

BOARD VIEW

Reliance Governance Layer - Governing AI Itself

Suite: Reliance Governance Layer **World:** Reliance Layer **Outcome:** ALLOW WITH BOUNDS

What happened (execution intent)

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

Status quo	Organisations govern models, not reliance; 'the model said so' becomes invisible authority.
AI shift	Model outputs increasingly become automated inputs to downstream systems - turning probabilistic output into deterministic action.
Judgement Spine difference	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.
Impact	Prevents institutional liability from unowned reliance while keeping AI useful; proof exists before consequence.

Decision summary

Outcome	ALLOW WITH BOUNDS
Bounds enforced	<ul style="list-style-type: none">- disable auto-lockout- allow reversible mitigations- require human confirmation for lock- attach provenance + chain
Escalation path	If confidence drops or provenance incomplete, SAFE DEGRADE to monitoring-only + notify account owner.

Proof you can produce in 60 seconds

1. Open WOW_PACK/OPEN_ME.html (offline).
2. Select this scenario and click Replay (90 seconds).
3. Open Regulator View and Dispute Pack PDFs.
4. Run verify_manifest.py to confirm integrity (hashes + signature if available).

Evidence Pack ID: JSP-PUBLIC-20260302-RELIANCE_GOV-0001 **Control plane demos:**
<https://judgementspine.com/control-plane-demos>



JUDGEMENT SPINE

Note: This is a public, redacted, illustrative pack. The structural claim is about runtime authority, bounded autonomy, and proof written before consequence.