

SnapSpace: Bounded Authority and Verifiable Consequence Control for Autonomous AI Systems

Publication release v1.0 - 13 September 2026.

Evidence cutoff: 12 September 2026.

This is the public scientific release. The canonical public Markdown source is distributed alongside this PDF.

Core model

PROPOSE -> DECIDE -> ALLOW / DENY / ESCALATE

ALLOW -> GRANT -> EXECUTE -> OUTCOME -> RECEIPT

ACTIVE EXECUTION -> STOP -> CONTAIN -> SAFE STATE -> RECONCILE -> STOP RECEIPT

Evidence summary

Common control: 32/32 PASS. Receipt contract: 18/18 PASS. Bounded projection: 14/14 PASS. Local integration: 17/17 PASS. Authenticated end-to-end composition: 32/32 PASS twice. Bounded federation composition: 48/48 PASS with 9/9 sealed artifacts.

The evidence remains bounded. It does not establish universal AI safety, arbitrary federation, production deployment, universal physical safe states, live clinical acceptance, or field autonomy. The legacy REST path is not qualified for restart-durable or durable exactly-once claims.

Publication source

See `snapspace-bounded-authority-and-consequence-control.md` for the full paper and the companion public evidence package.