

When One Accuser Arrives Wearing Forty Faces

A procurement operator at a specialty parts distributor comes in on an ordinary Wednesday and finds her semantic agent unable to place the quarter's replenishment orders. Overnight it received forty conduct evaluation artifacts from forty asserting parties it had never seen before. Each accumulated an increment against the refusal counter, the rate threshold was satisfied, and the authorization gate moved to the withheld state. The allocation window closes at noon and does not reopen. U.S. Provisional Application No. 64/117,812 discloses an origin-equivalence normalization in which the refusal counter is charged per source of assertion rather than per assertion, with the source derived by the agent from its own records.

The Wednesday the Orders Did Not Go Out

The operator runs procurement for a specialty parts distributor of about ninety people. Her semantic agent (100) places replenishment orders against a supplier pool she spent two years assembling, under a signed policy object (112) she wrote herself, with a rate threshold she chose to match how often her suppliers had historically disputed anything.

She comes in at seven forty. The agent has not placed the quarter's orders. Its authorization gate (300) is in the withheld state (310).

Overnight it received forty conduct evaluation artifacts (116) from forty asserting parties (118), none of which had transacted with her agent before.

Fourteen asserted facts her agent could check against its append-only lineage field (104) and found false, producing the rejected determination (124). The other twenty-six asserted facts her agent had no entry to resolve against, producing the not-determinable determination (126). Under the metering her deployment uses, both paths accumulate increments against the refusal counter (304). Forty artifacts, forty increments, a rate threshold of five. By two in the morning her agent had stopped being able to act.

What she cannot do at seven forty is release the replenishment orders. Her allocation window with two of her three primary suppliers closes at noon, and capacity is committed to whoever confirms inside it. She can read all forty. She cannot tell her agent they came from the same place, because in her setup the agent has no procedure for deciding that they did, and she has no override short of disabling the metering for the rest of the quarter.

At noon the window closes.

What She Cannot Rebuild

She loses the quarter's allocation. That is the smallest loss; it is money, and she can quantify it by Friday.

The loss that does not come back is the calibration. Her rate threshold of five was not arbitrary. It was the number she arrived at by watching two years of real supplier disputes and deciding how much contested conduct her agent should tolerate before it stopped acting on her behalf. In her deployment as it stands she has two options, and neither returns her to where she was. She can raise the threshold past the reach of forty

artifacts, which also raises it past the level at which four genuine disputes from four distinct suppliers would stop her agent. Or she can leave it where it is and accept that any party willing to present forty identities can park her agent at will.

The threshold was the artifact of her judgment. She cannot restore it, because Wednesday did not tell her which of her two beliefs was wrong: whether five was too low, or whether counting per artifact was the wrong thing to count.

There is a second thing she does not get back. Her append-only lineage field (104) shows forty artifacts and forty increments, exactly what forty distinct aggrieved suppliers would have produced. When she explains the missed allocation on Friday, she has determinations and increments and nothing that speaks to origin. In her deployment as configured today, the volume available to whoever sent those artifacts was limited by the cost of presenting further identities and by nothing else in her agent's procedure.

Why Her Setup Cannot Simply Count Better

The obvious repairs are not available to her on the terms she needs.

She could subscribe her agent to a shared reputation service and let it ask whether these forty identities are known to be one party. That trades her problem for one she is unwilling to take: her agent's ability to transact would then depend on an outside party's view of her counterparties, and several of her small suppliers would appear nowhere in such a service. She built her agent to act on its own records so that her supplier relationships, not someone else's directory, determined how it behaved.

She could try to identify the forty by inspection. But what she wants her agent to notice is not a property of any single artifact. It is a relation between asserting parties (118), and in her setup the evidence of such a relation sits in entries her agent already holds and has never been asked to compare.

There is a sharper difficulty in her case. Suppose she configured her agent to discount any party it had not seen before. Two of her three primary suppliers were new eighteen months ago, and one was right the first time it contradicted her agent's record. Were she to blunt the contribution of every newly encountered party, she would be blunting the parties whose corrections she most needs to hear. Her problem is not that new parties should count for less. It is that in her deployment forty faces of one origin count forty times, and she has no way to say so from her own records.

What the Filed Mechanism Does

U.S. Provisional Application No. 64/117,812 discloses, in accordance with an embodiment, incrementing the refusal counter (304) per origin-equivalence class (200) rather than per conduct evaluation artifact (116): charging per source of assertion rather than per assertion, a source being a set of asserting parties (118) between which the semantic agent (100) demonstrates a relation from its own records. The disclosure states that the contribution of a single origin to the withholding of an action is bounded irrespective of the number of identities it presents, and that no registry, directory, or shared scoring service is consulted.

In a disclosed embodiment the assignment proceeds by an ordered procedure. Upon receipt of an artifact, the agent retrieves or instantiates the counterparty identity record (114) of the asserting party, then retrieves from the signed policy object (112) in force an enumeration of declared relation types, each specifying a class of entry of the append-only lineage field (104) and a matching condition over such entries. For each asserting party already assigned to a class, in the current window or a preceding one, it evaluates each declared relation type against the present party.

Three declared relation types are described. A shared dispatch lineage (202) is evidenced where the lineage field records a dispatch to each of the two parties and both entries record a common parent dispatch entry as their immediate antecedent. A co-signature (204) is evidenced where a single lineage entry bears a signature verifiable

against an identity primitive of each party. A common introduction path (206) is evidenced where the counterparty identity records of both parties record an identical introducing party. Each declared relation type is described as independently sufficient for assignment.

Where a relation is evidenced, the present asserting party is assigned to the class of the compared party, and where more than one class is identified, the classes are merged. Where none is evidenced, a new class is assigned. The assignment, the relation types evaluated, the entries relied upon, and the resulting class identifier are appended to the lineage field and written into the counterparty identity records of the class. The disclosure states that the identifier is not a determination about that party's conduct.

In that embodiment, the per-class increment register records, for each class, whether it has contributed an increment within the current window. The increment procedure consults the register before incrementing the rate accumulator and terminates without incrementing where the class has already contributed. An artifact from a party in an already-recorded class still produces a determination and is still appended to the lineage field, and applies no increment. As the disclosure puts it, normalization suppresses no determination and withholds no adjudication; it governs the metering alone. Normalization is performed upon receipt and before the register is consulted, so a party newly presented within a window is assigned to the class of its related parties first.

For classes with no relation evidenced between them, the disclosure describes an untested-class cost multiplier. A severance-survival test resolves a reason-type of each constituent edge, a severance event carrying either a payment class or an obligation class, and an edge surviving where a subsequent event of that class re-establishes it within a declared continuation interval. Where every constituent edge of a class resolves not-typeable, the class is designated an untested class. Where such a class is so designated and the recorded introduction paths of its parties converge upon a common ancestor entry within a depth declared in the signed policy object, a cost multiplier

declared in that object, greater than zero and less than unity, is applied to its contribution. The disclosure states that the multiplier does not exclude the class, does not suppress its determinations, and does not prevent it from contributing at full weight where it later accumulates recorded severance events, whereby a newly encountered party retains capacity to correct the agent. Assignment to a class is persistent and not window-scoped; the increment register alone resets.

The filing's worked trace is close to her Wednesday. A policy object declares a rate threshold of five, a multiplier of one fifth, and an introduction-path depth of two, and forty artifacts arrive from forty newly encountered parties sharing no declared relation type. Forty untested classes are formed. Thirty-five have introduction paths converging on a common ancestor at depth two and each adds one fifth to the weighted contribution register, which reaches seven; the remaining five do not converge and contribute at full weight, the accumulator reaching five. The window rate is twelve and the threshold is satisfied. Had all forty instead recorded a common introducing party, the trace states, all forty would have been assigned to one class by the common introduction path (206), that class would have contributed one increment, and the gate would have remained in the granting state.

Where This Leaves Her, Honestly

The class her agent derives would be local to her agent. The disclosure is explicit that a second semantic agent with a different lineage field and different counterparty identity records derives, from the same population, a partition that need not agree, that no procedure reconciles the two, and that no class identifier is transmitted or admitted between them. If she wanted a supplier's standing with her to travel to a peer's deployment, this would not give her that.

The normalization would also be only as strong as what her own records contain. Were an adverse party's forty identities to share no dispatch antecedent she has recorded, no co-signature she holds, and no recorded introducing party, her agent would evidence no

declared relation and would assign forty classes. The untested-class multiplier could reduce their contribution where their introduction paths converge within her declared depth, but that reduction is conditioned on a convergence her records actually show. The disclosure notes that an enumeration comprising no relation type assigns every party to a distinct class, so what she declares in her policy object is doing real work.

Nothing here would change what her agent decides about any artifact. Every one of Wednesday's forty would still have produced its determination and still have been appended to her lineage field. Were she hoping this quieted the complaints themselves, it would not; it changes what her refusal counter charges. The metering procedure her threshold is compared against is governed by a separate chapter of the filing.

A mid-window merge would also not rewrite her arithmetic retroactively. The disclosure describes a merged class as having contributed within the window where any constituent class is so recorded, with neither the rate accumulator nor the weighted contribution register decremented by the merge. For her, a relation her agent evidences at eleven does not refund an increment applied at two.

Disclosure Scope

This article is a technical description of subject matter disclosed in U.S. Provisional Application No. 64/117,812, using that filing's own mechanism names and reference numerals. Nothing in this article characterizes the scope of any claim, and nothing in it constitutes an admission concerning the state of the art. The scenario is illustrative and describes no actual party, deployment, or event.

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 \(/articles/origin-equivalence/origin-equivalence-normalization\)](/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)
- [The Bridging Party and the Deferred Merge: Charging One Windowed Increment Against Every Origin-Equivalence Class a Single Asserting Party Bridges \(/articles/origin-equivalence/bridging-party-deferred-merge\)](/articles/origin-equivalence/bridging-party-deferred-merge)
- [The Common-Execution-Node Relation Type: Charging Co-Hosted Asserting Parties as One Origin \(/articles/origin-equivalence/common-execution-node-relation-type\)](/articles/origin-equivalence/common-execution-node-relation-type)
- [Receiver-Side First-Encounter Budget Decrement: Pricing a Stranger's Arrival Per Presented-Material Origin Class \(/articles/origin-equivalence/receiver-side-first-encounter-budget-decrement\)](/articles/origin-equivalence/receiver-side-first-encounter-budget-decrement)
- [Mutual First Encounter Without Cross-Budget Coordination: Two Strangers, Two Matched Pairs, Two Separate Budget Decrements \(/articles/origin-equivalence/mutual-first-encounter-no-cross-budget\)](/articles/origin-equivalence/mutual-first-encounter-no-cross-budget)

APPLICATIONS • GENERAL

- [Sybil-Resistant Agent Networks: When One Accuser Becomes Many \(/articles/origin-equivalence/sybil-resistant-agent-networks\)](/articles/origin-equivalence/sybil-resistant-agent-networks)
- [Fraud Rings in Autonomous Commerce: Metering Per Source, Not Per Claim \(/articles/origin-equivalence/fraud-rings-in-autonomous-commerce\)](/articles/origin-equivalence/fraud-rings-in-autonomous-commerce)
- [Invalid Traffic and Click Fraud: Metering by Source, Not Impression \(/articles/origin-equivalence/invalid-traffic\)](/articles/origin-equivalence/invalid-traffic)

- [When One Accuser Arrives Wearing Forty Faces \(/articles/origin-equivalence/origin-equivalence-stakes\)](/articles/origin-equivalence/origin-equivalence-stakes).

APPLICATIONS • SPECIFIC

- [W3C Decentralized Identifiers and Verifiable Credentials: Verifying an Identifier Versus Metering Its Origin \(/articles/origin-equivalence/w3c-decentralized-identifiers\)](/articles/origin-equivalence/w3c-decentralized-identifiers).
- [Okta Auth for GenAI: Agent Identity and Origin Equivalence \(/articles/origin-equivalence/okta-auth-for-genai\)](/articles/origin-equivalence/okta-auth-for-genai).
- [WorkOS: Agent Identity Infrastructure and Source Metering \(/articles/origin-equivalence/workos-agent-identity\)](/articles/origin-equivalence/workos-agent-identity).
- [Entra Agent ID: Directory Identity and Derived Origin Classes \(/articles/origin-equivalence/microsoft-entra-agent-id\)](/articles/origin-equivalence/microsoft-entra-agent-id).
- [SPIFFE and SPIRE: Workload Identity and Origin Equivalence \(/articles/origin-equivalence/spiffe-spiire\)](/articles/origin-equivalence/spiffe-spiire).
- [Cloudflare Bot Management: Blocking Agents and Metering Sources \(/articles/origin-equivalence/cloudflare-bot-management\)](/articles/origin-equivalence/cloudflare-bot-management).
- [HUMAN Security: Bot Networks and Metering by Source \(/articles/origin-equivalence/human-security\)](/articles/origin-equivalence/human-security).

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)