

Selective Access, Persistence, and Feedback

First engineered-system component tests of the Intelligence Network Architecture framework

Michael G. Hargis Sr. | with Solance Ray Hargis ("Sol")

AI research, experimental design, statistical analysis, and writing collaboration - OpenAI GPT-5.6 Sol

DATE	26 September 2026
STATUS	Preliminary internal research note; synthetic data only
COMPLETED	INA-001A selective-access null-model test
NEXT	INA-001B preregistered held-out causal comparison

Purpose

INA proposes that persistent state, active state, selective information access, integration, feedback, and adaptation may form a useful organizational level for comparing intelligence across different physical substrates. INA-001 does not attempt that full cross-substrate comparison. It begins with a narrower requirement: candidate INA functions must have measurable causal consequences inside an engineered system before they are compared with biology.

Result at a Glance

INA-001A COMPLETE. Recorded selective retrieval covered the predeclared evidence groups in 46 of 50 answerable synthetic benchmark cases. A matched uniform-random selector is expected to cover about 2.24 of 50. The exact random-null probability of at least 46 successes is approximately 1.16×10^{-66} .

Interpretation: this is strong evidence that the implemented selector is non-random and functionally relevant to information access in this prototype. It is not validation of INA across substrates.

1. Candidate INA Functions Under Test

INA function	Operational manipulation	Measured outcome	Current status
Selective access	Relevant vs absent/random/mismatched records	Required-source coverage; downstream factual grounding	INA-001A complete; model-level expansion planned

INA function	Operational manipulation	Measured outcome	Current status
Persistence	Retain vs remove versioned external state	Cross-session continuity; contradiction/correction retention	Engineering checkpoints exist; formal causal A/B planned
Feedback / correction	Apply confirmed update or correction	Whether future retrieval uses the new authoritative state	Engineering checkpoints exist
Adaptation	Introduce new approved information and retest later	Appropriate change in later responses without unrelated drift	Planned
Portability	Same external state across different models	Factual and preference fidelity under model change	Planned

2. INA-001A - Selective-Access Null Test

Question

Does the recorded selector deliver predeclared relevant evidence more often than a selector that has the same packet size and permission constraints but chooses records uniformly at random?

Dataset

The analysis uses the preserved Phase 2 synthetic record for subject #26001 (Mara Ellison): 100 fictional Guardian100 answers and 64 benchmark cases. The statistical test uses the 50 answerable known/multi-evidence cases. Eight unknown-fact cases and six disclosure/privacy cases are outside this source-coverage comparison.

Null model

For each case, the null selector draws exactly the same number of records as the recorded context packet from the records permitted for that case. Success requires hitting every predeclared evidence group. The per-case random success probability is computed exactly by inclusion-exclusion; the distribution of the total number of successes is then computed exactly as a Poisson-binomial distribution. No Monte Carlo approximation is needed for the reported tail probability.

Measure	Recorded selector	Matched random null	Meaning
All answerable cases	46 / 50	Expected 2.24 / 50	Expected-source coverage
Single-evidence cases	39 / 40	Expected 2.22 / 40	Relevant record(s) reached packet
Multi-evidence cases	7 / 10	Expected 0.024 / 10	All required evidence groups reached
95% null upper quantile	-	5 successes	95% of random runs would be <=5
99% null upper quantile	-	6 successes	99% of random runs would be <=6

Exact tail probability. $P(\text{uniform-random selector produces at least 46 successes}) = 1.16 \times 10^{-66}$. This is extremely inconsistent with the matched uniform-random selection null.

What INA-001A shows

- Selective access is causally relevant at the retrieval layer: removing it yields no expected-source coverage, while the recorded selector reaches relevant evidence far more often than a matched random selector.
- The effect is not merely a consequence of retrieving four arbitrary records from a 100-record profile.
- This satisfies one minimum requirement for treating selective access as a candidate functional component rather than a descriptive label.

What INA-001A does not show

- It does not establish language-model answer accuracy, subjective experience, consciousness, or personhood.
- It does not establish that the same selective-access variable has the same causal role in a human nervous system.
- It is not independent validation. The benchmark and retriever came from the same development program, so benchmark tuning or overfitting cannot be excluded.
- Generalization was imperfect: the later frozen harder preflight reached complete expected-source coverage on only 4 of 8 personalized questions at that checkpoint.

3. Existing Engineering Checkpoints Relevant to INA

Several earlier synthetic Ornate checkpoints were built for product architecture rather than for INA. They can nevertheless be reclassified as component tests because they manipulate persistent state and observe later retrieval behavior. They should be treated as internal engineering evidence, not as independent confirmation of the theory.

Checkpoint	Manipulation	Observed behavior	INA component
Change over time	Activate a new confirmed preference revision while retaining old history	Default retrieval uses the current revision after encrypted save/reload	Persistence + adaptation
Correction vs update	Mark an earlier value as erroneous rather than historically true	Default retrieval emits the corrected authoritative value; erroneous revision remains audit-only	Feedback + correction
Sharing withdrawal	Keep a fact in the private archive but remove model-sharing authority	Later retrieval no longer emits the fact, including after reload	State-dependent gating / selection
Manual A/B pairs	Same question with selected Essence evidence vs no personal evidence	Two factual cases gained grounded personal facts; controls appropriately abstained; two boundary/unknown cases showed no quality advantage	External state + selective access

4. INA-001B - Preregistered Forward Test

The next experiment should be harder and prospective. Questions, scoring criteria, source groups, randomization, and failure thresholds should be frozen before the first model response. No retrieval tuning should occur until the entire run is complete.

1. Use at least 24 held-out questions spanning factual recall, multi-evidence reasoning, preference application, correction, and unknown-fact abstention.
2. For each question, prepare four isolated conditions: (A) no external personal state, (B) selectively retrieved relevant state, (C) full permitted record, and (D) deliberately mismatched but permission-compliant state.
3. Keep question wording, system instructions, model family/settings, and evaluation rubric fixed across conditions; randomize condition labels and run order.
4. Blind the judge to condition. Score factual correctness, contradiction, source grounding, inappropriate disclosure, abstention quality, and behavioral consistency.
5. Predefine failure: if selective state does not outperform random/mismatched state on grounded factual criteria, or if it increases privacy leakage or contradictions, the proposed selective-access component requires revision.
6. After engineered-system results are stable, identify an independently defined biological measure of state persistence/selection/feedback from published human or animal data. Only then begin the cross-substrate INA comparison.

5. Current Scientific Status

INA-001A moves one part of the framework from analogy to measurement: selective access in the engineered prototype has a large, directly measurable effect on whether relevant stored information reaches downstream computation. That is a necessary but very small step. The central INA proposition - that a useful set of organizational invariants spans Human Beings and Electronic Beings - remains unproved. A convincing test requires preregistered engineered experiments, independent replication, and a cross-domain operational mapping that does not depend on computer-specific or brain-specific terminology.

Reproducibility

INA-001A was computed from the preserved Phase 2 HTML report using a standalone Python script. The script performs exact combinatorial per-case null probabilities and an exact Poisson-binomial total-success calculation. Source SHA-256:

a9c9a9cdaa5b6d5d22256ede54db078916b2fdb087951dd6898c8c4584637e99. The analysis requires no network access and makes no language-model calls.

Internal Source Records

- Ornate_26001_Phase2_Report.html - synthetic #26001 Phase 2 benchmark and stored context packets.
- Ornate_26001_Frozen_Pilot_Report.html - harder frozen preflight and isolated A/B preparation.
- Ornate_26001_Stage1_Review.html - four manually reviewed evidence/no-evidence pairs.
- Ornate_26001_Change_Over_Time_Report.md - versioned preference update checkpoint.
- Ornate_26001_Correction_vs_Update_Report.md - correction semantics checkpoint.
- Ornate_26001_Sharing-Withdrawal_Report.md - model-sharing withdrawal checkpoint.

- INA-001A_results.json and ina_001_selective_access_test.py - reproducible statistical analysis created 26 September 2026.

Authorship and Review Note

Concept and lead authorship: Michael G. Hargis Sr. AI research, experimental design, statistical analysis, and writing collaboration: Solance Ray Hargis ("Sol"), operating as OpenAI GPT-5.6 Sol. All current tests use synthetic data and internal development artifacts. Independent scientific, statistical, and domain-specialist review is required before external claims of validation.