

Value-Scope Resolution: Testing Whether a Declared Value Is Even Implicated

An agent that admits assertions about its own conduct meets a prior question before it reaches the evidence: does the asserted conduct implicate anything its principal declared? Value-scope resolution, disclosed in U.S. Provisional Application No. 64/117,812, settles that question before the record is consulted. Each declared value of the signed policy object carries a value-scope tuple: a set of action-class identifiers, a set of scope-partition identifiers, and an empathy-scope designation selected from a personal, an interpersonal, and a global scope. Before testing whether the asserted conduct occurred, the admission evaluator extracts a conduct descriptor from the conduct evaluation artifact and applies a three-part membership test against each tuple. A declared value is implicated only where all three parts hold. Where none is implicated, the evaluator produces the not-applicable determination at no cost, and the appended entry carries the tuples tested and the failing element of each.

1. The Gap

The filed background states the need plainly: no existing system admits at runtime an assertion by an external party concerning a persistent agent's own conduct, tests it against an append-only record of the agent's executed actions, and returns a

determination distinguishing contradiction of the assertion from inability to resolve it and from inapplicability to any value the agent has declared.

Those findings are not interchangeable. "My principal never undertook this" speaks to the declared value set; "the record contradicts this" speaks to entries. The filed architecture keeps them apart: the rejected determination (124) issues only where a retrieved entry affirmatively contradicts the conduct evaluation artifact (116) on a recorded field, while the not-applicable determination (128) issues where the conduct descriptor implicates no declared value at all.

Then there is the question of who fixes the boundary. An asserting party characterizing conduct in its own terms would otherwise select the values it is adjudicated against, and a responding agent free to revise its declared values after receipt would select them too. The filing closes both directions.

2. Mechanism

What the tuple holds. The signed policy object (112) is authored by a principal to which the semantic agent (100) is bound and covered by that principal's signature; the agent does not author it and modifies it only by admitting a successor. Each declared value of its declared value set carries a value-scope tuple: a set of action-class identifiers, a set of scope-partition identifiers, and an empathy-scope designation selected from a personal, an interpersonal, and a global scope. Each also carries a declared-value identifier assigned by the principal, covered by the same signature, and stable across successors admitted under one canonical alias.

Inside the artifact. The conduct evaluation artifact (116) originates from an asserting party (118) other than the principal and asserts an evaluation of conduct performed by the semantic agent (100) itself, as distinguished from an output, a task result, or a third party. Its conduct descriptor comprises an action-class identifier, a scope-partition identifier, and an affected-party class, identifying conduct by structural elements

recorded in the append-only lineage field (104) rather than by natural-language characterization. No inference of the asserting party's state, intent, or affect is drawn from the content, and no language model or sentiment classifier is applied to it.

Applicability before evidence. The admission evaluator (120) performs the value-scope test first. For each declared value, the descriptor's action-class identifier is tested for membership in the tuple's action-class set, its scope-partition identifier for membership in the tuple's scope-partition set, and its affected-party class for mapping to the tuple's empathy-scope designation, a counterparty identified in a counterparty identity record (114) mapping to interpersonal, the agent itself to personal, and a party neither identified nor recorded to global. The value is implicated where all three parts hold, and not otherwise. Where the descriptor implicates no declared value, the evaluator produces the not-applicable determination (128) and the procedure terminates. Entries bearing on the conduct are retrieved only where at least one value is implicated.

Who controls the third part. Affected-party class is resolved from counterparty identity records (114) held by the agent, not taken from the artifact, so an asserting party (118) is incapable of selecting the empathy-scope designation against which its assertion is tested. That class is no input to entry retrieval either, so the asserting party cannot narrow the entries retrieved against its own assertion.

Cost and state. The not-applicable determination (128) is produced at no cost. No increment reaches the refusal counter (304), and its run accumulator is not reset. State movement is conditioned on the accepted determination (122), upon which a state modifier modifies the scoped integrity vector (106) and the self-esteem aggregate (108). Metering elsewhere stays live: an issuing agent takes the decrement to its authorization budget (404) at issuance, one and the same for any resolution, and a receiving agent's budget is replenished by one and the same amount for an artifact resolved to any of the four determination classes, upon exposure to an uncounted origin-equivalence class (200) and nothing else.

The record it leaves. Each determination is appended to the lineage field (104) with the artifact, the entries retrieved, the class produced, and the ground of the determination. For the not-applicable determination (128) that ground is the tuples tested and the failing element of each.

Narrowing backward. An assertion is resolved against the signed policy object (112) in force at the recorded assertion time, the policy reference field (110) resolving by canonical alias to candidates then filtered by validity window, by revocation state as of that time, and by an anti-rollback floor held in the memory field (102). A successor admitted after the assertion time bears a later validity window and is not retained, whereby the agent is incapable of narrowing its declared value set after receipt so as to render that artifact inapplicable.

Narrowing forward. From the lineage field the architecture detects an admission of a successor signed policy object (112) that excludes a declared value, or that drops an action-class, scope-partition, or empathy-scope element from a tuple corresponding by declared-value identifier, where the admission follows an as-yet-undetermined artifact implicating that value or element and occurs within a window declared in the policy object in force at the artifact's assertion time rather than in the successor. Responsive to detection the architecture scores a deflection event and increments the refusal counter (304), without adjudicating whether the contraction was well founded. A successor that adds a value or element is scored as no deflection and given no retroactive effect: contraction under a pending assertion is priced, broadening is neither priced nor retroactive. Exclusion and contraction are found by identifier correspondence, not by comparing full value sets.

3. Operating Parameters

Structure and closure are what the filing declares; magnitudes are not. Each bound below is declared in the signed policy object (112) in force.

- Empathy-scope designation is selected from three scopes: personal, interpersonal, and global.
- The membership test is conjunctive; satisfaction of fewer than three parts does not implicate the value.
- The determination set is closed at four members, and the admission evaluator (120) produces no scalar confidence, probability, or graded weight in place of a determination.
- Policy resolution retains only a candidate whose validity window encloses the recorded assertion time and that sits at or above the anti-rollback floor.
- The deflection window is declared in the policy object in force at the artifact's assertion time, not in the successor being admitted.
- The run accumulator of the refusal counter (304) is reset by the accepted determination (122) for the action class the conduct descriptor identifies, and by no other event.
- Where the policy object in force at the recorded assertion time cannot be resolved, no determination issues at all: the artifact is appended as pending and a structured inquiry requesting the policy object is generated, the consequence foreclosed being the not-applicable determination (128).
- The authorization gate (300) is maintained per action class and per scope partition, that partition being declared in the signed policy object (112).

4. Composition

Value-scope resolution is the first stage of the closed determination procedure; the record-grounded test is the second. The accepted determination (122) issues where the descriptor implicates at least one declared value and no retrieved entry affirmatively contradicts the artifact, so the state modification, the deviation recomputation, and the write of the authorization gate (300) to the withheld state (310) all sit behind the tuple.

Empathy-scope is reused rather than duplicated. The scoped integrity vector (106) comprises a personal, an interpersonal, and a global integrity component corresponding to the empathy-scope designations of the tuple, each modified responsive to recorded events whose conduct descriptor or affected-party class resolves to its designation. Counterparty resolution for reparation runs the same mapping: where no held record resolves as a member, the outcome is a no-identified-counterparty outcome and the affected party is designated the structurally silent party.

The scope-partition identifier composes with partition containment. The scoped integrity vector (106), self-esteem aggregate (108), refusal counter (304), authorization budget (404), and a per-partition lineage segment are keyed by a partition identifier declared in the signed policy object (112). State movement in a first partition propagates to a second only where the first partition's coupling declaration enumerates that partition and the empathy-scope designation implicated; a partition whose coupling declaration enumerates nothing is fully confined.

Tuples also travel between agents. Two semantic agents (100) resident concurrently at one stateless execution node each compute, under their own anti-rollback floors, the intersection of the two policies' tuple sets for a common action class and confine execution to it; on an empty intersection the agent produces no determination of the closed set, emitting a structured inquiry and appending a mutual abstention record. A delegate receiving an artifact that implicates a value the delegator declares and its own policy does not produces the not-applicable determination (128) and appends an unreconciled-value record, carried in the return record and barred from exclusion or attenuation. A co-signed mutual admission compact enumerates reciprocal tuples between two parties; where a compact tuple and a party's own tuple disagree, the union applies, so a compact cannot subtract from a declared value set.

5. Prior-Art Distinction

Attribute-based access control. XACML-conformant policy decision points already produce a not-applicable result, and the vocabulary is genuinely close. The structure is not. There the result is a target-match outcome returned to the requesting party, producing no persistent state in the decider and constraining nothing it does next. Here it is a determination the agent produces over its own record about an assertion concerning its own past conduct, appended with the tuples tested and the failing element of each, inside a metering structure that gates the agent's execution authorization.

Behavioral integrity and reputation systems. Two issued patents cited in the filed background, U.S. Patent No. 12,563,045 and U.S. Patent No. 12,526,244, compute an integrity or reputation quantity from observed execution against an expected baseline and modify capability or standing on a threshold comparison. The background characterizes both as conformity detectors, divergence from the baseline being itself the fault signal. Neither admits an assertion by an external party concerning the agent's conduct, nor supplies a procedure by which the agent tests such an assertion against its own record. There is accordingly no admitted assertion for an applicability stage to precede.

Ethical governor architectures. The ethical adaptor of Arkin and Ulam and related governors interpose an evaluative component between deliberation and actuation and restrict the permissible action set where a candidate action violates a constraint. The filed background characterizes such governors as credulous, admitting an evaluative input without a record-grounded procedure for testing it, and monotonic, the governed action set only ever contracting. The mechanism here can terminate at the first stage without state movement, and its bars price a contraction of declared scope.

Externally held determinations. Article 23(2) of Regulation (EU) 2022/2065, Section 623(a)(8)(F) of the Fair Credit Reporting Act, and bonded challenge procedures each attach a consequence to assertive conduct only after a merits determination by an

institution other than the penalized party. PeerReview and remote attestation architectures such as RFC 9334 locate the verdict in witness nodes, auditors, or a verifier. Here the determination, inapplicability included, is produced by the agent over its own record, and the deflection scoring that guards the tuple proceeds without adjudicating whether a contraction was well founded.

These are different structures addressed to different problems. Nothing here is a statement about any party's practices.

6. Disclosure Scope

What is disclosed is the subject matter of U.S. Provisional Application No. 64/117,812 at Section 10.1, paragraph [0329], "Value-scope resolution and the value-scope tuple," together with the base-architecture recitations it composes on at [0022] through [0025], [0032], [0033], [0036], [0041] through [0049], [0093], [0106], and [0145]; the declared-value identifier and the narrowing bars at [0330] through [0332]; the graceful-degradation path at [0152]; the compositions at [0231], [0353], [0354], [0416], [0417], [0422], [0469], and [0471]; the affected-party resolution at [0483]; and the prior-art characterizations at [0005] through [0010].

Disclosed specifically: the value-scope tuple; the conjunctive three-part membership test applied before testing whether the asserted conduct occurred; the empathy-scope mapping of an identified counterparty to interpersonal, the agent itself to personal, and an unrecorded party to global; production of the not-applicable determination (128) terminating the procedure at no cost; the appending of the tuples tested and the failing element of each; resolution of the affected-party class from held records; and the narrowing bars with asymmetric deflection scoring keyed by declared-value identifier.

Disclaimed, in the sense of not asserted here: any numeric value for any window, threshold, or magnitude, each being declared in the signed policy object (112) in force; any characterization of any system, standard, regulation, or patent named above as

infringing, and any statement about any party's practices; and any characterization of the applications commonly owned with the filing as prior art to this subject matter.

This article is a defensive technical disclosure describing embodiments as filed. The scope of protection is defined by the claims as they may issue, and not by this description.

Record-Grounded Conduct Admission (</conduct-admission>)

[All 49 steps → \(/inventive-steps\)](/inventive-steps)

Conduct judged against the agent's own record, not an outside authority.

Provisional application

PRIMARY TECHNICAL DISCLOSURE

- [Record-Grounded Conduct Admission: Governing How an Agent Answers an Accusation \(/articles/conduct-admission/record-grounded-conduct-admission\)](/articles/conduct-admission/record-grounded-conduct-admission)

SECONDARY TECHNICAL

- [Value-Scope Resolution: Testing Whether a Declared Value Is Even Implicated \(/articles/conduct-admission/value-scope-resolution\)](/articles/conduct-admission/value-scope-resolution)
- [The Retroactive Narrowing Bar: Why an Agent Cannot Edit Its Values After the Fact \(/articles/conduct-admission/retroactive-narrowing-bar\)](/articles/conduct-admission/retroactive-narrowing-bar)
- [The Delegator Floor: A Monotone Upward Ratchet on Delegated Authority \(/articles/conduct-admission/delegator-floor-ratchet\)](/articles/conduct-admission/delegator-floor-ratchet)
- [The Co-Signed Mutual Admission Compact: Two Agents Binding Each Other \(/articles/conduct-admission/co-signed-mutual-admission-compact\)](/articles/conduct-admission/co-signed-mutual-admission-compact)
- [The Inherited Governance Record: Fork- and Clone-Resistant Accountability \(/articles/conduct-admission/inherited-governance-record\)](/articles/conduct-admission/inherited-governance-record)

TERMINOLOGY

- [append-only lineage field \(/glossary#append-only-lineage-field\)](/glossary#append-only-lineage-field)
- [authorization gate \(/glossary#authorization-gate\)](/glossary#authorization-gate)

- [conduct evaluation artifact \(/glossary#conduct-evaluation-artifact\)](/glossary#conduct-evaluation-artifact)
- [principal \(/glossary#principal\)](/glossary#principal)
- [semantic agent \(/glossary#semantic-agent\)](/glossary#semantic-agent)
- [signed policy object \(/glossary#signed-policy-object\)](/glossary#signed-policy-object)

[Record-Grounded Conduct Admission overview → \(/conduct-admission\)](/conduct-admission)