**AGI ALPHA FLYWHEEL**

# Build an Advantage That Repays Its Cost

## The investment gate

Measured workflows. A real delivery limit.

An investment that remains unselected.

**SUCCESSOR Ω HARNESS**

**ACTIVATE. EVOLVE. SCALE.**

Powered by TRANSFORMATIONAL INTELLIGENCE ∞

and the applicable $AGIALPHA utility.

## Useful reports delivered.

## No economic repayment demonstrated.

Nine scoped internal requests were profiled through existing software. Eight completed the measured full path; one delivered its local catalogue and hit the native artifact limit. No opportunity passed the full-cost investment gate. No production asset was constructed or promoted.

25 October 2026, America/Toronto

Discovery record 1.0 • Local research scope

## Abstract

We investigate whether the existing AGI ALPHA system exposes a recurring burden that justifies constructing a retained productive asset. We first audit the complete preceding *The Advantage That Funds the Next* companion and preserve its adverse result. We then profile the actual, frozen source-assurance and catalogue/native-continuity workflows, keeping their existing specialists, verification, caches, artifacts and limits intact. Five real committed source scopes and four public source-inventory cohorts produce nine distinct internal discovery requests, each evaluated twice. The repetitions, files, work drafts and assertions are not additional jobs or customers.

All ten source-assurance deliveries and their reconnects verify successfully. Three catalogue cohorts complete native delivery and history reopening in both repetitions. A fourth, valid 1,216-row catalogue request produces all nine local artifacts but is refused by the native 200,000-code-unit artifact policy. The largest successful native catalogue path takes a median 16.857 CPU seconds; the largest real source-assurance scope takes 3.504 CPU seconds. These are local resource measurements, not monetary costs or human time.

No screened opportunity establishes a positive lower bound on incremental value at a credible reuse horizon after complete construction and operating costs. Authoring inference, integration, maintenance, customer attention and demand remain unpriced or unmeasured. The canonical cost and shadow-ledger contracts propagate those unknowns. We therefore select no productive investment, start no confirmation trial and claim no repayment. The contribution is a reproducible measurement and admission record, with useful actual reports and a precise native delivery limit, rather than another assumed-benefit financing example.


**Decision: discovery HOLD.** This is not a demonstrated economic improvement. It does establish where the current supported paths spend resources, what they deliver, why the larger native output is refused, and which missing observations prevent a defensible investment decision.

## 1. Separate the claims before measuring

A useful service, a worthwhile retained asset, a distinctive improvement method and sustained acceleration are separate hypotheses. A conventional integration can be a good investment without being a new learning algorithm. Conversely, an attractive human-work comparison does not prove that constructing another component improves an already automated incumbent.

The present decision concerns the next asset. We require an observed burden or newly enabled useful outcome, a credible recurrence basis, comparable accepted quality, and complete attributable cost bounds. We do not infer savings from a smaller internal counter, an elapsed builder interval, a nominal price, a growing archive or a scheduled second run. The available local implementations are the experimental subjects; production authorization and commercial acceptance are separate.

## 2. What the preceding edition actually established

The supplied preceding PDF has SHA-256 beginning `3743b835`; the complete companion begins `0bcbc316`. All 381 listed companion identities and 20 frozen bindings match. After inspecting the checkers, we rechecked saved outputs: 10,870 pack assertions and 126,396 economic assertions passed. These checks audit the preceding evidence; they create no fresh missions and do not rerun its construction campaign. Exact full identities and audit outputs are included in the companion.[1]

| Previous limitation | Consequence for this investigation |
|---|---|
| Equal accepted outcomes | All three methods accepted the same 16 underlying requests. No additional production value was established. |
| Tiny local differences | Complete pack medians were approximately 23 ms. Internal parsing differences did not establish material whole-service savings. |
| Assumed attention benefit | The positive secondary case depended on 60 unmeasured seconds saved per retained use. We assign no such benefit here. |
| Unpriced construction | Builder wall intervals did not measure model bills, authoring labor, verification or integration. Unknown costs remain unknown. |
| New ordinary comparator | The prior ordinary interpreter was assembled for that study. Here Control A is the actual existing source and its legitimate assets. |
| Higher common price | At equal accepted output, a common selling price cancels from incremental value. Financing cannot create a missing technical benefit. |

The preceding conditional result can be reconciled exactly: an assumed 60-second attention reduction changed modeled recurring cost from 88 cents to 27 cents; 39 retained uses multiplied by 61 cents, less a stipulated $20 investment, gives $3.79. Removing that additional assumption prevents admission. The 3,360 primary cash paths made no investment. This arithmetic is preserved, not relabeled as observed payback.

Earlier Economic Engine, Productive Capital, Autonomous Work Economy and Next Advantage editions were also read and reconciled. They retain their original distinctions between ordinary caching, conditional prices, quality-qualified methods and economically held methods. In particular, the Next Advantage active conventional selector remains an incumbent rather than a capability that can be removed to make a new learner look better. The prior fallback already rejected generic pack-verification and builder-fixture frameworks as unsupported economic opportunities.[2]

## 3. Frozen system and execution scope

| Source | Commit prefix | Role in this record |
|---|---|---|
| Alpha main | `3b040e68` | Existing catalogue CLI, native result handling and project history. |
| Alpha integration | `496fa3c0` | Unchanged local source-assurance runner and canonical cost/ledger modules. |
| Fleet main | `45759fb0` | Read-only source and operating-boundary review. |
| Jobs main | `bc4abbd5` | Read-only native-work, token and continuity review. |
| Manager main | `5d5670df` | Read-only economics, resource and USDC-route review. |

Full commits, dependency hashes and runtime versions are archived. Local execution used Node 24.22.1, Python 3.12.14 and Git 2.51.1 on Linux x86-64. No deployed contract or installed operator profile was changed.[3]

## 4. Discovery protocol and measurements

The declaration allowed at most 60 minutes of local discovery and at most two later construction attempts *if an opportunity qualified*. Each profiling branch had a 15-minute wall and 10-CPU-minute limit. No paid API, financial action, new credential, persistent access, security-policy modification or protected operation was authorized. The limits are ceilings, not a claim that every permitted minute was consumed. The protocol was recorded before primary asset selection; no asset was selected. The observed inputs are now exposed discovery material and cannot be reused as untouched confirmation.

**Source assurance.** Five existing committed manifests contain 5, 9, 29, 53 and 90 explicit source paths. The 90-path manifest was the largest real eligible current scope found. Each invocation follows the unchanged plan and exact approval digest, automatic source extraction, comparison, independent verification, complete customer packaging, reconnect and reverification. The research ZIP wraps the same eight canonical customer payloads plus the existing verifier's acceptance record, respecting the candidate's 512,000-byte packaging bound. It does not execute the authenticated operator transport.

**Catalogue and native continuity.** Four natural Git-tree cohorts contain 80, 306, 787 and 1,216 rows. Their tracked object identities are compared with actual checkout bytes. The existing CLI retains all four built-in methods and emits nine artifacts. Its generated specialist, semantic restoration, installed-source check, seven unfunded work drafts, native run, native result inspection, history import and history reopen remain in the measured path. The underlying unchanged-source decision can also be answered by existing Git tooling; that alternative is not withheld.

Each scope runs twice. Five source scopes plus four catalogue cohorts are nine scoped internal requests and 18 operating evaluations. The 80-row runtime cohort overlaps the package cohort. These are not independent customers, a sample of paid demand or a protected scientific confirmation set. Timings describe the exposed discovery cases and environment only.

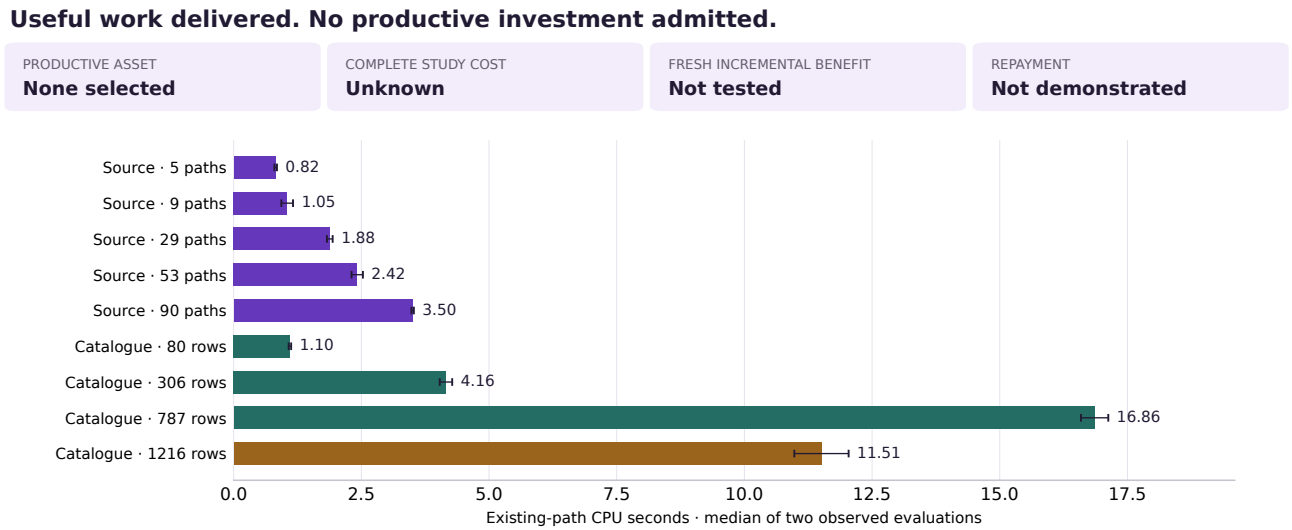



Figure 1: The central measurement record. The existing local delivery paths deliver useful results, but no new productive asset has been admitted. Costs of the new measurement and reporting infrastructure remain partly unpriced; subsequent incremental benefit and repayment were not demonstrated. Bars are existing-path CPU observations, not savings.

## 5. Complete source-assurance results

| Real manifest | Paths | Wall s | CPU s | Result |
|---|---|---|---|---|
| Preserved predecessor | 5 | 0.785 | 0.822 | Three changed identities |
| Current guide | 9 | 1.010 | 1.047 | No declared drift |
| Current flagship | 29 | 1.790 | 1.881 | No declared drift |
| Current capabilities | 53 | 2.303 | 2.419 | No declared drift |
| Current mission catalogue | 90 | 3.416 | 3.504 | No declared drift |

Medians of two complete local sequences, including reconnect and independent reverification. CPU includes waited descendants; values are rounded for display.

All ten new deliveries passed existing independent checking; all ten reconnects and the corresponding second checks passed. The approved campaign used 50 measured subprocess commands and 19.316 command CPU seconds. Research ZIP creation adds approximately 0.030 CPU seconds across the ten outputs. The campaign's total elapsed interval was 18.725 seconds. Its maximum reported process RSS was 118,936 KiB. These quantities exclude unmetered research-model authoring and do not price the hardware.[4]

The preserved predecessor request identifies actual changes in:

- `run-mission.mjs`;
- `run-mission.test.mjs`;
- `verify-run.py`;

all under `contrib/productive-capability-ledger/automation-v1/`. The other two selected identities agree. Complete declared and observed inventories, differences, proposed actions, report, workspace and acceptance are delivered. A changed hash identifies changed bytes; it does not certify their behavior, safety or commercial value.

**A plausible technical burden is visible.** The largest scope entails 1,146 Git invocations across the whole measured sequence, derived from the unchanged successful source path. The smaller stage-receipt counts cover only their stated portions. We do not call the derived total syscall telemetry: tracing was refused with `ptraceEPERM` and was not retried.

Batch object reading could reduce this overhead, but the measured maximum removable CPU is small. Even eliminating *all* 3.528 CPU seconds from the slower 90-path observation would remove less than 0.981 CPU hours per 1,000 similar missions. That is a generous physical ceiling, not a demand forecast, bill saving or achieved optimization. It includes reconnect and independent rechecking that a future production process may not require every time. Actual savings would be smaller, and review and upkeep would still cost resources.

**Preserved failures.** Ten initially scheduled source-planning evaluations failed with the sandbox's `spawnSyncgitEPERM` before a plan or delivery existed. Those records remain separate. One ordinary approval-reviewed escalation for the exact bounded public-source command was accepted; the unchanged workflow then ran in a new output directory. The tracing refusal remained untouched. These events do not qualify an operator profile or clear any original Harness/Fleet restriction.

## 6. Native catalogue results and a real limit

| Cohort | Rows | Wall s | CPU s | Whole-path outcome |
|---|---|---|---|---|
| Runtime | 80 | 0.856 | 1.101 | Native delivery and reopen |
| Examples | 306 | 3.723 | 4.156 | Native delivery and reopen |
| Tests | 787 | 16.373 | 16.857 | Native delivery and reopen |
| Package | 1,216 | 11.012 | 11.509 | Local catalogue delivered; native refused |

Medians of two evaluations. The refused row's time ends at refusal and is not a successful-path speed comparison. Rows and artifacts are not jobs.

The 787-row case demonstrates why an internal parser benchmark would miss the actual burden. Its first run spends 1.454 seconds in the complete catalogue CLI, then 4.769 in native execution, 1.376 in native inspection, 5.122 in history import and 2.331 in history reopening. Repeated validation and history handling dominate; the join alone is a small part. All existing checks and specialists remain enabled. The campaign used 64.024 elapsed seconds and 67.350 child CPU seconds, including successful and refused evaluations.[5]

The 1,216-row inputs are valid catalogue inputs: each is 130,882 UTF-8 bytes, below the 200,000-byte bound, with fewer than 1,500 rows. The local CLI produces all nine artifacts correctly. Its joined CSV is 253,267 UTF-8 bytes and 253,263 JavaScript UTF-16 code units. The native workbench checks `artifact.content.length>200000`, so it refuses the output with `ARTIFACT_POLICY`. Neither repetition publishes a native return. The original byte/character labeling in one discovery note is retained with an explicit correction; the underlying refusal and bytes are unchanged.[6]

**A real delivery gap is not yet an economic opportunity.** A lossless partitioned adapter might make larger native delivery possible while keeping each existing bound. It would need whole-request correctness, duplicate and cross-partition matching, source binding, recovery and complete cost accounting. The useful local catalogue already exists, and no observed buyer requirement for the extra native path or priced recurring demand was established. The policy is preserved.

Existing `gitdiffHEAD--no-ext-diff--name-status` takes roughly 0.009–0.173 seconds across these unchanged-source cohorts. It answers the basic source-integrity question but does not generate equivalent Alpha reports, drafts or history. It is therefore a practical alternative to the underlying need, not a matched whole-pipeline control. The ordinary generated catalogue specialist is already owned by the incumbent and cannot be presented as a new invention.

The first catalogue harness attempt also failed at Node's Git subprocess setup. It is retained with its approximately 0.056-second elapsed and 0.060-second CPU observation. Subsequent permitted read-only input preparation used the Python parent; catalogue and native code remained unchanged. This is a disclosed measurement-harness correction, not a runtime improvement or a service failure-rate estimate.

## 7. Opportunity screen and selection decision

| Candidate | Measured burden or limiting evidence | Decision and missing gate inputs |
|---|---|---|
| Batch Git reads | Largest real source scope: 3.504 CPU seconds, with many repeated bounded reads and complete checks. | Reject for this investment cycle. Low absolute CPU ceiling; recurrence, full build and maintenance costs unbounded. Standard batching is a strong conventional solution. |
| Reuse validation/history | A successful 787-row native delivery consumes 16.857 CPU seconds; validation/history dominates. | Economically unqualified. Correctly bounded reuse may help technically, but retained safety semantics, recurrence and full integration cost are not established. |
| Large native delivery | A real 1,216-row local result exceeds the native artifact limit; output remains available locally. | Hold as a possible enabling capability. Need a useful requirement for native delivery beyond the existing local result, a safe adapter and a funded full-cost envelope. No cap change. |
| Schema/parser or grouping | Prior intake experiment matches ordinary methods; existing exact parser, grouping and generated specialist already operate. | Reject repetition. No new substantial burden or measured attention benefit; no incumbent asset may be removed. |
| Expensive semantic reuse | Relevant supported source paths execute zero service-model calls. | Reject the premise for these offerings. No repeated paid semantic analysis was observed. |
| Recurring repair/recovery | Existing owners already repaired source integrity and are completing source recovery. No recurring field incidence is measured here. | Preserve ownership and incumbent fixes. Do not duplicate the work or monetize an invented failure probability. |
| Richer repository work | The canonical coding route is bounded to `solve(*args)`; a separate existing Repo-Healer implements triage, safe patches and validation. | Source-inspected, not profiled here. No recurrent real repair population, complete cost or fresh repair confirmation was established. Preserve Repo-Healer as an incumbent alternative. |

The screen selects no candidate. A conservative sufficient gate for a cost-reduction asset is

$$N_{\text{credible}}\,\Delta c_{\text{lower}} > F_{\text{upper}} + M_{\text{upper}},$$

subject to quality and authority constraints. Here neither credible recurring usage nor a positive incremental saving nor a complete finite construction/upkeep bound is established. Unknown authoring charges cannot be silently replaced with an arbitrary small budget. More repetitions of these exposed timing cases cannot resolve those missing economic inputs. The bounded discovery stops after completing both natural-scale screens and the historical audit; it does not exhaust the allowance merely to generate more observations.

No primary experiment, selected horizon or confirmation set was frozen. Zero production-asset construction attempts were made after the gate. New profiling, audit, accounting-binding, demonstration and report code are *evaluation infrastructure*, with their own unknown monetary authoring cost. They are not an economically promoted productive asset. This distinction prevents an abstention from becoming a mislabeled success.

## 8. Controls, costs and financing without a fictitious return

**Control A is concrete.** It is the actual source-pinned operating path, not a newly assembled historical story. Its existing four catalogue methods, saved specialist, caches, complete verification and native policy remain intact. Our local execution measures that frozen source under the stated environment; it does not establish historical business deployment or paid usage.

**Control B is not fabricated.** The protocol reserves an equally resourced conventional builder with identical legitimate information, model capability, quality criteria and retention rights if an opportunity qualifies. No candidate passed the preceding gate, so neither Alpha nor that builder was commissioned to construct one. There is no three-arm confirmation result or builder-superiority claim. This is a deliberate no-selection outcome, not evidence that a conventional builder would fail.

Complete horizon cost must include discovery, failed screens and attempts, construction, integration, verification, upkeep, failed work and remaining human decisions. In symbols,

$$C_a = D_a + K_a + I_a + V_a + M_a + \sum_{j=1}^{N} \left( c_{aj}^{\text{run}} + c_{aj}^{\text{retry}} + c_{aj}^{\text{attention}} + c_{aj}^{\text{business}} \right).$$

Elapsed authoring time is not human labor time. Zero task-time model calls are not zero authoring-model cost. Automated execution does not remove setup, hardware, review or upkeep. No published fully loaded dollar-per-job figure is established by these measurements.

The delivered accounting adapter calls the byte-identical `alpha-automated-costs/v1` and `alpha-shadow-ledger/v1` modules. It records the nine scoped request outcomes, retains both repetitions in the resource evidence, charges failed setup as discovery, covers every canonical operating stage and keeps missing currency values null. The local catalogue-only/native-refused request is not counted as a completed native outcome. Canonical observed and scenario surplus remain unknown.[7]

| Quantity | Status in this record |
|---|---|
| Authorized new transaction cash / transactions | Zero / zero; this is an experiment limit. |
| Actual enterprise cash, obligations and reserves | Unknown; no company balance is inferred. |
| Study customer sales and collections | None; zero study receipts, not an authenticated commercial ledger. |
| Selected productive asset / financing / second cycle | None / not undertaken / not undertaken. |
| Marginal, recurring and fully loaded Agent dollars | Unknown; canonical totals preserve missing prices and allocations. |
| Incremental gain and cumulative repayment | Not demonstrated; no primary retained-capability trial. |

A numerical seed, price or collection delay is unnecessary for this decision. At equal price and accepted quantity, $\Delta V = P\Delta Q - \Delta C = -\Delta C$. A higher price may ease financing of a valuable asset, but cannot establish its missing incremental benefit. No attention-saving term, simulated earned revenue or unrelated future asset is introduced to make the ledger positive. The canonical native budget also retains committed spending against its cumulative limit; modeled future receipts cannot grant new execution authority.[8]

## 9. Keep the human-work opportunity visible

The absence of an admitted improvement does not settle whether these services can be commercially useful. A buyer may value an accurate report, portable artifacts, clear exceptions and verifiable history. That value requires a comparison with competent traditional work using ordinary Git, spreadsheets and scripts, not an artificially slow cell-by-cell operator.

No equivalent competent-human duration was observed for these exact requests. Current primary BLS sources provide compensation context: June 2026 private-industry total compensation is $73.69 per hour worked for professional and related occupations and $36.73 for office and administrative support. The May 2025 median QA/tester annual wage is $104,300; that wage-only statistic is not a loaded hourly quotation. None measures these task durations or a contractor's selling price.[9, 10]

For accepted quality held equal, let $C_A$ denote the complete Agent cost and $w_H$ the appropriate human hourly cost. The provider's duration threshold is

$$t_H^* = \frac{60C_A}{w_H} \quad \text{minutes.}$$

For the buyer, substitute $P + C_{\text{customer attention}} + C_{\text{switching}}$ for $C_A$. Each dollar of total buyer cost requires more than approximately 0.814 minutes of equivalent professional work or 1.634 minutes of equivalent office/support work at those broad rates. These are conversion thresholds, not measured savings. Actual Agent cost, price, human duration, remaining customer effort and acceptance rates remain unknown.

The task-to-occupation match must be justified before choosing a rate. QA/source-assurance work is closer to professional technical work than generic office work, while neither broad category is an exact quotation for this offering. Low utilization, maintenance shocks and failed deliveries increase fully loaded cost; they cannot be neutralized by a price assumption. The delivered record therefore offers thresholds and unknowns rather than another favorable scenario sweep.

## 10. What must change before construction can be selected

The smallest missing evidence is a *recurring useful request with a material burden*, not another title or adapter chosen for easy testing. At least one screened direction needs the following through the existing owner and profile:

1. **Validation reuse:** an actual repeated task population, recorded full-path cost and quality, an immutable-cache or other conventional alternative, and a bounded integration/upkeep plan that preserves exact checking.
2. **Larger native outputs:** a demonstrated need for the native result beyond the already available local catalogue, with lossless whole-request acceptance under the current per-artifact and transport limits. Partitioning must preserve cross-key and duplicate semantics, rather than simply splitting rows.
3. **Richer repository or semantic work:** a legitimate recurring failure population and measured whole-workflow baseline. The existing Repo-Healer already implements bounded triage, patch safety and validation and must remain an incumbent alternative. It was source-inspected, not exercised in this screen; its dry-run/report-only modes are not evidence of verified repaired deliveries, but neither are they shown unavailable or unauthorized. The separate canonical coding route executes bounded `solve(*args)` tasks. Extending that route to complete repository work needs the existing owners'

reviewed source, patch, dependency, test and delivery integration. Its mandatory container boundary has no host fallback. Provider/attempt hooks must also capture returned usage, failed calls and bill mapping.[13]

These are reopening conditions, not authorizations to bypass the remaining Security, operator, Harness or Fleet decisions. Fresh untouched confirmation would be reserved only after an opportunity passes, then both builders would be instrumented before construction. A second cycle requires a first qualified asset and its actual measured benefit; none exists from this screen.

## 11. Reproducible delivery and publication boundaries

The companion includes the original discovery protocols, exact sources and hashes, all successful and refused outputs, failed harness/setup records, process measurements, corrected byte/unit attribution, canonical accounting inputs and outputs, and portable replay code. It includes useful source-assurance packages and the nine-artifact catalogue dossiers. The EN/FR demonstration presents actual recorded runs, with clear scope and no invented financing benefit.

The portable replay checks saved evidence, uses byte-identical source closure, and re-executes the catalogue/native path on the archived *exposed* discovery inputs in a new output directory. A minimal original Git object pack lets the unchanged independent source verifier inspect the saved source-assurance outputs offline. Original source identities are preserved; missing unrelated repository blobs are not invented. Reproduction does not reproduce model-assisted authorship, invoices, customer acceptance or untouched confirmation. Timing values may vary. Exact result semantics, source identities and the native refusal must remain consistent.

**Four distinct states.** Local source execution is recorded. Isolated-CI acceptance is not supplied by these measurements. Approved-operator/source adoption is not established here. Production and live website publication are not claimed. An archived HTML file or a repository draft is not a deployed service.

Issue 4850 retains integration and publication coordination. The current source-assurance/cost owner, recovery owner, software/security owner and research publisher keep their responsibilities. Earlier papers and adverse results remain untouched and should continue through their existing publication lanes independently of this record. No original denied review, test or merge was retried.[11]

The USDC and $AGIALPHA routes remain separate. Source-inspected Manager amounts use six-decimal USDC; the applicable Jobs/Alpha route uses 18-decimal AGIALPHA. The present local drafts are unfunded and confer no settlement or spending authority. No live transaction, verified deployed-contract compatibility, token appreciation, holder distribution or invented financial right is part of this experiment.[12]

## 12. Conclusion and claim ledger

| Claim | Evidence-supported disposition |
|---|---|
| Existing useful work can be delivered locally | Supported for eight complete scoped requests; a ninth has a useful local catalogue and refused native output. |
| A recurring implementation burden exists | Local repeated reads and semantic/history validation are measured; business recurrence is not. |
| A new asset repays complete costs | Not established; no asset selected or constructed after the gate. |
| Alpha constructs improvements better than a conventional builder | Not tested; no qualified two-builder experiment. |
| The service beats equivalent competent human work in total cost | Not established; sourced rates support thresholds, while duration and full Agent cost remain unknown. |
| The system finances and accelerates a second cycle | Not undertaken and not demonstrated. |

The requested positive Flywheel result was not achieved. The measured screen gives a concrete reason: the profiled paths are already automated, their remaining local cost is modest, the larger native outcome is constrained by an existing policy, and the information needed to value and fund an improvement is missing. This bounded screen is not proof that every existing AGI ALPHA workflow lacks a valuable opportunity. In particular, a real recurring Repo-Healer workload remains an unmeasured direction. A primary investment experiment should begin only when its specific missing conditions are supplied. The evidence, useful outputs and refusal remain reproducible.

# References and exact evidence map

## References


[1] AGI ALPHA. *The Advantage That Funds the Next*. Original 14-page PDF and 382-member ZIP (381 listed payload identities), supplied through the user's files. Full hashes and recheck results: `companion/evidence/prior-review/`. Original public source handoff: https://github.com/MontrealAI/AGI-Alpha-Agent-v0/pull/4916; passive paper: https://github.com/MontrealAI/AGI-Alpha-Agent-v0/pull/4917.

[2] AGI ALPHA. Original Economic Engine, Productive Capital, Autonomous Work Economy and Next Advantage manuscripts at frozen main `3b040e68adae8196f7bda4c65f92af92ee6f7335`, `docs/papers/`. Detailed reconciliation: `companion/evidence/prior-review/PRIOR_STUDY_AUDIT.md`.

[3] MontrealAI. Frozen source records: Alpha main `3b040e68adae8196f7bda4c65f92af92ee6f7335`; integration `496fa3c00dc7498173a896404ec756dfdeb94033`; Fleet `45759fb096faf8988f58ae5ea71248d0f35a97c7`; Jobs `bc4abbd5c28f3672275b13300e40ea26c820b8ee`; Manager `5d5670df9dbbc79da115e0a44b71989f34df58c1`. `companion/SOURCE_MANIFEST.json` binds copied runtime bytes.

[4] This discovery record. `companion/evidence/source-profile/approved-process-run/`: `protocol.json`, `results.json`, `process-observations.json`, `summary.json`, exact run artifacts and verifier outputs. Original setup failures remain under `source-profile/`.

[5] This discovery record. `companion/evidence/catalog-profile/`: `summary.json`, `process-accounting.json`, `existing-git-alternative.json`, initial failure, archived CSVs, local outputs, native outputs and history. Two evaluations per four exposed requests.

[6] Alpha main `3b040e68`, `examples/successor-workbench/engine.mjs`, artifact validation; `docs/agent/SUCCESSOR_ALPHA_CONTINUITY.md`, documented stricter native bounds. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/3b040e68adae8196f7bda4c65f92af92ee6f7335/examples/successor-workbench/engine.mjs.

[7] Alpha integration `496fa3c0`, `examples/successor-alpha/automated-costs.mjs` and `shadow-ledger.mjs`. Byte-identical copied modules and final bindings: `accounting/`. Unknown values remain null.

[8] Alpha main `3b040e68`, `alpha_factory_v1/core/runtime/budget.py`; Manager `5d5670df`, `scripts/economics/project-economics.cjs` and `capability-report.cjs`. Native reservations and conditional accounting do not grant financial authority.

[9] U.S. Bureau of Labor Statistics. *Software Developers, Quality Assurance Analysts, and Testers*, May 2025 wage observations. Accessed 25 October 2026. https://www.bls.gov/ooh/computer-and-information-technology/software-developers.htm.

[10] U.S. Bureau of Labor Statistics. *Employer Costs for Employee Compensation*, Table 4, private industry occupational groups, June 2026. Accessed 25 October 2026. https://www.bls.gov/news.release/ecec.t04.htm.

[11] MontrealAI. Issue 4850 and original-action record `6386746278`; this discovery declaration `6405754805`. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/issues/4850. Earlier editions and original holds remain independently owned.

[12] Manager `5d5670df`, `docs/USDC_PAYOUT_SPLIT.md`; Jobs `bc4abbd5`, `docs/successor/TRANSFORMATIONAL_WORK.md`. Source-inspected separate USDC and AGIALPHA routes; no live chain or financial action in this study.

[13] This discovery record. `analysis/runtime-capability-boundaries.json` binds 16 exact source files and their Git identities across the four repositories. Existing Repo-Healer: Alpha `3b040e68`, `alpha_factory_v1/demos/self_healing_repo/repo_healer_v1/engine.py`. Canonical solve-function and computer isolation are distinct from that demo's workspace validation. Source inspection is not execution or admission.


Prepared at Vincent Boucher's direction. Model-assisted research, implementation and review occurred in one collaborative session. Role separation does not establish independent institutional review.