

Candidate-Evidence Discrepancy Triage

id: page.candidate-evidence-discrepancy-triage title: Candidate-Evidence Discrepancy Triage visibility: public candidate_visibility: candidate-only candidate_only: true audiences: - agent-orch maintainers - release reviewers reading_modes: - executive - engineer - ai claim_refs: [] source_refs: - source.naming-decision - source.documentation-boundary - source.platform-delivery-contract

Candidate-Evidence Discrepancy Triage

Purpose

This is a read-only candidate-evidence triage for agent-orch maintainers and release reviewers. It maps declaration drift, fixture rejection cases, and the supplied platform observations to the exact manifest, model, fixture, authority owner, reproducible check, and escalation owner. It distinguishes a shape rejection from a runtime observation and a product conclusion from a platform attestation. The page is candidate-only: it creates no approval, release pointer, promotion, publication, schedule, rollback, activation, or live-directory effect. Missing recurrence, product impact, or authority is kept unresolved rather than completed with plausible prose.

Evidence Matrix

The matrix is a traceability aid, not a resolution ledger. The manifest named in every row is `journeys/user_journeys_manifest.json`; the pinned comparison manifest is `journeys/candidate-evidence-handoff-manifest.json`; and the machine-readable model is `models/candidate_evidence_handoff_contract.json`. The human-readable contract at `docs/candidate-evidence-preflight-contract.md` remains a separate five-check declaration. The three-check manifest delivered with this packet is synchronized to the executable/model route, but that synchronization does not canonize or resolve the five-versus-three declaration discrepancy.

Discrepancy and observed fact	Manifest	Model
D-1 — Declaration cardinality and route drift: the human contract names five checks and five canonical journeys, while the executable handoff route names AC-1 through AC-3 and three journeys.	<code>journeys/user_journeys_manifest.json</code> ; compare <code>journeys/candidate-evidence-handoff-manifest.json</code> .	<code>models/candidate_evidence_handoff_contract.jso</code> including its three pinned journey declarations.
D-2 — Journey-name mismatch: <code>invalid-journey-name.json</code> substitutes a shortened name for a pinned journey name.	The exact <code>name</code> and <code>traces_to</code> values in <code>journeys/user_journeys_manifest.json</code> .	The model's <code>journeys</code> array pins the byte-for-byte name.
D-3 — Command allowlist mismatch: <code>disallowed-command.json</code> claims a command outside the exact allowlist.	The manifest allowlist contains only <code>python3 -B tests/check_documentation_smoke.py</code> and <code>python3 -B -m pytest -p no:cacheprovider tests/test_product_contract.py</code> .	The model repeats that exact allowlist and the journey command bindings.
D-4 — Missing observed result: <code>malformed-evidence.json</code> omits the required command observation, so a declaration cannot be mistaken for execution evidence.	The manifest's evidence schema requires <code>observed_result</code> for each command claim.	The model requires <code>command</code> , <code>exit_code</code> , and <code>observed_result</code> .

Discrepancy and observed fact	Manifest	Model
D-5 — Independent-review path gap: <code>missing-review-path.json</code> omits the required review-verdict path, while the manifest requires <code>code-reviews/content-review.verdict.json</code>.	<code>journeys/user_journeys_manifest.json</code> requires the path and accepted verdict shape.	The model repeats the same required review path and verdict constraints.
D-6 — Reader-gate command/result mismatch: the supplied platform observation records <code>python3 tests/check_documentation_smoke.py</code>, exit <code>1</code>, and <code>Operation not permitted</code>, while the pinned journey declares <code>python3 -B tests/check_documentation_smoke.py</code>; localhost access was denied.	The third journey in <code>journeys/user_journeys_manifest.json</code> declares the exact <code>-B</code> command, expected exit <code>1</code>, and AC-3 trace.	The model's third journey carries the same command and expected exit.

Discrepancy and observed fact	Manifest	Model
D-7 — Preserved-validator representation boundary: the authoritative records supply <code>python3 -m pytest -p no:cacheprovider tests/test_product_contract.py</code> as their command array and separately record the bytecode-disabled declared command, while the journey declaration uses the exact <code>-B</code> form.	The first two journeys in <code>journeys/user_journeys_manifest.json</code> declare the <code>-B</code> form and expected exit <code>0</code>.	The model carries those exact journey declarations.

Discrepancy Triage

Triage starts by classifying the row, not by assigning a repair. D-1 is a declaration conflict between the five-check human contract and the pinned three-check executable/model route. The delivery keeps the user manifest on the pinned three-check route because the executable preflight explicitly pins that shape, while leaving the human-contract conflict open for its owner. D-2 through D-5 are intentionally negative fixtures: they show how the preflight rejects altered identity, unallowlisted commands, missing observations, and a missing review path. Their rejection shape must not be reported as a live reader failure or a resolved product defect.

D-6 is the only supplied runtime-style discrepancy. One platform-owned reader-gate attempt exited `1` with `Operation not permitted` because the evaluator environment denied localhost socket access. The command in that record lacks the manifest's `-B` spelling, so the record must be preserved as observed rather than silently rewritten. One denied attempt does not establish recurrence, a retry count, an affected claim, a product harness defect, or a repair target. D-7 records the same kind of boundary for the five preserved validator results: each is authoritative captured evidence, but each command array and declared command must not be conflated with the manifest's reader claim. No discrepancy in this report is called resolved without current owner-supplied evidence that proves the relevant fact.

Ownership and Escalation

The author owns only this product-side packet: the candidate page, its registry identity, the candidate change record, and the reader-journey manifest. The author may describe evidence and unresolved gaps, but may not change `sources/authority.json`, private governance, the validator authority, the governing contract, or a platform-owned reader record. The agent-orch contract owner owns reconciliation of the five-check versus three-check declarations, exact journey identity, and allowlist changes. A change that would alter governing authority, routing, validator authority, activation, or publication policy requires the separate governance authority; it is not an authoring decision.

The platform evaluator owns the reader-gate execution, environment cause, and any permitted follow-up result. The system-validator authority owns the preserved product-contract record and its evidence hash. The independent release reviewer owns the verdict at the required review path; the presence of a path in a manifest is not a verdict. Lee's bounded human decision is whether to retain the unresolved environment-specific blocker and route a permitted follow-up evaluation. Lee is not being asked by this page to approve, promote, publish, schedule, activate, roll back, create or change a release pointer, or write a live directory. Until the named owner supplies the missing evidence, the candidate stays parked and the discrepancy stays open.

Acceptance Checks

- **AC-1 — Traceability:** A maintainer or release reviewer can start at this page and map every listed discrepancy to `journeys/user_journeys_manifest.json`, the pinned handoff manifest, `models/candidate_evidence_handoff_contract.json`, the named fixture or preserved artifact, the three public-safe source IDs, and an explicit owner without treating a declaration as a run result.
- **AC-2 — Reproducibility:** A reviewer can replay the exact non-empty commands declared by the user manifest, compare names and commands byte-for-byte, and preserve the supplied exit status, output, counts, and evidence hashes. All five preserved validator results are read as authoritative evidence and are not rerun or replaced in this handoff.
- **AC-3 — Authority ownership:** Each discrepancy has a named authority and escalation route; Lee and governance authority decisions are distinguished from author work, independent review, platform evaluation, and validator attestations. The reader retains the candidate-only boundary and leaves recurrence, affected claims, repair targets, approval, promotion, publication, scheduling, rollback, pointers, and live-directory actions unresolved unless separately proved.

This report is complete as a candidate-only triage record, not as a repair, review verdict, release decision, or publication authorization. A later owner may add evidence to the relevant route, but may not rewrite the original observation to make the matrix appear closed.

Provenance: Candidate packet candidate-evidence-discrepancy-triage (candidate-evidence-discrepancy-triage.json, step: step_03_author_triage_report)

Registry Source: `content/public/candidate-evidence-discrepancy-triage.md`