A Five-Tier Oracle Pipeline with Bounded Deferral

Promotion, Dispute Bonding, and Consensus-Checked Finality

Raeez Lorgat

April 2026

Abstract

An event-native chain must resolve heterogeneous external events inside consensus. It must do so without allowing one faulty source, one disputed report, or one under-collateralized escalation step to stall block production or weaken finality. This paper gives a standalone formalization of a five-tier oracle pipeline with states T_0 through T_4 ranging from raw evidence to consensus acceptance. The first contribution is architectural: the five tiers correspond to five distinct invariants-syntactic validity, authentication, corroboration, bonded disputability, and consensus acceptance-and the promotion and demotion maps are explicit monotone partial functions. The second contribution is a source-local circuit breaker: a single-source fault marks only the events that depend on that source with the defer flag, Phase 1 skips those events, and Phase 2 clears only the resolved subset. The third contribution is an economic rule for dispute collateral,

$$B(E) = \max\left\{B_{\text{esc}}(E),\ 0,\ \frac{V(E) + U(E) - R(E)}{q(E)\,(2p(E) - 1)}\right\},$$

where $V(E)$ is event value, $U(E)$ is settlement-unwind cost, $q(E)$ is dispute probability for a false report, $p(E)$ is adjudicator correctness, and $R(E)$ is the honest winner reward. Under locally stated assumptions, truthful submission weakly dominates dishonest submission, and strictly dominates when the bond exceeds the floor. The fourth contribution is an escalation-monotone bond floor enforced at consensus. Every accepted outcome carries forward a floor at least as large as the prior tier's floor, so oracle disputes cannot unwind settled blocks below that carried floor.

Contents

1 Introduction

Oracle resolution for discrete events has two constraints that operate simultaneously. The chain must resolve many events inside consensus fast enough to keep block production live, and it must defend against adversaries whose gains can exceed the face value of the event itself once clearing and settlement are taken into account. A pipeline that treats every event identically wastes latency on simple cases and under-prices disputes on hard cases. A pipeline that escalates without preserving a monotone economic floor creates a cheaper attack path at a later tier than at an earlier tier. A pipeline that halts an entire block when one source fails turns a local data fault into a consensus fault.

This paper isolates the event-resolution problem and formalizes a five-tier oracle pipeline with two load-bearing properties:

(i) **Source-local failure handling.** A single-source failure defers the affected events; it does not halt the block. Events resolve in Phase 1 at consensus, Phase 2 clears only the resolved subset, and the defer cascade is bounded by the number of events that directly depend on the failed source in that round.

(ii) **Escalation-monotone economic protection.** Each escalation step carries a bond or slashable-capital floor at least as large as the prior tier's floor. Consensus checks that floor before accepting the outcome into a settled block, so a later oracle dispute cannot unwind that block below the carried floor.

The paper makes four formal claims. Section 3 defines the five tiers T_0 through T_4, the promotion map π, the demotion map δ, and proves that five distinct states are minimal for the invariants the pipeline needs to preserve. Section 4 defines a source-local circuit breaker and proves bounded deferral plus a bounded-time Phase 1 guarantee. Section 5 defines the dispute bond as a function of event value, dispute probability, and unwind cost, and proves truthful submission under that bond. Section 6 defines the escalation-monotone bond floor and proves that consensus enforcement prevents low-capital unwind of settled blocks. Section 7 states the adversarial assumptions precisely and gives Sybil, data-feed, and collusion bounds.

2 System Model

Definition 2.1 (Event)**.** An *event* is a tuple

$$E = (Q, \Omega, S(E), V(E), U(E)),$$

where Q is the query to be resolved, Ω is a finite admissible outcome set, $S(E)$ is the finite set of sources consulted by the oracle for the event, $V(E) \in \mathbb{R}_{\geq 0}$ is the maximum event-local value extractable from a false resolution before repair, and $U(E) \in \mathbb{R}_{\geq 0}$ is the cost of unwinding or compensating a false outcome after settlement has propagated.

Definition 2.2 (Source Integration Classes)**.** Every source adapter belongs to exactly one of the following abstract classes:

(a) *reputation-weighted rate sources*, which emit signed numerical or categorical reports and carry reputation or deposit loss under detectable misreporting;

(b) *cryptographic-attestation feeds*, which emit authenticated statements whose validity is checked against registered keys or attestations;

(c) *threshold-signed committees*, which emit a report only after a threshold of eligible members signs the same payload;

(d) *physical-sensor IoT*, which emit measurements coupled to device identity, freshness, and sensor-attestation checks;

(e) *multi-source aggregation*, which combine outputs from at least two distinct source classes under a published aggregation rule.

Definition 2.3 (Evidence Bundle)**.** For an event E in round r, an *evidence bundle* is a finite multiset

$$X_r(E) = \{(s, o, \sigma, \tau, \kappa)\},$$

where $s \in S(E)$ is the source identity, $o \in \Omega$ is the reported outcome, σ is the source's authenticity witness, τ is the report timestamp, and κ is the source class from Definition 2.2.

Assumption 2.4 (Consensus Round Structure)**.** Each consensus round r admits a finite candidate set $\mathcal{E}_r$ of oracle events. Phase 1 has a protocol deadline $\tau_1(r)$ by which each event in $\mathcal{E}_r$ must be labeled either accept or defer. Phase 2 begins immediately after Phase 1 and consumes only the Phase 1 accepted subset.

Assumption 2.5 ((t, n) Honest-Majority Committees)**.** Whenever a threshold-signed committee of size n is used at a corroboration or dispute tier, the committee size satisfies $n \geq 3t + 1$, at most t members are Byzantine, and the signature threshold q satisfies $q \geq 2t + 1$. For stake-weighted committees, honest participants control strictly more than one half of admissible weight.

Assumption 2.6 (Sybil Resistance by Weight)**.** Every eligible source identity or voting identity carries a non-zero admission cost or locked capital proportional to the voting weight it contributes. Creating additional identities does not create voting weight for free.

3 Five-Tier Architecture

The pipeline's states are ordered on the finite chain

$$T_0 \prec T_1 \prec T_2 \prec T_3 \prec T_4,$$

with the tier index used as the order parameter.

Definition 3.1 (Tier Invariants)**.** For an event certificate $C(E)$, define the following invariants.

I_0 the bundle is syntactically valid and references only admissible outcomes;

I_1 every retained report is authenticated, fresh, and class-typed;

I_2 a corroboration rule has produced a candidate outcome and witness;

I_3 the candidate outcome is backed by collateral $b \geq B(E)$ and an open dispute relation;

I_4 Phase 1 consensus has accepted the certificate into the block.

Definition 3.2 (Five Tiers)**.** The pipeline assigns each event certificate to exactly one of the following states.

T1: Raw. T_0 contains evidence bundles satisfying only I_0.

T2: Authenticated. T_1 contains evidence bundles satisfying I_0 and I_1.

T3: Corroborated. T_2 contains certificates satisfying I_0, I_1, and I_2. Corroboration may arise from cross-source agreement, a threshold-signed committee, or a multi-source aggregation witness.

T4: Bonded-Disputable. T_3 contains certificates satisfying I_0 through I_3.

T5: Consensus-Accepted. T_4 contains certificates satisfying I_0 through I_4. Outcomes in T_4 are the Phase 1 outputs seen by Phase 2.

Definition 3.3 (Promotion Map)**.** For $k \in \{0, 1, 2, 3\}$, the promotion map

$$\pi_k : T_k \rightharpoonup T_{k+1}$$

is the partial function defined by

$$\pi_k(x) = \begin{cases} x' & \text{if } x \text{ satisfies } I_k \text{ and the guard } G_{k\to k+1}(x) \text{ holds,} \\ \bot & \text{otherwise.} \end{cases}$$

The guards are:

$G_{0\to1}$: authentication, freshness, and source-class checks pass;
$G_{1\to2}$: corroboration succeeds under the admissible aggregation rule;
$G_{2\to3}$: a designated bond-proposer locks $b \geq B(E)$ and opens the dispute relation;
$G_{3\to4}$: Phase 1 consensus validates the certificate and admits it into the block.

The *bond-proposer* is the participant bound to the certificate by the protocol at tier T_3; in a submitter-posts-bond regime this is the reporting submitter, and in a delegated-bond regime it is the delegate-of-record.

Definition 3.4 (Demotion Map)**.** For $k \in \{1, 2, 3\}$, the intra-round demotion map

$$\delta_k : T_k \to \{T_0, T_1, T_2, \mathsf{defer}\}$$

returns the highest lower tier whose invariant remains valid in the current round after the failure that triggered review. For $k = 4$, the post-acceptance follow-on map

$$\delta_4 : T_4 \to T_3^{(r+1)}$$

emits a fresh T_3 dispute record for round $r + 1$ without rewriting the settled block. Concretely:

(a) δ_1 returns defer if authentication or freshness fails;

(b) δ_2 returns T_1 if reports remain authenticated but corroboration fails;

(c) δ_3 returns T_2 if the evidence remains corroborated but the bond or dispute precondition fails;

(d) δ_4 emits a fresh T_3 dispute record in round $r + 1$ when a post-acceptance contradiction is raised. The settled block in round r is not rewritten.

Proposition 3.5 (Promotion Monotonicity and Intra-Round Termination)**.** *If $\pi_k(x)$ is defined, then the tier index strictly increases by one. If $\delta_k(x) \neq$ defer for $k \in \{1, 2, 3\}$, the returned tier index is strictly smaller than k. Within a single consensus round r, each event undergoes at most four successive promotions and at most three successive intra-round demotions before receiving a terminal* accept *or* defer *label. Post-acceptance maps δ_4 are not intra-round transitions and cannot trigger a sequence that revisits round r.*

Proof. By Definition 3.3, π_k maps from T_k to T_{k+1} and nowhere else, so the tier index increases by one whenever promotion succeeds. By Definition 3.4, each intra-round δ_k with $k \in \{1, 2, 3\}$ returns either defer or a strictly lower tier. The intra-round tier index ranges over the finite chain $\{0, 1, 2, 3, 4\}$, so no infinite monotone sequence exists inside the round. The post-acceptance map δ_4 writes into round $r + 1$ by Definition 3.4, so it does not create a cycle that reopens round r. Intra-round termination follows. □

Proposition 3.6 (Five Tiers Are Minimal for the Stated Invariants)**.** *Fix the five invariants $I_0, \ldots, I_4$ of Definition 3.1. Any state machine whose reachable states each carry a monotone sub-sequence of those invariants, and which separates the five properties of heterogeneous raw ingestion, authentication, corroboration, bonded disputes, and consensus-checked finality, must admit at least five distinct states. The chain of Definition 3.2 realizes this minimum.*

Proof. The five invariants in Definition 3.1 are pairwise distinct and each load-bearing for a separate property.

(a) If I_0 and I_1 are merged into a single state, malformed raw input and authenticated input become indistinguishable, and a parser failure can no longer be quarantined before source authentication is evaluated.

(b) If I_1 and I_2 are merged, a single authenticated source can bypass corroboration because the system no longer distinguishes "authenticated" from "corroborated."

(c) If I_2 and I_3 are merged, evidence sufficiency and collateral sufficiency become indistinguishable, and dispute logic cannot tell whether failure arose from weak evidence or weak bonding.

(d) If I_3 and I_4 are merged, an open dispute state and a settled consensus state become indistinguishable, and finality cannot carry a checked economic floor into the block.

Under the hypothesis of the proposition, each reachable state is labelled by a downward-closed subset of the invariant chain, so distinct invariants induce distinct states. Any alternative decomposition that separates the same five properties is related to Definition 3.1 by relabeling and therefore also requires five states. □

Remark 3.7 (Scope of Minimality)**.** Proposition 3.6 bounds the number of distinct states below by five *under the fixed invariant decomposition*. It does not preclude alternative decompositions that use a different set of invariants, for example a pipeline that drops consensus-checked finality or fuses corroboration into dispute adjudication. The claim is therefore that, given the stated invariant set, five is both necessary and sufficient.

4 Circuit Breaker, Deferral, and Consensus Phases

Definition 4.1 (Affected-Event Set)**.** For a source s and round r, define

$$A_r(s) = \{E \in \mathcal{E}_r : s \in S(E)\}.$$

This is the set of candidate events in round r that depend directly on source s.

Definition 4.2 (Source-Local Circuit Breaker)**.** The circuit breaker for source s in round r is the Boolean predicate

$$\mathrm{CB}_r(s) = 1$$

if at least one of the following holds for the source's contribution to that round: authentication failure, freshness failure, threshold-certificate failure, or class-specific coherence failure. An event $E \in \mathcal{E}_r$ receives the defer flag in round r precisely when

$$\mathsf{defer}_r(E) = 1 \iff \Big(\exists s \in S(E) : \mathrm{CB}_r(s) = 1\Big) \vee \Big(\pi_1 \text{ or } \pi_2 \text{ fails before } \tau_1(r)\Big).$$

Theorem 4.3 (Single-Source Failure Is Local)**.** *Suppose exactly one source $s^\star$ fails in round r, so $\mathrm{CB}_r(s^\star) = 1$ and $\mathrm{CB}_r(s) = 0$ for all $s \neq s^\star$. Then every event in $\mathcal{E}_r \setminus A_r(s^\star)$ remains eligible for promotion independently of the failed source. Only events in $A_r(s^\star)$ can be deferred due to that fault. In particular, a single-source failure defers the affected events and does not halt the block.*

Proof. By Definition 4.1, a source fault influences an event only through membership of that source in the event's dependency set. For any $E \notin A_r(s^\star)$, the failed source does not appear in $S(E)$, so the first disjunct in Definition 4.2 is false for that event. All promotion guards for that event are evaluated from its own evidence bundle and own collateral state. The failed source therefore changes no predicate for E. Only events in $A_r(s^\star)$ inherit the source fault. Since Phase 1 labels events individually under Assumption 2.4, block production continues on the unaffected events. □

Proposition 4.4 (Bounded Deferral Cascade)**.** *For every round r,*

$$N_r := \left|\{E \in \mathcal{E}_r : \pi_1(E) \text{ or } \pi_2(E) \text{ fails before } \tau_1(r) \text{ for reasons not caused by a source circuit breaker}\}\right|.$$

Then

$$\left|\{E \in \mathcal{E}_r : \mathsf{defer}_r(E) = 1\}\right| \leq \left|\bigcup_{s:\mathrm{CB}_r(s)=1} A_r(s)\right| + N_r \leq \sum_{s:\mathrm{CB}_r(s)=1} |A_r(s)| + N_r.$$

If exactly one source fails and no non-source promotion guard fails before $\tau_1(r)$, then

$$\left|\{E \in \mathcal{E}_r : \mathsf{defer}_r(E) = 1\}\right| \leq |A_r(s^\star)|.$$

No recursive defer cascade occurs inside the same round.

Proof. An event is deferred only if it depends on at least one failed source or if corroboration or bonding does not complete by the round deadline. The source-triggered contribution is contained in the union of the affected-event sets; the remaining contribution is exactly N_r. Taking cardinalities yields the bound. There is no recursive cascade because defer is terminal for the current round: a deferred event is skipped by Phase 1 rather than reused as evidence for other events in the same round. □

Definition 4.5 (Phase Outputs)**.** At the end of Phase 1 in round r, the candidate set $\mathcal{E}_r$ is partitioned as

$$\mathcal{E}_r = \mathcal{R}_r \sqcup \mathcal{D}_r,$$

where $\mathcal{R}_r$ is the set of events promoted to T_4 and $\mathcal{D}_r$ is the set of events marked defer. Phase 2 receives only $\mathcal{R}_r$.

Theorem 4.6 (Phase 1 Bounded-Time Guarantee)**.** *Under Assumption 2.4, Phase 1 terminates within one consensus round and outputs the partition of Definition 4.5. Every event in $\mathcal{R}_r$ resolves at consensus in round r, every event in $\mathcal{D}_r$ is skipped in round r, and Phase 2 runs only on the resolved subset $\mathcal{R}_r$. Moreover, each event in $\mathcal{E}_r$ requires at most four promotion-guard evaluations plus at most one circuit-breaker evaluation before it is labeled, so the per-event work is bounded independently of the deadline.*

Proof. The candidate set $\mathcal{E}_r$ is finite. For each event $E \in \mathcal{E}_r$, the labeling procedure evaluates the promotion guards $G_{0\to1}$, $G_{1\to2}$, $G_{2\to3}$, $G_{3\to4}$ in sequence, or the circuit-breaker predicate CB_r, with at most one invocation of each. This is a bounded number of steps per event by Definitions 3.3 and 4.2, so the labeling algorithm terminates for each event in bounded work.

By Assumption 2.4, each event must carry either the accept label or the defer label by the deadline $\tau_1(r)$. The accept label corresponds to successful promotion through T_4 under Definition 3.3; the defer label corresponds to Definition 4.2. No third state exists at the deadline. Thus the partition of Definition 4.5 exists and is complete. Because Phase 2 begins only after Phase 1 closes and consumes only $\mathcal{R}_r$, deferred events are skipped rather than blocking the round. □

5 Bond Economics

Definition 5.1 (Dispute Parameters)**.** For an event E, let

- $q(E) \in (0, 1]$ be the probability that a false report is disputed within the challenge window;
- $p(E) \in (1/2, 1]$ be the probability that the dispute adjudicator rules correctly once the dispute is raised;
- $R(E) \in \mathbb{R}_{\geq 0}$ be the honest winner reward for correct submission and successful completion.

Definition 5.2 (Base Bond Function)**.** The *base dispute bond* for event E is

$$B_{\text{base}}(E) = \max\left\{0,\ \frac{V(E) + U(E) - R(E)}{q(E)\,(2p(E) - 1)}\right\}.$$

The numerator of the fraction is the dishonest gross gain net of the honest reward. The denominator is the expected penalty multiplier induced by dispute probability and adjudicator correctness, and is strictly positive by the parameter ranges in Definition 5.1. The outer maximum with zero handles the case $R(E) \geq V(E) + U(E)$, in which the honest reward already dominates every possible dishonest payoff and no positive bond is required for local dominance. All monetary quantities $V(E)$, $U(E)$, $R(E)$ and the bond itself are denominated in a single slashable-capital unit.

Remark 5.3 (Interpretation of the Bond Function)**.** Definition 5.2 is the formal statement of the bond economics. The bond is a function of event value $V(E)$, dispute probability $q(E)$, settlement-unwind cost $U(E)$, honest winner reward $R(E)$, and adjudicator correctness $p(E)$. The honest winner reward $R(E)$ lowers the bond needed to make truthful submission dominant because it raises the truthful branch's payoff without strengthening the false branch.

Theorem 5.4 (Truthful Submission Dominates Dishonest Submission)**.** *Fix an event E and restrict the submitter's local strategy set to* $\{\text{truthful}, \text{false}, \text{abstain}\}$*, with payoffs $R(E)$, the false-branch expected payoff computed below, and* 0*, respectively. Suppose the protocol requires collateral $b \geq B_{\text{base}}(E)$, where $B_{\text{base}}(E)$ is given by Definition 5.2. Then truthful submission weakly dominates* false *and weakly dominates* abstain*. If $b > B_{\text{base}}(E)$ and the fraction in Definition 5.2 is positive, truthful submission strictly dominates* false*; if $R(E) > 0$, truthful submission strictly dominates* abstain*.*

Proof. The dispute is structured as a two-sided bond game: the submitter posts collateral b, and any honest challenger posts equal collateral b; on adjudication the losing side forfeits its bond to the winner. The truthful payoff is $R(E)$: the bond is returned because a correct report is not successfully slashed. The abstain payoff is 0, since the submitter neither posts bond nor earns a reward. Truthful therefore weakly dominates abstain whenever $R(E) \geq 0$, which holds by Definition 5.1, and strictly dominates when $R(E) > 0$.

For the false branch, the submitter posts bond b and gains at most $V(E) + U(E)$ gross. A dispute arises with probability $q(E)$. Conditional on dispute, the adjudicator is correct with probability $p(E)$, in which case the submitter loses bond b to the challenger, and incorrect with probability $1 - p(E)$, in which case the submitter wins the challenger's bond. The false-branch expected payoff is therefore bounded above by

$$V(E) + U(E) - q(E)\big(p(E) - (1 - p(E))\big)b = V(E) + U(E) - q(E)(2p(E) - 1)b.$$

If $b \geq B_{\text{base}}(E) \geq \frac{V(E)+U(E)-R(E)}{q(E)(2p(E)-1)}$, this quantity is at most $R(E)$. If the inequality is strict and $V(E) + U(E) > R(E)$, the bound is strictly below $R(E)$. When $R(E) \geq V(E) + U(E)$, the fraction is non-positive, so even $b = 0$ already yields false-branch payoff $\leq R(E)$. Truthful therefore weakly dominates false in every case and strictly dominates when $b > B_{\text{base}}(E)$ and $V(E) + U(E) > R(E)$. □

Corollary 5.5 (Incentive Compatibility Under Honest Challenges)**.** *If honest challengers act whenever a false report is observed, so that $q(E) > 0$, and the adjudicator satisfies $p(E) > 1/2$, then a submitter cannot rationally improve by replacing a truthful report with a false report once the bond floor of Theorem 5.4 is enforced.*

6 Escalation-Monotone Bond Floor and Consensus Enforcement

Definition 6.1 (Tier Floors)**.** For each event E, define the tier floor sequence recursively by

$$\begin{aligned}
\text{bond}_0(E) &= 0,\\
\text{bond}_1(E) &= C_1(E),\\
\text{bond}_2(E) &= \max\{\text{bond}_1(E), C_2(E)\},\\
\text{bond}_3(E) &= \max\{\text{bond}_2(E), B_{\text{base}}(E)\},\\
\text{bond}_4(E) &= \text{bond}_3(E).
\end{aligned}$$

Here $C_1(E)$ denotes the accountability-capital floor at the authenticated-source tier, measured as the slashable capital attributable to the source set $S(E)$. The quantity $C_2(E)$ denotes the minimum cost of corrupting the corroboration quorum or threshold certificate for the event, measured by pricing each unit of required admissible weight at its admission or acquisition cost. All three floors $C_1(E)$, $C_2(E)$, and $B_{\text{base}}(E)$ are denominated in the same slashable-capital unit declared in Definition 5.2, so the max operations compare like quantities.

Definition 6.2 (Effective Bond)**.** The effective dispute bond required at T_3 is

$$B(E) = \max\{\text{bond}_2(E), B_{\text{base}}(E)\} = \text{bond}_3(E).$$

Equivalently,

$$B(E) = \max\left\{B_{\text{esc}}(E),\ \frac{V(E) + U(E) - R(E)}{q(E)\,(2p(E) - 1)}\right\}$$

with the escalation floor identified as $B_{\text{esc}}(E) = \text{bond}_2(E)$.

Definition 6.3 (Consensus Acceptance Rule)**.** Phase 1 accepts an event certificate into T_4 only if the certificate contains a witness that the carried collateral or slashable-capital floor is at least $\text{bond}_4(E)$. Certificates below that floor are invalid for block inclusion.

Theorem 6.4 (Escalation-Monotone Bond Floor)**.** *For every event E and every escalation step $k \in \{1, 2, 3, 4\}$,*

$$\text{bond}_k(E) \geq \text{bond}_{k-1}(E).$$

If an event reaches T_4, then the accepted block carries a checked floor of at least $\text{bond}_4(E)$. Consequently no later oracle dispute can unwind the settled block below that floor.

Proof. The recursive definition in Definition 6.1 uses the maximum of the previous tier's floor and the current tier's intrinsic corruption cost, so $\text{bond}_k(E) \geq \text{bond}_{k-1}(E)$ follows immediately for each step.

For the finality claim, Definition 6.3 makes the floor part of the Phase 1 validity condition. An under-collateralized certificate cannot enter T_4. Once a certificate is accepted into a settled block, later disagreement is represented through a new T_3 dispute record in a later round by Definition 3.4; the accepted block is not rewritten. Any follow-on challenge must therefore meet or exceed the carried floor already checked at T_4. The block cannot be unwound below $\text{bond}_4(E)$ through the oracle layer. □

Remark 6.5 (Load-Bearing Rationale)**.** Theorem 6.4 is the formal reason escalation cannot create a cheaper attack at a later tier than at an earlier tier. Without the max recursion in Definition 6.1, an adversary could prefer the later dispute tier whenever its required bond fell below the earlier tier's quorum-corruption cost. The recursion removes that gap.

7 Adversarial Analysis

Definition 7.1 (Adversarial Weight Deficit)**.** For a weighted committee or voting pool with acceptance threshold q and adversarial admissible weight w_A, define the *weight deficit*

$$\Delta_w = \max\{0, q - w_A\}.$$

This is the additional admissible weight the adversary must obtain or bribe to force acceptance.

Proposition 7.2 (Sybil Resistance via Stake Weight)**.** *Under Assumption 2.6, forcing acceptance at any weighted corroboration or dispute tier requires the adversary to pay for at least Δ_w additional admissible weight. The cost of a Sybil attack is therefore bounded below by the admission cost of that missing weight.*

Proof. By Assumption 2.6, weight is not created for free. If the adversary already has weight at least q, then $\Delta_w = 0$ and no Sybil attack is needed. If not, any successful attack must acquire or bribe at least $q - w_A$ additional admissible weight. The admission or collateral cost of that weight lower-bounds the attack. □

Proposition 7.3 (Single-Source Compromise Response)**.** *If one source instance is compromised, then at least one of the following holds, and the three cases together cover all outcomes:*

(a) the source fails an authentication, freshness, or coherence check and the circuit breaker defers the events in its affected-event set; or

(b) the source remains admissible but contributes only minority weight, in which case corroboration at T_2 or dispute at T_3 rejects the false outcome under Assumption 2.5; or

(c) the adversary purchases enough additional admissible weight to close the deficit Δ_w, in which case the cost is lower-bounded by Proposition 7.2.

Proof. If the compromised source fails any check listed in Definition 4.2, case (a) applies. Otherwise the source remains admissible. If its weight alone does not meet the threshold required for corroboration or dispute acceptance, then under Assumption 2.5 the honest complement controls the decision, giving case (b). If the source does control sufficient weight to force acceptance, Assumption 2.6 implies the adversary has paid for that admissible weight, which yields case (c) via Proposition 7.2. The three cases together are exhaustive; (b) and (c) are mutually exclusive because they partition the admissible-weight control, and (a) is exclusive of both because the source has been rejected before any weight contribution is scored. □

Theorem 7.4 (Collusion Bound)**.** *Consider an event E resolved through a threshold-signed corroboration tier or a bonded dispute tier. Let $\lambda(E)$ be a lower bound on the bribe or acquisition cost per unit of admissible honest weight for the relevant committee. Then every collusion strategy that produces false acceptance into T_4 incurs total expected cost at least*

$$\min\{\lambda(E)\Delta_w,\ B(E)\},$$

including mixed strategies that combine partial weight acquisition with partial bond posting. Any strategy that attempts to avoid the weight-acquisition cost by forcing escalation encounters the effective bond floor $B(E)$ of Definition 6.2.

Proof. Any successful path to false acceptance must pass through either T_2 (false corroboration) or T_3 (false bonded dispute that resolves in the adversary's favor). Let the adversary's actual strategy close a fraction $\alpha \in [0,1]$ of the weight deficit and post collateral fraction $\beta \in [0,1]$ of the effective bond toward the dispute route.

If $\alpha < 1$, the adversary cannot force a valid corroboration certificate: by Assumption 2.5, forming a threshold-signed certificate requires honest weight at least q to concur, which the adversary has not secured. The only remaining path is the bonded dispute tier, which by Definition 6.3 rejects any certificate whose collateral is below $\text{bond}_4(E) \geq B(E)$. Therefore $\beta = 1$ is required, and the dispute-route cost is at least $B(E)$.

If $\alpha = 1$, the adversary has closed the full weight deficit Δ_w, which by Proposition 7.2 costs at least $\lambda(E)\Delta_w$.

The adversary's total cost is at least $\alpha\lambda(E)\Delta_w + \beta B(E)$, minimized over the admissible pairs $(\alpha = 1, \beta \geq 0)$ and $(\alpha < 1, \beta = 1)$. The minimum of the admissible region is $\min\{\lambda(E)\Delta_w,\ B(E)\}$; no mixed strategy with $\alpha < 1$ and $\beta < 1$ yields false acceptance. □

Corollary 7.5 (Honest-Majority Interpretation)**.** *Under Assumption 2.5, if the adversary does not close the weight deficit and does not post the effective bond, then the pipeline's only admissible response to disagreement is bounded deferral. False acceptance is excluded.*

Proof. If the adversary fails both conditions, it cannot produce a valid corroboration certificate and cannot finance a winning bonded dispute. The only remaining path is the defer branch of Definition 4.2 together with Theorem 4.6. □

8 Conclusion

The five-tier pipeline separates five distinct invariants that event resolution needs: raw validity, authentication, corroboration, bonded disputability, and consensus acceptance. The promotion and demotion maps make those transitions explicit and monotone. The source-local circuit breaker prevents a local source failure from turning into a block-wide halt by deferring only the affected events. Phase 1 resolves or defers every candidate event within one consensus round, and Phase 2 operates only on the resolved subset. The bond function ties oracle collateral directly to event value, dispute probability, and unwind cost, and the escalation-monotone floor carries that protection into consensus finality. Under the stated honest-majority and Sybil-resistance assumptions, single-source compromise, Sybil expansion, and collusion all reduce to bounded deferral, admissible-weight acquisition, or the effective bond floor.

References

[1] M. Castro and B. Liskov, "Practical Byzantine fault tolerance," *Proceedings of the Third Symposium on Operating Systems Design and Implementation*, pp. 173-186, 1999.

[2] V. Shoup, "Practical threshold signatures," *Advances in Cryptology - EUROCRYPT 2000*, pp. 207-220, 2000.

[3] R. Hanson, "Logarithmic market scoring rules for modular combinatorial information aggregation," *Journal of Prediction Markets*, vol. 1, no. 1, pp. 3-15, 2007.

[4] Y. Chen, X. A. Gao, R. Goldstein, and I. A. Kash, "Market manipulation with outside incentives," *Proceedings of AAAI 2012*, pp. 614-620, 2012.

[5] V. Buterin, "A minimal-trust universal data feed," 2014.

[6] J. Adler, R. Berryhill, A. Veneris, Z. Poulos, N. Veira, and A. Kastania, "Astraea: A decentralized blockchain oracle," *Proceedings of the 2018 IEEE International Conference on Blockchain (Blockchain)*, pp. 1145-1152, 2018.

[7] L. Breidenbach, C. Cachin, B. Chan, A. Coventry, S. Ellis, A. Juels, F. Koushanfar, A. Miller, B. Magauran, D. Moroz, S. Nazarov, A. Topliceanu, F. Tramer, and F. Zhang, "Chainlink 2.0: Next steps in the evolution of decentralized oracle networks," *Chainlink Whitepaper v2.0*, 2021.

[8] A. Juels, A. Kosba, and E. Shi, "The ring of Gyges: Investigating the future of criminal smart contracts," *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 283-295, 2016.