

Coalition Operations: Standing Without a Shared Authority

A coalition assembles from organizations that answer to different principals, keep separate records, and share no registry that ranks anyone. Automated participants still have to rely on determinations made inside other organizations, and the familiar answers tend to require either a common scoring authority or an unrecorded act of trust. U.S. Provisional Application No. 64/117,812 discloses a different arrangement: one agent executes under a second agent's recorded determination, and the counter-signature that governs the reliance decrements the followed agent's own authorization budget, the same budget that gates its ability to dispatch. No rating, score, standing quantity, or credential passes between them. What the deferring agent obtains is admissibility, the ability to accept correction from parties that previously corrected the agent it followed.

When Nobody Outranks Anybody

Picture a joint task force assembled from several contributing organizations, each bringing people, systems, and automated components that act on their own recognizance inside a narrow scope. Every contributor answers to its own principal,

holds its own records, and hands over neither its internal evaluations nor authority over its systems. There is a liaison structure, no single commander over all of them, and no registry any of them will treat as the arbiter of who is competent at what.

The operational problem appears the first time a component in one organization must act on a determination reached inside another. A partner's system has already evaluated a piece of conduct, and its evaluation is the only one available inside the decision window. Acting on it means relying on a judgment produced under someone else's policy, by a system you cannot inspect, on records you cannot retrieve. Situations like this are commonly resolved by people: a liaison vouches, someone senior accepts the risk verbally, and the reliance leaves no durable machine-readable trace. When the operation dissolves, so does much of what was learned about who could be relied upon for what.

Federated Trust Models Stall at the Organizational Seam

The first instinct is usually to build the missing registry: publish a directory of participants with attributes, competencies, and assurance levels, and let systems consult it. The difficulty there is governance rather than engineering. Whoever operates the registry becomes the authority the coalition was formed without, and a published score becomes something to manage rather than something to measure.

Identity federation answers a narrower question, whether an identity asserted in one domain can be recognized in another, which says little about whether a determination reached under a foreign policy should be relied upon.

Capability delegation architectures come closer and then invert the incentive. Authority passes downward and is attenuated on the way, so the delegate receives a narrower capability than the delegator held and the delegator is protected. As the filing characterizes such designs, the cost falls upon the relying party and the party relied

upon is protected. Applied to a coalition, a contributor whose determinations become widely adopted would absorb no consequence from that adoption, while the organizations relying on it absorb it.

Closing that gap with a reputation quantity, whether a score, a standing, or a stake, reintroduces the original problem. The quantity has to live somewhere and be computed by someone, it invites optimization by the parties it measures, and it amounts to the shared authority the coalition could not agree to create.

The Disclosed Mechanism: Deference That Costs the Followed Agent

Chapter 9 of the filing discloses a mechanism by which one semantic agent (100) executes under the recorded evaluative determination of a second semantic agent (100), with no rating, score, standing quantity, or credential passing between them. Only an agent already authorized to decide the matter for itself can defer: its dispatch-authority predicate recomputes as satisfied from its own memory field (102), and it elects to execute under the followed agent's determination rather than one its own admission evaluator (120) would produce. An agent whose predicate fails, whose authorization gate (300) stands in the withheld state (310), or which has entered the non-executing cognitive mode (302) does not defer; its holding is a withholding.

Before any deference takes effect the deferring agent runs an origin-disjointness test on its own records. If it and the followed agent are assignable to a common origin-equivalence class (200), the record is appended as an attempted deference and given no effect: nothing transmitted, nothing decremented, nothing enrolled, nothing adverse appended. It then performs a retrievability-form verification of three ordered conjuncts on the followed determination record: the presented authority credential and continuity hash field against its counterparty identity record (114), the tested-entry enumeration

against a lineage commitment over the followed agent's ordered lineage entries, and the determination as a member of a closed set defined earlier in the filing, carrying a matching action class. That verification is of form and retrievability rather than merits.

On both tests passing it appends a deference record (900) carrying the followed determination record, the tested-entry enumeration unaltered, a deference abstention entry naming what was withheld from its own admission evaluator (120), and a counter-signature window field (912) declared in its own signed policy object (112). Electing to counter-sign, the followed agent appends a deference counter-signature (902) and by that append decrements its own authorization budget (404), the budget gating its own non-advisory dispatch, under a deference decrement schedule declared in its signed policy object (112). That schedule declares an amount not less than the amount declared for that agent's issuance of a conduct evaluation artifact (116), monotonically non-decreasing in the count of counter-signatures it has appended within a declared window, and bounded above by a declared maximum. A step form and a form linear in the count are both disclosed, the applicable one declared in policy, so that the marginal cost of being followed rises with the rate of adoption.

That decrement happens without regard to merits: no adjudication of whether the deference was well founded, or the reliance reasonable, is performed by either agent or any other party as a condition of it. Responsive to the budget satisfying a floor declared in the followed agent's own policy object, its authorization gate (300) is written to the withheld state (310) for an enumerated set of action classes and an escalation record is emitted to its principal. The budget replenishes by the general procedure disclosed elsewhere in the filing and by no other path, so being followed spends a replenishable quantity rather than terminating the agent.

What the deferring agent receives is not standing. Responsive to receipt of the counter-signature, and not otherwise, it retrieves a corrector-class enumeration (904): the origin-equivalence classes (200) that both corrected the followed agent in that action class and whose corrections that agent thereafter conformed to, admitted only where

both parts of a two-part conformance filter (906) hold. Correction means an accepted determination (122) the followed agent produced upon an artifact from that class; conformity means a later conduct record of that agent, within an observation interval, recorded as compliant with that determination. A class satisfying the first part and not the second is excluded, and each admitted class is carried by a class-defining reference. Those classes enter the deferring agent's replenishment register (402) as admissible-but-unexercised entries with two effects and no third: an artifact from a party assignable to the class becomes admissible to its admission evaluator (120) with no prior encounter history, and while unexercised the entry is not counted in renewal.

What Changes Inside a Joint Operation

Return to the task force with the mechanism in place. A component in Organization A holds a pending dispatch and is authorized to decide it; Organization C's component has already recorded a determination on the same conduct descriptor. A's component runs origin-disjointness first, and where A and C are separate contributors sharing no origin-equivalence class (200) on A's records, the test passes. Deference to a component assignable to a class A shares would instead be given no effect, so deference takes effect only across an origin boundary, which is the boundary a coalition is assembled from.

A's component then appends a deference record (900) and transmits it for counter-signature. C faces a real decision rather than a courtesy. Counter-signing decrements its authorization budget (404) in the units that gate its own dispatch, and under the disclosed schedule the amount is non-decreasing in the count of counter-signatures it has already appended within the declared window. A contributor whose determinations many parties adopt would feel that adoption in its own capacity to act.

C may decline, expressly or by letting the carried window lapse. A declination is not a denial of A: it is recorded as a not-determinable determination (126) naming the counter-signature as an unavailable input, increments no counter of either agent, and is

treated at no stage as a rejected determination (124) or an accepted determination (122). A may still execute, appending an ungoverned reliance record naming the deference record (900), the absence of the counter-signature, the window applied, and the execution performed, together with a reliance-bound narrowing record that narrows its own counterparty scope object for C by the categories its policy declares. That narrowing is not a sanction and is released upon a counter-signature within the reinstatement window. The ungoverned reliance sits on A's own record.

Where C does counter-sign, what A gains is worth stating precisely. No rating, standing quantity, or transferable credential passes to A. What A enrolls is C's conformance-filtered corrector classes, written into its replenishment register (402). To each enumerated class A dispatches its own execution records in that action class, in the categories of lineage record disclosure its counterparty scope object admits, with a declaration that it admits conduct evaluation artifacts (116) from the class. Nothing transmitted describes C's standing, and A receives no rating from those parties.

That enrollment is where the workflow diverges from ordinary reliance, in which adopting a partner's judgment transfers the judgment alone and leaves the correction channel inside the partner. Under the disclosed arrangement, when one of those correcting parties later sends A an artifact, A appends an exercised transition record (908) binding the arriving correction to the exposure edge record (910), the deference record (900), the counter-signature (902), C's identity, and the determination A originally adopted. Only upon that arrival, and where the class is absent from the register for the declared window, does A's own authorization budget (404) replenish, without regard to whether the artifact was favorable to A. Enrollment by itself renews nothing; the exposure has to be exercised.

Review after the fact is served by a further record. A deference correction linkage record, keyed on the adopted determination, lets A enumerate the borrowed judgment, its source, the counter-signature that governed the reliance, and each correction that reached A through the exposure that deference enrolled. That record is subject to the

conversion bar (502): no stage converts it, or a count of such records held for a followed agent, into a scalar value, a rating, or an operand of a threshold comparison, and no admission or dispatch decision consumes such a count as a standing quantity.

Chained reliance is addressed too. Where a third participant defers to A upon a determination A itself adopted, A presents a chained deference record naming the originating followed agent, a chain depth one greater, and every agent already in the chain; upon counter-signature A decrements its own budget, and the originating agent is not decremented a second time. Should two partners present determinations of differing outcome class on one conduct descriptor, the deferring agent appends a deference conflict record, adopts neither, and either decides the matter itself or withholds as a positive abstention.

Fielding It, and What It Does Not Solve

Integration is per-participant. Each contributor runs its own agent, holds its own records, and declares its own signed policy object (112). No coalition-wide service has to be stood up, and no onboarding step assigns a new participant a level. The origin-disjointness test, the ordered assignability test applied to arriving artifacts, and the exercised transition are each evaluated on the deferring agent's own records, expressly without reference to a centralized registry and without coordination with a further execution node.

Most of the quantities are left to policy. The base decrement amount, the multiple or increment in the schedule, the declared maximum, the budget floor, the counter-signature and reinstatement windows, the observation interval, and the maximum chain depth are each declared in a participant's own signed policy object; the filing gives figures only inside an illustrative embodiment. A coalition adopting the arrangement has calibration work ahead of it; the mechanism supplies a vocabulary for that negotiation, not an answer to it.

Several limits are deliberate. The architecture does not tell a deferring agent whether a determination is correct, since verification is of form and retrievability only and a well-formed determination drawn from a verifiable lineage can still be wrong; what it binds is the consequence of borrowing a judgment. No participant obtains a coalition-wide picture, only its own edges. Nothing compels a participant to counter-sign, and a refusal narrows scope on the deferring side alone. The chapter supplies no identity infrastructure and does not speak to classification or releasability policy beyond respecting the disclosure categories a counterparty scope object already admits. Separate chapters of the filing govern the origin-equivalence class, the authorization budget and its replenishment, the conversion bar, and the matched-pair transmission form.

Disclosure Scope

This article describes subject matter disclosed in Chapter 9, "Deference and Earned Exposure," of U.S. Provisional Application No. 64/117,812, part of the Adaptive Query portfolio, and is published as a general-applications description and a timestamped public disclosure. Architectural statements above are drawn from that chapter alone; mechanisms governed by other chapters of the same filing are referenced only by name. The application is pending, and nothing here asserts that any system infringes or that any license is required. The coalition framing illustrates where the mechanism applies and describes no deployment.

Deference and Earned Exposure ([/earn](#) [All 49 steps](#) → [\(/inventive-steps\)](#) [ed-exposure](#))

Standing earned through recorded conduct; deference that survives across agents.

Provisional application

PRIMARY TECHNICAL DISCLOSURE

- [Earned Exposure and Deference Records: Standing Built From Recorded Conduct \(/articles/earned-exposure/earned-exposure-and-deference-records\)](/articles/earned-exposure/earned-exposure-and-deference-records)

SECONDARY TECHNICAL

- [The Valuation-Free Ordinal Standing Quantity \(/articles/earned-exposure/valuation-free-ordinal-standing\)](/articles/earned-exposure/valuation-free-ordinal-standing)
- [Determination-Invariance Under an Agent's Own Policy Succession \(/articles/earned-exposure/determination-invariance-policy-succession\)](/articles/earned-exposure/determination-invariance-policy-succession)
- [Elective Exposure Enrollment Without Deference \(/articles/earned-exposure/elective-exposure-without-deference\)](/articles/earned-exposure/elective-exposure-without-deference)
- [Interventional Attribution by Receipt-Path Control \(/articles/earned-exposure/interventional-attribution-receipt-path\)](/articles/earned-exposure/interventional-attribution-receipt-path)
- [The Portable Exoneration Record \(/articles/earned-exposure/portable-exoneration-record\)](/articles/earned-exposure/portable-exoneration-record)
- [Attribution, Exoneration, and the Domain Baseline: An Ordered Branch Over the Latency Differential \(/articles/earned-exposure/attribution-versus-exoneration-branch-domain-baseline\)](/articles/earned-exposure/attribution-versus-exoneration-branch-domain-baseline)
- [Bounded Count of Dispatches Per Deference Record: Capping How Many Executions One Counter-Signed Adopted Determination Authorizes \(/articles/earned-exposure/bounded-count-dispatches-deference-record\)](/articles/earned-exposure/bounded-count-dispatches-deference-record)
- [Coarse Lineage-Segment Granularity: Verifying a Tested-Entry Enumeration Without Disclosing Which Lineage Entries Were Tested \(/articles/earned-exposure/coarse-lineage-segment-granularity-tested-entry\)](/articles/earned-exposure/coarse-lineage-segment-granularity-tested-entry)
- [Correction-Exposure Decay: An Authorization Quantity That Depletes With Agent Operation and Renews Only on Recorded External Evaluation \(/articles/earned-exposure/correction-exposure-decay-authorization-quantity\)](/articles/earned-exposure/correction-exposure-decay-authorization-quantity)
- [Cause-Keyed Gate Restoration: A Distinct Return Path for Decay-Induced Withholding of Agent Authorization \(/articles/earned-exposure/distinct-return-path-decay-induced-withholding\)](/articles/earned-exposure/distinct-return-path-decay-induced-withholding)
- [Edge Severance Events and the Three-Valued Survival State: Resolving Whether a Counterparty Relationship Outlived a Payment-Class or Obligation-Class Termination \(/articles/earned-exposure/edge-severance-events-three-valued-survival\)](/articles/earned-exposure/edge-severance-events-three-valued-survival)
- [Emitting the Deference Correction Linkage Record to the Followed Agent: Reporting a Downstream Correction Without Producing a Rating \(/articles/earned-exposure/emission-deference-correction-linkage-record-followed\)](/articles/earned-exposure/emission-deference-correction-linkage-record-followed)
- [The No-Address Omission Branch: What a Deferring Agent Records When an Exposure Dispatch Reaches Nobody \(/articles/earned-exposure/exposure-dispatch-no-address-omission-branch\)](/articles/earned-exposure/exposure-dispatch-no-address-omission-branch)
- [Channel-Baseline Normalization: Taking Channel Overhead Out of the Latency Differential \(/articles/earned-exposure/latency-differential-channel-baseline-normalization\)](/articles/earned-exposure/latency-differential-channel-baseline-normalization)

- [Reason-Type Resolution From Severance-Survival States: Typing an Agent's Relationship Edges Without Asking the Counterparty](/articles/earned-exposure/reason-type-resolution-survival-states) (/articles/earned-exposure/reason-type-resolution-survival-states).
- [Reason-Type Retyping by a Superseding-Type Entry: Recomputing an Edge's Reason-Type on Later Severance Evidence Without Disturbing Modifications Already Applied](/articles/earned-exposure/reason-type-retyping-superseding-type-entry) (/articles/earned-exposure/reason-type-retyping-superseding-type-entry).
- [Renewal-Side Replenishment of the Authorization Budget: Restoring an Agent's Metered Spending Capacity Only on Exposure to an Uncounted Origin-Equivalence Class](/articles/earned-exposure/renewal-side-replenishment-authorization-budget) (/articles/earned-exposure/renewal-side-replenishment-authorization-budget).
- [Following Does Not Count as Being Tested: Retention of the Short Decay Constant When an Agent Defers, and the Action Class That Stays Untested](/articles/earned-exposure/retention-short-decay-constant-deference-class) (/articles/earned-exposure/retention-short-decay-constant-deference-class).
- [Standing-Determination Designation: Substituting a Verified Counterparty Determination for One Conjunct of the Dispatch-Authority Predicate](/articles/earned-exposure/standing-determination-designation-substitution-inside-dispatch) (/articles/earned-exposure/standing-determination-designation-substitution-inside-dispatch).
- [Uncounted-Class Renewal: Gating Agent Authorization on a Windowed Register of Origin-Equivalence Classes](/articles/earned-exposure/uncounted-class-renewal-gated-windowed-register) (/articles/earned-exposure/uncounted-class-renewal-gated-windowed-register).

APPLICATIONS • GENERAL

- [Vendor Qualification Without Ratings: Standing Earned From Recorded Conduct](/articles/earned-exposure/vendor-qualification-without-ratings) (/articles/earned-exposure/vendor-qualification-without-ratings).
- [Credentialing Autonomous Practitioners: Deference That Survives Across Agents](/articles/earned-exposure/professional-credentialing-agents) (/articles/earned-exposure/professional-credentialing-agents).
- [Coalition Operations: Standing Without a Shared Authority](/articles/earned-exposure/coalition-operations) (/articles/earned-exposure/coalition-operations).

APPLICATIONS • SPECIFIC

- [Marketplace Seller Reputation Systems: Presented Scores Versus Earned Admissibility](/articles/earned-exposure/marketplace-seller-reputation) (/articles/earned-exposure/marketplace-seller-reputation).
- [Letta: Agent Memory and Standing Earned From Conduct](/articles/earned-exposure/letta-agent-memory) (/articles/earned-exposure/letta-agent-memory).
- [Trustpilot Signals and Conduct-Derived Admissibility](/articles/earned-exposure/trustpilot-signals) (/articles/earned-exposure/trustpilot-signals).
- [Copilot Coding Agents: Review Gates and Earned Exposure](/articles/earned-exposure/github-copilot-agent) (/articles/earned-exposure/github-copilot-agent).
- [Supplier Risk Scores and Recorded Counterparty Conduct](/articles/earned-exposure/supplier-risk-scoring) (/articles/earned-exposure/supplier-risk-scoring).
- [Verified Profiles and Standing Built From Recorded Conduct](/articles/earned-exposure/linkedin-verification) (/articles/earned-exposure/linkedin-verification).

TERMINOLOGY

- [counterparty identity record \(/glossary#counterparty-identity-record\)](#)
- [deference record \(/glossary#deference-record\)](#)
- [edge \(/glossary#edge\)](#)
- [reason-type \(/glossary#reason-type\)](#)
- [severance event \(/glossary#severance-event\)](#)
- [standing quantity \(/glossary#standing-quantity\)](#)
- [untested class \(/glossary#untested-class\)](#)

[Deference and Earned Exposure overview → \(/earned-exposure\)](#)