bound: Pricing a Flood of Freshly Minted Identities \(/articles/origin-equivalence/untested-origin-weighted-saturation-bound\)**](/articles/origin-equivalence/untested-origin-weighted-saturation-bound)
- [Counterparty-Side Fork Detection: Recognizing a Split Identity From Outside \(/articles/origin-equivalence/counterparty-side-fork-detection\)](/articles/origin-equivalence/counterparty-side-fork-detection)
- [The Attestation Origin-Equivalence Gate: An Anti-Collusion Bound \(/articles/origin-equivalence/attestation-anti-collusion-bound\)](/articles/origin-equivalence/attestation-anti-collusion-bound)

TERMINOLOGY

- [asserting party \(/glossary#asserting-party\)](/glossary#asserting-party)
- [origin-equivalence class \(/glossary#origin-equivalence-class\)](/glossary#origin-equivalence-class)

[Origin-Equivalence Normalization overview → \(/origin-equivalence\)](/origin-equivalence)