



DISPUTE PACK

Reliance Governance Layer - Governing AI Itself

World: Reliance Layer **Outcome:** ALLOW WITH BOUNDS **Evidence Pack ID:**
JSP-PUBLIC-20260302-RELIANCE_GOV-0001

1. What was claimed / requested

Model emits high-risk score; agent wants to lock an account.

2. What the system allowed vs blocked

Allowed

- Automation proceeds inside explicit bounds.
- Reversible mitigations permitted.
- Escalation wired if confidence drops.

Blocked

- Any irreversible action outside bounds.
- Any execution without required evidence or authority.

3. Why (governance rationale)

Judgement Spine governs the model-to-action boundary: capture reliance intent, verify provenance and chain integrity, and allow automation only with bounds (draft recommendation + require human confirm). If confidence shifts, escalation is structural.

This pack is designed to survive dispute: it links decision, authority, bounds, and integrity in one exportable object.

Evidence references

4. Evidence signals

- model_score
- model_version
- prompt_provenance
- data_lineage
- tool_call_trace
- confidence
- autonomy_threshold

5. Evidence artefacts in this bundle

policy_trace.ndjson	Which doctrine checks fired and why.
audit_export.ndjson	Execution-time event record.
provenance_chain.json	Illustrative provenance and chain integrity.
authority_contract.json	Named authority + bounds contract.
manifest.json + signature	Integrity references (tamper-evident).

6. Integrity verification

Run `verify_manifest.py` from the root folder and retain the output as part of the dispute file.

If the hashes and signature verify, the pack is unmodified since it was sealed.

Contact: founder@judgementspine.com