Nick Clark · Published July 29, 2026 · PDF

Gig Platform Deactivations: Answering an Accusation From the Record

A worker or seller on a marketplace can lose an income stream to a complaint they never see in full, filed by a counterparty they cannot question, decided in a queue they cannot inspect. The accused party holds a record of what actually happened and has no structured way to put it against the assertion. U.S. Provisional Application No. 64/117,812 discloses a different arrangement: a signed conduct evaluation artifact arrives from a party other than the principal, over the agent's ordinary interaction channel, and the agent resolves it against its own append-only lineage field and the signed policy object in force into exactly one of four determination classes. No registry, arbitrator, scoring service, or platform sits in that loop, and the ground of each determination is recorded with it.

When a Complaint Quietly Ends an Income Stream

A courier finishes a route and later finds the account restricted. A seller finds a listing frozen, alongside a note saying a customer reported a problem. Each learns that someone said something, that a decision followed, and that there is a form to fill out. Neither learns which specific action was implicated, what the complainant actually asserted, or what the decision rested on.

The asymmetry here is not that platforms decide badly. It is that the accused party holds most of the relevant facts and has no channel that will accept them in a form the decision process can consume. The courier knows the handoff, the message thread, the condition of the item at dispatch, but a free-text appeal box reduces all of it to narrative competing with the complainant's narrative.

Cost is the second asymmetry. Filing is cheap and repeatable for whoever files, answering is not, and a first report from a given source typically looks no different from a repeated one. Portability is the third, at least where these systems are described publicly: a restriction imposed inside one marketplace tends to be a fact about the worker that the worker cannot carry, contest, or annotate elsewhere.

Why Better Appeal Systems Do Not Close the Gap

Proposals in this area tend to reach for a better adjudicator: an independent review body, a dispute registry, a model trained to read appeals more evenhandedly. Each of those adds a third party to a two-party problem and inherits the same defect, since the adjudicator was not present and holds no record of the conduct.

Free-text appeals compound that. When the input is narrative, the processing has to be interpretive, and interpretive processing yields a graded output: a confidence, a risk band, a score. A graded output cannot be contested on a specific point, because there is no specific point, and the only response available is more narrative.

Underneath both sits a quieter problem. Silence tends to work against the accused, so an appeal fails wherever the worker cannot produce evidence, and the failure looks the same whether the worker had nothing to say or no admissible way to say it.

How an Assertion Resolves Against the Agent's Own Record

The disclosed architecture puts a persistent semantic agent (100) on the side of the accused party. That agent carries a memory field (102), an append-only lineage field (104) recording executed actions and determinations, a scoped integrity vector (106), a self-esteem aggregate (108), and a policy reference field (110) resolving to a signed policy object (112). It also holds a counterparty identity record (114) per counterparty, carrying that party's identity primitive and the states it has attested.

The signed policy object (112) is authored by the principal to which the agent is bound, covered by that principal's signature, and carries a declared value set, a scope partition enumeration, and the declared bounds, weights, and coefficients the computations consume. The agent does not author it and modifies it only by admitting a successor. Which object applies is fixed by an ordered procedure over the policy reference field (110), which includes an anti-rollback monotonicity constraint so a superseded object is not reinstated.

An accusation arrives as a conduct evaluation artifact (116) originating from an asserting party (118) other than the principal, asserting an evaluation of conduct performed by the semantic agent (100) itself, as distinguished from an evaluation of an output, a task result, or a third party. It arrives over the agent's ordinary interaction channel, so no dedicated evaluation channel or centralized intake service is required and no registry of evaluations is referenced. The artifact carries an asserting-party identifier, a recorded assertion time, a conduct descriptor comprising an action-class identifier, a scope-partition identifier, and an affected-party class, a signature of the asserting party (118), and an attested assertion-cost counter (400) state with an epoch identifier of that party's dynamic agent hash chain. The descriptor identifies conduct by structural elements recorded in the lineage field, not by natural-language characterization, and the architecture performs no inference of the asserting party's intent and no classification of artifact content by a language model or sentiment classifier.

Admitted artifacts pass to the admission evaluator (120), which produces exactly one determination from a closed set: an accepted determination (122), a rejected determination (124), a not-determinable determination (126), or a not-applicable determination (128). It produces no scalar confidence, probability, or graded weight in place of a determination, and the determination is produced by the semantic agent (100) over its own lineage field against the declared value set of the policy in force at the recorded assertion time, not by an external authority, arbitrator, registry, or scoring service.

Walking a Restriction Through the Determination Path

Return to the courier, now operating with such an agent, and suppose a customer disputes a handoff. Instead of free text routed to a queue, the counterparty signs a conduct evaluation artifact (116) carrying that descriptor triple and attesting its own assertion-cost counter (400) state and hash chain epoch.

Admissibility runs first. The signature is verified against an identity primitive recorded in the counterparty identity record (114), and an artifact whose signature does not verify is appended to the lineage field and not admitted. An artifact whose attested counter state is absent, whose epoch identifier is not a valid successor of the epoch recorded for that party, or whose counter state is below a state that party previously attested, is likewise appended and not admitted. A non-admitted artifact produces no determination, moves no value of the scoped integrity vector (106), and increments no counter. Where the agent holds no record for the asserting party, the signature is verified against a provisional identity primitive constituted from the presented material, and the artifact waits in the pre-settlement inert state pending a promotion governed by a separate chapter. An unsigned or stale-epoch assertion fails on verification rather than on a moderator's judgment.

For an admitted artifact, the evaluator runs a value-scope test against each declared value: the action-class identifier is tested for membership in the value-scope tuple's action-class set, the scope-partition identifier for membership in its scope-partition set, and the affected-party class for mapping to its empathy-scope designation, which is personal, interpersonal, or global depending on whether the affected party is the agent itself, a recorded counterparty, or a party neither identified nor recorded. A declared value is implicated only where all three parts are satisfied. Where the descriptor implicates none, the result is the not-applicable determination (128), with the tuples tested and the failing element of each recorded as the ground. Conduct the principal's policy does not govern is disposed of on the record.

Where at least one declared value is implicated, the evaluator retrieves entries bearing on the identified conduct, an entry bearing where it records an action of that action-class identifier within that scope partition. One detail of that retrieval is worth naming: the affected-party class is not an input to the retrieval and is an input to the value-scope test alone, so an asserting party is incapable of narrowing the set of entries retrieved against its own assertion by the class it declares.

The rejected determination (124) is produced only where a retrieved entry affirmatively contradicts the artifact on a recorded field, a contradiction being affirmative where that field holds a value inconsistent with the conduct asserted, and not affirmative where the entry is silent as to that field. This is the courier's actual answer, and it is structural: the appended entry carries an identifier of the contradicting entry and the field on which the contradiction was found.

Where the record holds no bearing entry, an entry that does not resolve the descriptor, or an entry whose fields are incomplete, the result is the not-determinable determination (126). Absence of evidence resolves to this class and no other, so the agent is incapable of refusing a conduct evaluation artifact (116) by reason of its own record being silent, incomplete, or unavailable. Thin records cannot make an accusation vanish, and the gap is itself recorded.

The accepted determination (122) follows where at least one declared value is implicated and no retrieved entry affirmatively contradicts the artifact. What follows runs on the agent's own side. A state modifier modifies the scoped integrity vector (106) and the self-esteem aggregate (108), the magnitude of each modification being scaled by an entropy-weighted harm coefficient computed under a separate section. A deviation engine then recomputes the deviation likelihood (706) as the quotient whose numerator is the difference between a need quantity (700) and a dynamic ethical threshold (702) resolved for the implicated integrity scope, and whose denominator is the product of an empathy weighting (704) and the self-esteem aggregate (108); those operands, and the bound against which the quotient is compared, are disclosed in a separate section. Where the recomputed quantity satisfies the policy-declared bound, the authorization gate (300) transitions to the withheld state (310) for the affected action class and the agent enters the non-executing cognitive mode (302). Nothing in that path makes withholding automatic on an accepted determination; the bound has to be satisfied first.

A rejected determination is not costless either. Each one increments a refusal counter (304), and that counter writes the authorization gate (300) upon satisfaction of a threshold recited in a separate section.

Two consequences follow. First, restriction is state the principal's own policy produces over the principal's own record: the authorization gate (300) is a persistent element of the agent's memory field, not state held by a counterparty, and an escalation emitter emits a record to the principal upon a transition to the withheld state (310), enumerating the action class, the scope partition, and the entries on which the transition was computed. Second, lifting a restriction is procedural. A restoration controller returns the gate toward the granting state only upon a procedure appended to the lineage field, and no elapse of time, and no payment, transfer, or consideration by any counterparty, returns it.

Dispatch is handled with the same concreteness. The agent recomputes a dispatch-authority predicate on each request to dispatch an action, from state then carried in the memory field (102), and not from an authorization token, a cached result, or a session grant, so a gate write or a lineage append takes effect at the next dispatch request without revocation infrastructure. Satisfaction requires a three-stamp conjunction in full, and fewer than three stamps produces a deterministic denial that is appended naming the stamp that did not resolve and is not converted into a determination concerning any party.

Integration Realities and What Stays Unsolved

Adoption does not require a platform to move first. Intake runs over the agent's ordinary interaction channel, so an operator can stand up the artifact format and the determination path between its own agent and willing counterparties, with no shared registry and no coordination with a further execution node.

The limits are real. A counterparty who will not sign produces no admissible artifact, and non-admission is a disposition rather than a resolution, so the underlying grievance stays outside the system. Nothing here compels a marketplace to honor a determination the worker's agent produced, or displaces an operator's right to restrict access on its own terms. The filing describes governance of an agent bound to a principal, not a claim on anyone else's enforcement.

Quality of determination is bounded by quality of record. Because absence of evidence resolves to the not-determinable determination (126), an agent with a thin lineage field accumulates that class rather than contradictions, and repetition of that class against one asserted conduct event carries consequences under a separate section. Authoring the policy object is its own burden: declared values, value-scope tuples, action classes, scope partitions, bounds, and coefficients all have to be signed by the principal first. Counterparty identity has to be established too, since the pre-settlement inert state means an assertion from a stranger waits rather than resolving.

Note finally how narrow the decision is. It resolves whether an assertion about the agent's own conduct is contradicted, unresolvable, out of scope, or uncontradicted by the record. It does not determine truth, assign fault between parties, or compute damages, and it produces no score. That narrowness is what makes the output contestable on a named field.

Disclosure Scope

The architecture described here is disclosed in Chapter 1, Overview and Base Architecture, of U.S. Provisional Application No. 64/117,812, covering the persistent semantic agent and its carried fields, the signed policy object in force, the authorization gate, the dispatch-authority predicate and three-stamp conjunction, the conduct evaluation artifact, admissibility, the closed set of determination classes, and the recording of determinations and their consequences. Mechanisms referenced above as governed elsewhere in that filing, including the entropy-weighted harm coefficient, the refusal-counter threshold, matched-pair settlement, and assertion-cost counter treatment, are disclosed in their own chapters and are not described here. This article is a general description of an application domain, published as a technical disclosure establishing a timestamped record. It names no company, product, or service, and it makes no assertion that any person or system infringes any claim. The filing referenced is pending.

Record-Grounded Conduct Admission (/conduct-admission)

All 49 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)

SECONDARY TECHNICAL

- Value-Scope Resolution: Testing Whether a Declared Value Is Even Implicated (/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)
- The Delegator Floor: A Monotone Upward Ratchet on Delegated Authority (/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)
- The Inherited Governance Record: Fork- and Clone-Resistant Accountability (/articles/conduct-admission/inherited-governance-record)
- The Corrective Encounter: A Commission-Versus-Omission Compliance Test for Whether an Agent Actually Changed (/articles/conduct-admission/corrective-encounter-commission-omission-compliance-test)
- Continuity-Break Demotion of a Counterparty Identity Record: Reversing First-Encounter Promotion Without a Conduct Determination (/articles/conduct-admission/demotion-persistent-tier-continuity-break)
- Ratchet-Freeze on Uncomplied Correction: Why One Conforming Act Cannot Restore an Agent's Earned Authorization Persistence (/articles/conduct-admission/ratchet-freeze-uncomplied-correction)
- Three-Way Corroboration Classification: Corroborating, Contradicting, or Neither, and the Failed-Corroboration Reject Trigger (/articles/conduct-admission/three-way-corroboration-classification-reject-trigger)
- Untested-Class Authorization Decay: Metering Agent Permission Persistence by Recorded Corrective Encounters (/articles/conduct-admission/untested-class-authorization-decay)

APPLICATIONS · GENERAL

- Agent Marketplaces: Handling a Complaint Against an Autonomous Seller (/articles/conduct-admission/agent-marketplace-dispute-handling)
- Clinical AI Agents: Answering an Accusation Without a Central Registry (/articles/conduct-admission/clinical-agent-accountability)
- **Gig Platform Deactivations: Answering an Accusation From the Record (/articles/conduct-admission/gig-platform-deactivation)**

- AI Agent Accountability: What Happens When an Agent Is Accused (/articles/conduct-admission/ai-agent-accountability-stakes)

APPLICATIONS · SPECIFIC

- The EU AI Act and Inter-Agent Conduct: Where Record-Keeping and Oversight Duties Stop (/articles/conduct-admission/eu-ai-act-accountability-records)
- Sierra AI Agents and the Record Behind a Customer Complaint (/articles/conduct-admission/sierra-ai)
- Decagon: Support Agents and Record-Grounded Accountability (/articles/conduct-admission/decagon-ai)
- Intercom Fin: Resolution Rates and the Conduct Record (/articles/conduct-admission/intercom-fin)
- Zendesk AI Agents: Escalation, Audit, and Agent Conduct (/articles/conduct-admission/zendesk-ai-agents)
- Agentforce and the Accusation an Agent Must Answer (/articles/conduct-admission/salesforce-agentforce-conduct)

TERMINOLOGY

- append-only lineage field (/glossary#append-only-lineage-field)
- authorization gate (/glossary#authorization-gate)
- conduct evaluation artifact (/glossary#conduct-evaluation-artifact)
- principal (/glossary#principal)
- semantic agent (/glossary#semantic-agent)
- signed policy object (/glossary#signed-policy-object)

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