

Receipted Actions: A Reproducible Audit Capsule for Rewarded-Action Payout Adjudication

Noah Erlwein¹

¹*Invariant Systems Inc.*

(Dated: May 2026)

Commercial decision systems increasingly act on externally claimed events: a rewarded action completed, an install attributed, or a control gate passed. These actions often trigger money movement or operational state changes, yet the evidence chain behind the decision is commonly scattered across logs, vendor dashboards, and informal review queues. We present a reproducible audit capsule for claim-based commercial adjudication, organized around a practical five-surface pattern we call *receipted actions*: typed claim, evidence references, machine-readable verdict, receipt surface, and appeal path. We instantiate the capsule on rewarded-action payout adjudication, where mobile acquisition and offerwall claims can be evaluated offline before production payout enforcement. The released artifact uses typed signal and postback contracts, replayable verdict envelopes, reason codes, manifests, checksum records, a public schema file, a release-local validator, and a synthetic appeal sample. The public evidence bundle contains a 10,150-event deterministic synthetic conformance artifact with 10,150 rewarded-action events, 10,150 verdict envelopes, a manifest, checksums, and a synthetic appeal sample. Public verification is exact and byte-for-byte for the released README, manifest, event, verdict, and synthetic appeal files, and the release-local validator checks the public verdict JSONL against the shipped schema contract and manifest totals while validating the synthetic appeal-sample shape. The current public capsule exposes receipt binding at the bundle level through manifests, checksum records, and publication receipts; it does not expose a separate public receipt hash for each verdict row. Partner-specific event data, live connector credentials, tuned scoring thresholds, and production heuristics are deliberately excluded. The result is not a claim of live fraud-detection efficacy; it is a reproducible audit surface for claim-based commercial adjudication.

I. INTRODUCTION

Commercial software increasingly clears or blocks actions that are claimed by another party. In rewarded-action advertising, a network claims that a user completed an install, tutorial, purchase, or reward path and a payout should clear. Invalid-traffic guidance recognizes that such claims sit inside a taxonomy of detectable and disputed traffic quality problems[1]. In AI systems, an operator claims that a model response was emitted under particular parameters. In compliance workflows, a team claims that a control gate passed. The common problem is not merely prediction accuracy. The common problem is accountable action: what was claimed, what evidence supported it, who or what issued the verdict, what receipt binds the decision, and how the affected party can challenge it.

We call this pattern a *receipted action*. A receipted action has five surfaces:

1. a typed claim;
2. evidence references;
3. a verdict with reason codes;
4. a receipt or content-addressed decision record; and
5. an appeal or override path.

The contribution of this paper is the released audit capsule and its public verification boundary. The five-surface receipted-action pattern is the framing used to

organize that capsule, not a claim that decision logging itself is newly invented.

This paper uses rewarded-action payout adjudication as the first concrete case study. The domain is commercially important and technically constrained: attribution claims are noisy, fraud and replay are common, and overblocking creates partner disputes. A useful system must therefore be more than a classifier. It must emit reviewable decisions.

II. RECEIPTED-ACTION NOTATION AND VALIDITY BOUNDARY

We use compact notation to distinguish the optional per-verdict receipt from the public bundle receipt actually released in v1. Let the externally asserted claim for row i be denoted by c_i , the available evidence references by E_i , and the adjudicator output by V_i . A receipted action is

$$A_i = (c_i, E_i, V_i, \rho_i, a_i), \quad (1)$$

where the receipt surface ρ_i can take one of two forms. Optional per-verdict receipt binding is

$$r_i = H(c_i, E_i, V_i, m_i), \quad (2)$$

while the public bundle receipt exposed in the current release is

$$R_B = (H(M_B), H(S_B), m_B), \quad (3)$$

with M_B the released manifest, S_B the released checksum set, and m_B release metadata such as release state, date, and artifact location. Here $H(x)$ denotes the SHA-256 digest of the released byte representation of x . In the current public capsule, $\rho_i = R_B$ at release level rather than a separately published r_i for each verdict row. The notation is not offered as a proof calculus. Its purpose is to make explicit which fields must be inspectable in a released decision bundle.

For the current paper, a released bundle is publicly valid when three conditions hold: each public verdict envelope conforms to the released schema contract; the manifest and checksum file bind the released event, verdict, and appeal artifacts byte-for-byte; and the release record states the manifest digest, checksum digest, release identifier, and artifact location. We refer to that release record as a *publication receipt*. In the shipped artifact, R_B is therefore a structured publication receipt rather than a single composite public hash: it records digest fields for the manifest and checksum set together with release metadata, matching the released receipt surface. Operationally, the bundle also ships with a release-local validator that checks each public verdict JSONL row against the released verdict-envelope contract, reconciles verdict totals and reason-code histograms against the manifest, and validates that the synthetic appeal sample is a block verdict with a non-empty appeal path. Private or counterparty-specific deployments may additionally bind individual verdicts to per-decision receipt identifiers. The design does not require that all evidence be public. It is informed by provenance, policy, and supply-chain systems that bind reviewable artifacts without forcing every private input into the released surface[2–5]. It does require that the public or counterparty-facing object identify what evidence class was used, what verdict was returned, and what can be challenged.

For payout adjudication, verdicts are operational: allow, hold, or block. For slower governance loops, verdicts may be evaluative: pass, at-risk, or fail. The same envelope shape can carry both, provided the verifier speed and claim type are explicit.

A. Verdict-Envelope Schema

Listing 1 shows a compact public-safe verdict-envelope schema excerpt that matches the released JSON row shape. The design choice is deliberate: public artifacts expose typed claim fields, evidence references, verdict output, and an appeal locator without forcing the release of raw partner traffic, tuned thresholds, or operational source paths.

Listing 1. Compact public-safe verdict-envelope schema excerpt.

```
{
  "version": "kaleidos.verdict_envelope.v1",
  "verdict": "allow|hold|block",
```

```
  "claim": {
    "claim_type": "rewarded_action_completed",
    "subject": "com.example.app:evt-00000042",
    "context": {
      "timestamp_epoch": 1777075569.818111,
      "campaign_id": "camp-023",
      "partner_id": "partner-011",
      "geo": "BR",
      "payout_amount": 7.02,
      "payout_currency": "USD"
    }
  },
  "verifier": {
    "name": "public_release_adjudicator",
    "speed": "operator",
    "version": "kaleidos.public_release_adjudicator.v1"
  },
  "reason_codes": ["device_farm", "..."],
  "evidence_refs": ["device:..."],
  "partner_disposition": "allow",
  "appeal_path": "https://example.invalid/appeal-review"
}
```

The released rows do not embed the full appeal state machine or a per-row required-fields object. Those workflow details remain part of the surrounding review process and synthetic appeal sample rather than every published row. The exact required row fields are enforced by the shipped schema file and release-local validator rather than restated inline in every verdict object.

B. Receipt Levels

The formal model permits per-action receipt identifiers, but a public artifact can bind at different levels. Table I makes that distinction explicit so that the current paper is not misread as claiming a public per-verdict transparency log.

III. REWARDED-ACTION PAYOUT CASE STUDY

Rewarded-action advertising creates a compact test case for receipted actions. A payout claim has a known commercial consequence and an obvious dispute path. The system under study is intentionally not a bidding engine or a targeting engine. It is a payout adjudication layer.

The case-study pipeline has four contract surfaces:

1. a thin game or app signal payload;
2. a normalized mobile measurement partner postback;
3. a payout decision surface; and
4. a conformance artifact or pilot report.

The design keeps high-entropy scoring, benchmark mixes, and calibration logic server-side. Client or partner integrations emit observations, not final verdicts. This boundary is important: it allows a buyer to validate a small integration without receiving the production scoring model.

TABLE I. Receipt levels in the model and current release.

Level	Symbol	What is bound	Public?	Note
Per-verdict receipt	r_i	$\text{claim}_i + \text{refs}_i + \text{verdict}_i + \text{metadata}_i$	Optional	Not exposed in public v1
Public bundle receipt	R_B	manifest digest + checksum-set digest + release metadata	Yes	Current public surface
Private deployment	$\{r_i\}$ and R_B	individual verdict rows plus bundle receipt	Optional	Partner-facing option

A. Worked Appeal Example

The appeal path is intended to be operational rather than decorative. Table II shows a compact synthetic example aligned to the public bundle. The example is illustrative rather than evidentiary: it demonstrates the contestability fields that a released verdict exposes.

TABLE II. Worked synthetic appeal example.

Field	Example
Claim subject	Reward completion for campaign X
Verdict	Block
Traffic-quality code	device_farm
Adjudication meta-code	partner_disposition_disagrees
Evidence references	Burst timing cluster; postback mismatch
Appeal path	https://example.invalid/appeal-review
Possible outcomes	Uphold, override, request more evidence

B. Appeal Workflow Boundary

The released public capsule exposes enough information to locate and contest a decision, but not a full private review ledger. Each public verdict row carries a claim subject, evidence-reference classes, reason codes, and an appeal path. The synthetic appeal sample shows the same block-decision surface in single-record form. The public packet does not release the private reviewer state machine, internal evidence joins, or override ledger; those remain part of diligence or future work rather than the public v1 capsule.

IV. EVIDENCE TIERS

We separate three evidence tiers.

A. Deterministic conformance vectors

The first tier is a deterministic test suite. Fraud or dispute patterns are encoded as transformations over clean

postback-like events. The role of this tier is conformance: whether an implementation emits expected reason codes and stable verdict shapes on known perturbations. It is not evidence that live traffic is solved.

B. Public conformance and replay artifact

The second tier is replayable benchmark generation. A fixed seed produces event and verdict bundles, manifests, and checksums. This tier tests determinism, schema conformance, reason-code coverage, and receipt binding. The current public capsule includes this tier. Public adtech datasets may also be processed when licensing permits, but their role is calibration and stress testing rather than customer-specific lift measurement.

C. Design-partner offline pilot

The third tier is a partner-approved aggregate pilot. The intended first pilot processes anonymized historical rewarded-action events offline and emits per-event verdicts to the partner, together with an aggregate fraud-lift report and appeal sample. Public release of this tier is restricted to approved aggregate metrics. Raw partner events and per-event details remain private. This separation mirrors the caution required in online controlled experiments, where measurement discipline and operational boundaries matter as much as headline effect sizes[6].

V. PUBLIC CONFORMANCE AND REPLAY ARTIFACT

The public evidence bundle is a 10,150-event deterministic synthetic conformance artifact, not partner traffic. It contains 10,150 rewarded-action claims generated from a fixed public seed and processed by a controlled internal adjudication implementation that emits the released public verdict-envelope schema for this study. The app, campaign, partner, and device identifiers visible in the released rows are synthetic placeholders generated for this conformance artifact rather than redacted customer identifiers. Each output row carries a claim, verifier metadata, top-level verdict label, reason codes, evidence references, and an appeal path. The released benchmark

manifest, checksum records, and publication receipts provide the public receipt surface for the bundle; the public per-event envelopes do not expose a separate receipt hash.

TABLE III. Public 10K conformance artifact outcome.

Metric	Value
Events observed	10,150
Verdicts emitted	10,150
Allow verdicts	8,694
Hold verdicts	135
Block verdicts	1,321
Block rate	13.01%
Synthetic payout amount associated with blocked benchmark claims	USD 7,905.76

The released reason-code histogram mixes traffic-quality signals and adjudication-meta codes. Traffic-quality codes include replay detection (300), device farm (279), geo mismatch (200), click flooding (189), conversion hijacking (173), cookie stuffing (153), SDK spoofing (142), and click injection (88). A distinct meta-code, partner-disposition disagreement (1,321), records disagreement between the released adjudicator and the upstream partner disposition rather than a traffic-quality property. These counts are occurrence counts rather than mutually exclusive classes; a single blocked verdict may carry multiple reason codes. In the released bundle, all 1,321 block verdicts carry the meta-code by construction, so the traffic-quality codes should be read as additional, non-exclusive evidence labels on those same blocked rows rather than as alternative outcome classes. Within that structure, 1,256 blocked rows carry exactly one traffic-quality code, 65 carry multiple traffic-quality codes, and the average block verdict carries 2.05 total reason codes. The sample block verdict used for the public appeal envelope cites a device-farm burst together with a partner-disposition disagreement, then exposes the appeal path as part of the same verdict object.

The benchmark was re-generated after pinning the canonical v1 manifest timestamp. A fresh replay now matches README, manifest, event, and verdict files byte-for-byte. Targeted internal implementation tests covering benchmark generation and released-bundle validation passed with 60 tests and one non-blocking dependency warning. The shipped public validator additionally checks row shape, verdict totals, and reason-code histograms against the released manifest, and it validates that the synthetic appeal sample is itself a released block verdict with an appeal path. The appeal sample is therefore byte-bound by the manifest/checksum surface and shape-validated as a released verdict object, even though the private appeal state machine is not part of the public capsule. These results validate deterministic generation, schema conformance, reason-code emission, manifest and checksum stability, and appeal-envelope construction. They do not validate live fraud-detection ac-

curacy, lift, or production economic impact.

A. Public Verification Procedure

Third parties can verify the released public surfaces directly with the exact commands in Listing 2. Deterministic regeneration from the public seed remains part of the controlled-review surface rather than the public release. The public paper therefore exposes a real verification interface now and defers the full generator interface to diligence or later staged release.

Listing 2. Exact public verification commands.

```
cd invariant-systems-research
sha256sum -c \
  3-proof/3d-evidence/adtech/09-receipted-actions/benchmark_v1_10k/SHA256SUMS
python \
  3-proof/3p-deliverable/publications/arxiv/09-receipted-actions-adtech/anc/validate_public_verdict_envelopes.py \
  --schema \
  3-proof/3p-deliverable/publications/arxiv/09-receipted-actions-adtech/anc/verdict_envelope_public_release_v1.schema.json \
  --manifest \
  3-proof/3d-evidence/adtech/09-receipted-actions/benchmark_v1_10k/benchmark_v1_10k.manifest.json \
  --verdicts \
  3-proof/3d-evidence/adtech/09-receipted-actions/benchmark_v1_10k/benchmark_v1_10k.verdicts.jsonl \
  --appeal-sample \
  3-proof/3d-evidence/adtech/09-receipted-actions/benchmark_v1_10k/appeal_sample.synthetic.json
sha256sum -c \
  3-proof/3p-deliverable/publications/arxiv/09-receipted-actions-adtech/SHA256SUMS
just validate
```

B. Artifact Access and Fixed Public Identifiers

The current release is repo-native rather than DOI-backed. Table IV lists the release identifier, artifact pointers, and fixed benchmark digests that a reviewer can check directly. A later archive deposit can attach a public URL without changing these benchmark digests.

VI. CLAIM GATES

Publication is governed by explicit gates. The gate register is included as an ancillary file, but Table V summarizes the release-relevant result for the PDF reader.

Table V uses reader-facing pass/deferred summaries; the more specific internal gate-state tokens remain available in the ancillary gate register. The contract-surface gate checks that claim, evidence, verdict, receipt, and appeal fields are documented and schema-stable. The public-benchmark verification gate checks that the released bundle and packet match their published checksums. The public synthetic appeal-sample gate checks that a released verdict includes an appeal path. The disclosure-boundary gate checks that partner identifiers, raw evidence, thresholds, source paths, and live operational endpoints are absent from the public surface. The

TABLE IV. Repo-native artifact pointers and fixed public identifiers.

Field	Value
Release state	public-benchmark-ready (2026-05-03)
Packet path	3-proof/3p-deliverable/publications/arxiv/09-receipted-actions-adtech/
Benchmark path	3-proof/3d-evidence/adtech/09-receipted-actions/benchmark_v1.10k/
Benchmark schema	kaleidos.public_benchmark.v1
Verdict-envelope version	kaleidos.verdict_envelope.v1
Seed	3571126456339198
Manifest SHA-256	740fd321893b9f8b63cd12283ca44d65cfac108f3a138f3b5c8a39faf79d4906
Events SHA-256	875f2a4432b90fac52481b50785f9e342a7f6400ebaa419bbc0858cd79df58c6
Verdicts SHA-256	3456c71238ead22b0636b8033d2d2299101e2032c07f7b6544dedbc35ffd7fa7
Appeal sample SHA-256	39054daff1402a502dc0bfd00ddd8244b184616a843f882483727fab0a0cc8c
Schema file	anc/verdict_envelope_public_release_v1.schema.json
Validator script	anc/validate_public_verdict_envelopes.py
Validation transcript	3-proof/3d-evidence/adtech/09-receipted-actions/validation/public-benchmark-validation-2026-05-03.txt

TABLE V. Gate summary for the current public capsule.

Gate	Check	Status
G1	Contract surface documented and schema-stable	pass
G2	Deterministic export verified against checksums	pass
G3	Partner aggregate pilot approved for release	deferred
G4	Appeal surface present in released block sample	pass
G5	Disclosure boundary excludes private internals	pass

design-partner aggregate gate is explicitly deferred from the first public version and reserved for a later revision if partner-approved aggregate evidence becomes available. Failure of a future partner gate is part of the result. For example, a design-partner pilot may show low fraud lift; that would be a valid commercial finding and should not be converted into a positive claim by broadening the fraud layer after the fact.

VII. RELATED WORK AND POSITIONING

This paper sits at the intersection of invalid-traffic guidance, decision logging, provenance, contestability, transparency systems, and disciplined experimentation. Invalid-traffic guidance provides taxonomy pressure for disputed rewarded-action claims, but it does not by itself define a typed decision envelope with verdict, receipt, and appeal surfaces[1]. Policy-decision systems such as Open Policy Agent emphasize machine-readable policy evaluation and decision logs[2]; the receipted-action capsule narrows that pressure to externally consequential commercial claims with an explicit release boundary.

Provenance-oriented work emphasizes traceability of

claims and evidence across systems[7]. Contestability literature argues that affected parties need routes to engage, challenge, and remedy algorithmic decisions rather than receive only explanations[8, 9]; the appeal surface in a receipted action follows that pressure. Transparency systems such as certificate transparency, in-toto, and SLSA show how released artifacts can be bound to verifiable records and policy gates without publishing every private input[3–5]. Finally, controlled-experiment methodology matters for any later partner aggregate evidence: if a future revision makes live-performance claims, that evidence must satisfy stronger causal discipline than a public synthetic benchmark can provide[6]. The paper is therefore positioned as a systems and auditability contribution rather than an ad-fraud performance benchmark or a claim of novel fraud-detection efficacy.

VIII. THREATS TO VALIDITY AND NON-CLAIMS

A. Operational Threat Model

The threat model is practical rather than cryptographically exhaustive. The relevant threats for the released public capsule are replayed or duplicated claims, partner disputes over blocked payouts, tampering with released bundle files, schema drift, public reason-code gaming, verifier compromise, and privacy leakage through over-disclosure. The corresponding public mitigation surfaces are evidence references plus reason codes, the appeal path, checksum and manifest verification, explicit schema versioning, private thresholding and evidence boundaries, deterministic replay, and synthetic data in place of partner traffic.

B. Synthetic Benchmark Limitations

The public benchmark is deterministic synthetic data. It is suitable for conformance, replay, schema inspection, and bundle verification. It is not customer traffic and should not be read as a proxy for production incidence, false-positive rate, or marginal economic lift.

C. Absence of Public Partner-Traffic Lift

No partner-approved aggregate pilot result is included in the first public version. Claims about blocked dollars, review savings, overturn rates, or deployment efficacy on live traffic remain reserved for a later aggregate-only revision.

D. Evidence Visibility and Privacy Boundary

The public capsule exposes evidence references rather than raw evidence. That keeps partner data, tuned thresholds, and customer-specific rules out of scope, but it also limits how much of the private adjudication context an external reviewer can inspect directly.

E. Bundle-Level Versus Per-Action Receipt Binding

The formal model allows per-action receipt identifiers. The public v1 release binds the benchmark at bundle level through manifests, checksum records, and publication receipts. Reviewers should therefore read the current artifact as a release-level receipt surface, not as a public per-verdict transparency log. A natural v1.1 extension is a Merkle tree over released verdict envelopes with the root committed in the manifest.

F. Adversarial Adaptation and Reason-Code Gaming

Any published reason-code taxonomy can be studied or gamed by adversaries. The purpose of the paper is to show inspectable decision structure, not to freeze an immutable fraud-defense rule set. Operational thresholds, rule mixes, and exception logic remain private partly for that reason.

The contribution remains narrower: a typed, replayable, appealable decision envelope with manifest-backed public receipt binding for commercial actions.

IX. DISCUSSION

Receipted actions are useful when the cost of an unexplained decision is high and an affected party may reasonably challenge the outcome. Rewarded-action payout review is one such domain. Other domains may adopt the same pattern later, but this paper does not evaluate them; the contribution here is the rewarded-action capsule and its public verification boundary.

X. CONCLUSION

We present a reproducible audit capsule for claim-based commercial adjudication and use rewarded-action payout review as the case study. Receipted actions are the organizing pattern: a typed claim, evidence references, machine-readable verdict, receipt boundary, and appeal path. The current public capsule includes a compact public-safe verdict-envelope schema, a clarified receipt-level model, exact public verification commands, a deterministic 10K conformance artifact, checksum evidence, and a synthetic appeal sample. The first public version explicitly defers partner-approved aggregate pilot evidence to future work; if aggregate evidence is approved later, it should enter as a follow-on revision with an aggregate-only disclosure boundary.

-
- [1] Media Rating Council, Invalid traffic detection and filtration guidelines (2023), industry guidance for invalid traffic taxonomy and filtration practices.
 - [2] Open Policy Agent Authors, Open policy agent documentation (2026), <https://www.openpolicyagent.org/docs/latest/>.
 - [3] S. Torres-Arias, H. Afzali, T. K. Kuppusamy, R. Curtmola, and J. Cappos, in *USENIX Security Symposium* (2019).
 - [4] OpenSSF, Supply-chain levels for software artifacts (slsa) specification v1.0 (2023), <https://slsa.dev/spec/v1.0/>.
 - [5] B. Laurie, A. Langley, and E. Kasper, in *RFC 6962* (2013).
 - [6] R. Kohavi, R. Longbotham, D. Sommerfield, and R. M. Henne, *Data Mining and Knowledge Discovery* **18**, 140 (2009).
 - [7] T. Lebo, S. Sahoo, and D. L. McGuinness, *Prov-overview: An overview of the prov family of documents* (2013), w3C Recommendation.
 - [8] K. Vaccaro, K. Karahalios, D. K. Mulligan, D. Kluttz, and T. Hirsch, in *Companion Publication of the 2019 Conference on Computer Supported Cooperative Work and Social Computing* (2019).
 - [9] D. Kluttz, N. Kohli, and D. K. Mulligan, *SSRN Electronic Journal* 10.2139/ssrn.3311894 (2018).

TABLE VI. Claims and non-claims for the current public capsule.

Claim	Supported in current paper?
The public bundle can be checksum-verified and replayed at released-artifact level	Yes
Released verdict envelopes conform to a public schema contract	Yes
Bundle-level receipt binding works for the released benchmark bundle	Yes
Public per-verdict receipt hashes are released	No
The synthetic conformance artifact demonstrates live fraud-detection efficacy	No
The current paper demonstrates production payout savings or partner lift	No
The public capsule releases the full private appeal ledger and review state machine	No