Nick Clark · Published July 24, 2026 · PDF

W3C Decentralized Identifiers and Verifiable Credentials: Verifying an Identifier Versus Metering Its Origin

An autonomous agent that accepts assertions about its own conduct from parties it has never met can verify every signature it receives and still see its authorization gate driven to the withheld state, because a verified identifier says nothing about how many identifiers answer to one party. W3C Decentralized Identifiers and Verifiable Credentials, as publicly described, standardize how a subject controls an identifier and how an issuer attests claims about that subject, and they leave the number of identifiers a subject may hold unconstrained by design. U.S. Provisional Application No. 64/117,812 works at a different layer. Its origin-equivalence normalization charges a semantic agent's refusal counter per origin-equivalence class rather than per conduct evaluation artifact, the class being derived by the agent from its own append-only lineage field and the counterparty identity records it holds, with no registry, directory, or consensus procedure consulted.

Forty Valid Signatures, One Origin

A semantic agent operating between organizations takes in assertions about its own conduct from counterparties it did not select and cannot vet. Each arrives over the agent's ordinary interaction channel as a signed conduct evaluation artifact naming an

action class, a scope partition, and an affected-party class, and is tested against the append-only lineage field recording what the agent executed.

Verifying the signature is the tractable half of that exchange. Public key cryptography settles whether an artifact was produced by the holder of a particular key. What it cannot settle is how many of the keys filing artifacts answer to one party.

That gap has teeth, because the metering downstream is merit-independent by construction. An artifact the agent's own record affirmatively contradicts produces the rejected determination and applies an increment. An artifact the record cannot resolve produces the not-determinable determination, which applies an increment where the lineage field already records a prior such determination for the same asserted conduct event. Increments accumulate against a policy-declared rate threshold and a policy-declared run threshold, and the refusal counter writes the authorization gate where both comparisons are satisfied. Section 2.1 of the filed chapter states the exposure without softening it: volume alone drives the authorization gate to the withheld state, and the volume available to an adverse party is limited only by the cost of presenting further identities. Its worked trace runs the arithmetic at forty asserting parties.

The instinct at that point is to lean on identity infrastructure.

How the DID and Credential Model Is Publicly Described

Decentralized Identifiers and Verifiable Credentials are published by the World Wide Web Consortium, which develops specifications through an open, multi-stakeholder process and publishes them as Recommendations. What follows summarizes those public materials. As publicly described, a Decentralized Identifier is a globally unique identifier that its subject controls directly rather than one issued and revocable by a central registrar. It resolves, through a mechanism specific to its method, to a document publishing the verification material and service endpoints a relying party needs in order to authenticate the subject and interact with it.

A Verifiable Credential, as documented publicly, is a set of claims about a subject that an issuer signs cryptographically. A holder presents it to a verifier, and the verifier checks the proof and the credential's status without necessarily contacting the issuer at the moment of presentation. The model keeps the roles of issuer, holder, and verifier distinct and is publicly described as flexible about serialization and proof format, so different underlying mechanisms sit beneath one data model.

The method surface is heterogeneous on purpose. As publicly described, some methods resolve deterministically from the identifier itself, some anchor the resolution target under a web domain or to a distributed ledger, and some are intended for pairwise relationships in which the identifier never appears in a public registry.

None of this is a Sybil-resistance architecture, and the specifications do not present themselves as one. That is a design commitment, not an oversight. Correlation resistance is a publicly stated privacy objective of the work: a subject is meant to present different identifiers to different relying parties so those parties cannot join their records and track it, which is why pairwise identifiers exist at all. Making identifiers cheap and plural is how the model protects the party behind them. Who may issue which credential, and how much weight a verifier should give it, is left to trust frameworks layered above the data model.

So an agent that receives forty conduct evaluation artifacts from forty distinct identifiers, each carrying a valid proof, holds exactly forty valid proofs. Nothing has failed. The model answered its own question; the agent's remaining question is a different one.

Charging the Source Instead of the Claim

Chapter 2 of U.S. Provisional Application No. 64/117,812 does not try to determine who anyone is. It changes what gets charged. The refusal counter is incremented per origin-equivalence class rather than per conduct evaluation artifact, a class being a set of

asserting parties between which the semantic agent demonstrates a relation from its own records. The contribution of a single origin to the withholding of an action is thereby bounded irrespective of how many identities that origin presents.

Derivation runs as an ordered procedure on receipt of an artifact. The agent retrieves the counterparty identity record of the asserting party, or instantiates one, then reads from the signed policy object an enumeration of declared relation types, each specifying a class of lineage entry and a matching condition over such entries. Every declared type is evaluated between the present party and each party already assigned to a class, in the current window or a preceding one.

The chapter recites three relation types, each independently sufficient for assignment. A **shared dispatch lineage** is evidenced where the lineage field holds an entry recording a dispatch to the present party and a second recording a dispatch to the compared party, both recording a common parent dispatch entry as their immediate antecedent. A **co-signature** is evidenced where a single lineage entry bears a signature verifiable against an identity primitive of each of the two parties. A **common introduction path** is evidenced where the counterparty identity records of both parties each record an introducing party and the two are identical. Where any type is evidenced the present party joins the compared party's class; where more than one class is identified the classes merge; where nothing is evidenced a new class forms.

The evaluation runs on the agent's own lineage entries and the counterparty identity records it holds, 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. The class is strictly local: a second agent holding different records derives a partition that need not agree with the first, nothing reconciles the two, and neither admits a class identifier derived by the other.

Assignment persists rather than expiring with the window, and the per-class increment register alone is reset. On a merge, the merged class counts as having contributed where any constituent class did, so it contributes within one window the minimum of one increment and the number of constituent classes that contributed.

Section 2.5 addresses identities the agent has never encountered, which would otherwise each form a distinct class contributing at full weight. A class is designated an untested class where every constituent edge resolves not-typeable under the severance-survival test, meaning no severance of either the payment class or the obligation class has been recorded against that edge and survived within the declared continuation interval. Where a class is so designated **and** its recorded introduction paths converge upon a common ancestor entry within a depth declared in the signed policy object, a cost multiplier greater than zero and less than unity applies to that class's contribution, which then accumulates in a weighted contribution register holding a rational quantity. The rate of the window is the rate accumulator's integer count plus that register rounded down, rounding being performed at comparison time alone.

Normalization suppresses no determination and withholds no adjudication. An artifact from a party in a class already recorded still produces a determination and is still appended to the lineage field, only the contribution to the refusal counter being aggregated. The filed chapter puts the scope plainly: the normalization governs the metering alone.

Two Different Questions About the Same Counterparty

The two architectures answer questions that do not overlap. A Decentralized Identifier answers which subject this is and how to authenticate it. A Verifiable Credential answers what a trusted party has asserted about that subject. Origin-equivalence normalization answers neither: it answers how much a source may contribute, within one window, to the withholding of the agent's own execution authorization.

The evidence differs accordingly. Credential verification rests on cryptographic proof and issuer trust, both properties of material presented to the verifier. Class assignment rests on relations the agent evidences from entries it wrote itself: dispatches it made, signatures on its own lineage entries, introducing parties noted in its own counterparty identity records. Nothing presented at the moment of assertion establishes the relation. The filed chapter states the consequence in its own terms: presentation of a further identity within a window is incapable of yielding a further increment where the relation is evidenced from records the agent already holds.

Portability runs in opposite directions, and both are deliberate. A credential is portable by design, and portability is central to what the model offers, since a claim verified once can travel to the next verifier. A class identifier is not portable at all. The filing is explicit that it is not an identity attested by a third party, confers no portable standing, and records the class to which a party is assigned rather than any determination concerning that party's conduct.

The two converge in what they refuse. Both decline to place a central authority in the middle of a bilateral interaction. The identity model removes the registrar from identifier control; the filed mechanism removes the registry, directory, consensus procedure, and shared scoring service from the metering computation. Same instinct, different layers.

Both Layers in One Deployment

Decentralized identifiers and credentials establish that an asserting party is who it says it is and, where a trust framework applies, that an issuer has attested attributes about it. The filed mechanism then governs how much any one source may contribute to withholding the agent's execution authorization. Remove either and a gap opens: without some identity layer no signature can be verified at all, and without the metering layer verified signatures multiply without bound.

There is a concrete seam between them. The filing conditions admissibility on a signature verified against an identity primitive recorded in the counterparty identity record of the asserting party. Where the agent holds no such record, the signature is verified against a provisional identity primitive constituted from the presented identity material, and the artifact is held in the pre-settlement inert state rather than passed to the admission evaluator, until a counterparty identity record is promoted upon a matched-pair settlement. The filed chapter recites an identity primitive and does not prescribe its form.

Policy parameters carry the deployment judgment. The enumeration of declared relation types fixes what evidence of common origin the agent can see at all, and the filing requires at least one: an empty enumeration assigns every asserting party to a distinct class and restores the condition described above.

Five limits bound the claim, in both directions.

- **The mechanism authenticates no one.** It consumes a verified signature rather than producing one, and evaluates neither issuer trust nor credential status.
- **Unrelated identities are not drawn into one class.** A party presenting identities that evidence none of the declared relation types has each of them assigned to a distinct class, and each contributes an increment. The cost multiplier reduces a contribution without eliminating it, and applies only where the class is designated an untested class and its recorded introduction paths converge.
- **No merit is determined.** No step evaluates whether a refusal was well founded, and no party adjudicates.
- **Nothing portable is produced.** The partition is local and never reconciled, so no other agent inherits it as a ready-made defense. That is a commitment, not an omission.
- **The identity model is not asked to change.** Nothing here requires an identifier scheme to constrain how many identifiers a subject holds, which would cut against the correlation resistance that scheme exists to provide.

One layer makes an assertion attributable. The other bounds what an origin can do with attribution once it has it.

Disclosure Scope

The mechanism described here is disclosed in U.S. Provisional Application No. 64/117,812, Chapter 2, "Origin-Equivalence Normalization," principally at Sections 2.1 through 2.6, together with the admissibility procedure of Chapter 1 and the merit-independent metering structure of Chapter 3.

Disclosed: incrementing a refusal counter per origin-equivalence class rather than per conduct evaluation artifact; derivation of such a class by an ordered procedure over declared relation types read from a signed policy object, including shared dispatch lineage, co-signature, and common introduction path, each independently sufficient; persistence of the assignment beyond the window with the per-class increment register alone reset; class merge with reconciliation of that register; designation of an untested class by the severance-survival test; the introduction-path convergence test at a declared depth; and a cost multiplier greater than zero and less than unity accumulated in a weighted contribution register.

Disclaimed: any requirement for a centralized registry, directory, consensus procedure, or coordination with a further execution node; any reconciliation of class partitions between agents; any transmission or admission of a class identifier derived by another agent; any portable standing conferred on an asserting party; and any adjudication of the merits of an assertion.

References to W3C Decentralized Identifiers and Verifiable Credentials are to public materials and are used for comparison only; no relationship, endorsement, or infringement is asserted. The description of those specifications is external context and

forms no part of the filing. The application is pending. Nothing here asserts that any particular system practices the disclosed subject matter, and nothing here states that any party requires a license.

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)

APPLICATIONS · GENERAL

- Sybil-Resistant Agent Networks: When One Accuser Becomes Many (/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)

APPLICATIONS · SPECIFIC

- **W3C Decentralized Identifiers and Verifiable Credentials: Verifying an Identifier Versus Metering Its Origin (/articles/origin-equivalence/w3c-decentralized-identifiers)**

TERMINOLOGY

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

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