



AUTONOMOUS ENTERPRISES

ACTIVATE. EVOLVE. SCALE.

SUCCESSOR Ω Harness

Powered by TRANSFORMATIONAL INTELLIGENCE ∞ .

A Foundry of Discovery Instruments

Vincent Boucher

MONTREAL.AI · QUEBEC.AI

Montréal, Québec, Canada

22 October 2026 · Reproducible local research edition

BUILD THE MEANS OF BUILDING.

An enterprise whose product is a reusable means of discovery.

Abstract

A cooling developer and a heat-recovery investigator need different answers, yet neither can justify constructing the experimental means alone. We study a foundry whose deliverable is that reusable means: configurable thermal-fluid fixtures (K1), an assembly/calibration station (K2), and a qualified research instrument (K3). A supplied finite grammar yields 108 configurations, of which 2 qualify. A frozen response then supports a separately revealed pulse-calorimetry requirement through a hand-written, checked measurement interface. Three sensor/history realizations pass unchanged criteria; their largest internal heat-integral discrepancy is 0.1711%. A bounded calibration revision helps commission one distinct simulated module. A conventional specialist matches useful quality with fewer calibration readings; the conventional successor arm is equivalent-procedure reuse. This edition corrects initial-sample and command-boundary validation, preserves the original constant-inlet result, and adds an independently implemented high-accuracy reference, end-to-end refinement and fixed model-discrepancy challenges. A hidden flow/power mismatch defeats heat accuracy despite runtime acceptance. Under hypothetical budgets, capability construction and voluntary sharing jointly fund two investigation families; a centralized lab is cheaper and adverse participation or utilization defeats feasibility. The executable contribution combines construction, checked continuation, explicit inheritance and accountable local execution. Evidence comprises executed software and simulated physics, a source-derived contract model and the actual candidate harness's report/review/recovery path. Physical construction, customers, EVM execution and installed thermal execution are **NOT RUN**.

Keywords: autonomous enterprises; programmable instruments; thermal-fluid digital twins; drift-rejecting calibration; capability inheritance; voluntary commissioning; qualified procedure improvement.

Build the means of building

A small cooling team can buy a study. A heat-recovery team can buy another. Both would pay again for fixtures, trustworthy sensing and repeatable starting conditions. The foundry asks: *can they instead commission a shared instrument that makes an expanding class of investigations possible?*




Figure 1: The complete transformation. K1 specifies configurable fixtures; K2 constructs their calibration and conditioning procedure; K3 retains a checked dynamic response. A new pulse-calorimetry investigation reuses those means and builds a measurement method. A distinct simulated successor module reuses the fixture/reset architecture but rebuilds calibration, parameter fit and acceptance evidence. Green bands mark reuse; purple bands mark reconstructed work. No parent identifier transfers acceptance or spending authority. Evidence: [results/foundry_lineage.json](#).

Three contributions. A generated instrument construction with retained alternatives; a corrected continuation interface tested beyond its original constant-inlet mission; and a reproducible comparison that separates model-consistent transfer, numerical accuracy, model discrepancy and voluntary feasibility.

Evidence boundary. The companion executes locally and simulates the thermal apparatus. The native harness executes its existing public-data report worker. The thermal adapter and physical commissioning remain proposed. The downloadable product today is an instrument-design and qualification package.

Contents

1	The enterprise builds the means to ask better questions	4
2	Foundations, predecessors and conventional methods	5
3	The constructed foundry and its digital twin	6
4	Construction, calibration and qualification	8
5	One three-minute journey, with downloadable evidence	11
6	A fresh mission: from cooling response to pulse calorimetry	11
7	Numerical accuracy and the failure an acceptance policy misses	13
8	Qualified improvement of the construction procedure	16
9	Feasibility, voluntary commissioning and venture families	17
10	Ascension accountability and the operational token path	19
11	Extend the real harness, with a narrow demonstrated handoff	22
12	Matched comparisons, resource accounting and negative results	22
13	Limits and a practical deployment path	24
14	Conclusion	24
A	Reproduction, identity and the evidence docket	25
B	Ascension call and failure map	26
C	Mechanisms, inputs and checked consequences	28
D	Revision history and numerical evidence	29
E	Interpretation and attribution	30

1 The enterprise builds the means to ask better questions

A small organization may have a valuable thermal question and no economical way to ask it. A cooling team needs a reliable temperature response for a proposed device. A heat-recovery team needs to know whether the receiving circuit obtains heat at a useful temperature. A research group later needs a time-resolved calorimeter for programmable heat pulses. Each can commission an isolated study; each also needs fixtures, calibrated sensors, controlled boundary conditions, a repeatable starting state and evidence that the answer survives reuse.

The foundry asks a different production question: *can we construct an instrument that makes an expanding class of investigations possible?* Its product is an enabling experimental interface. Replaceable thermal cartridges share known mechanical and fluid ports, two hydraulically isolated water loops, independent control channels, an attachable calibration reference and a finite conditioning procedure. A qualified response model makes later protocols executable. The initial instrument does not predetermine a catalogue of successful businesses: requirements, physical scope, individual budgets and accepted terms screen candidate opportunities.

The central transformation preserves the useful objective. The temperature and accuracy requirements do not become easier when a candidate fails. A request at 65°C is refused, retained in the record and excluded from the success count; it is not silently renamed a low-temperature success. Similarly, a more capable calibration procedure receives no additional spending authority. This intent-preserving discipline follows the distinction between a purpose charter and a changeable formulation in Transformational Intelligence, §§2.1–2.3 [10].

What exists now. Downloadable Python source, numerical experiments, generated construction manifests, raw evidence, separate checking code, an offline five-step experience, a compiled manuscript, and a local demonstration of the candidate harness's maintained report adapter.

What has not been demonstrated. A physical thermal-fluid instrument, installation in a customer laboratory, certified calibration, warranties, paid customers, an external validator council, deployed thermal execution, or any transaction. The concrete proposed launch is a reproducible instrument-design and qualified-software service; real-world commissioning is a subsequent experiment.

1.1 A harness that retains means, not only results

AGI ALPHA organizes agents, tools, memory, bounded jobs and validation. TRANSFORMATIONAL INTELLIGENCE asks which restrictive conditions can change while the beneficiary's objective stays fixed. The six Out mechanisms construct specific observation, representation, resource, policy, composition and successor operations. SUCCESSOR Ω carries the accepted artifacts and their conditions into another mission. AUTONOMOUS ENTERPRISES activate the available means, evolve them through measured improvement and scale only within demonstrated value and capacity.

This changes the object retained by the harness. A result record says that one temperature trace passed. A capability record retains the command interface, fitted dynamics, calibration method, reset conditions, dependency identities, checker, scope and cumulative costs needed to obtain future results. In this companion, a new request is checked against those retained conditions before continuation; missing evidence or unsupported inputs invalidate reuse. The finite constructor and hand-written calorimeter composition remain bounded supplied methods. Their reusable consequence, rather than an open-ended invention claim, is the contribution.

1.2 Research questions and evidence levels

We ask four questions. First, can a constrained constructor produce a reusable instrument interface from alternatives while retaining its failures and total costs? Second, does the frozen interface support a genuinely different downstream measurement requirement? Third, can an improved construction procedure help commission another instrument without inherited acceptance or expanded authority? Fourth, do the capability and a voluntary sharing arrangement jointly create modeled feasibility under unchanged requirements?

Our answer is bounded. The plant, response fitter, predictor and conventional specialist initially call the same forward-simulator implementation. Their close agreement establishes computational consistency and parameter recovery in that family. Section 7 introduces a separately implemented numerical reference and discrepancy challenges; neither is physical laboratory validation. A separate agent authored the withheld mission and checking code inside this project. That separation reduces direct implementation coupling; it is not an independent institution, laboratory or audit firm. All reported numerical evidence is new to this companion. None of the measurements or counts from the earlier coalescer or interval-instrument manuscripts is transferred here.

Label	Meaning in this study
Executed software	Programs actually ran locally, with retained inputs and outputs.
Simulated physics	Results from the documented water-loop twin; no hardware observation.
Modeled contracts	An explicit Python state model derived from pinned Solidity; no EVM equivalence claim.
Proposed deployment	A construction or integration step that still requires implementation and acceptance.

Table 1: The four labels appear throughout the companion and its interface. A receipt does not upgrade the underlying evidence level.

2 Foundations, predecessors and conventional methods

We retrieved the nine requested foundation PDFs from Alpha revision 2b92933a5e3955440be770362acc18f06abf1708 on 22 October 2026. The source manifest records individual retrieval timestamps, SHA-256 values and page counts. All requested page references agree with this edition. Out-Evolve’s configurable fixtures → assembly/calibration station → thermal-fluid research instrument progression on pages 9–11 is proposed there; it is not evidence that the physical station already exists [5, §§4.1–4.5].

Table 12 in Appendix C records the checked source locations. Table 13 makes each mechanism executable: input, constructed output, acceptance and downstream consequence.

These are causal contributions rather than stage names. In particular, Out-Think requires a representation that supports the permitted *future operations*; present-answer agreement is insufficient. Out-Execute’s exact finite boundary-value result requires complete domains, attaining witnesses, factor ownership and input identity. Our continuous port-response artifact is an engineering analogy with a tested scope, not an instance of that finite theorem. Out-Design explicitly warns that pointwise equality does not preserve failure domains or continuous availability. The World as an Output explains why finite noisy observations cannot establish physical support zeros; the simulated hydraulic isolation is an imposed structural premise [4, 6, 9, 11].

The broader methods are conventional. Modelica.Fluid formalizes modular pressure, mass-flow and enthalpy ports with conservation laws [15]. Reversal and error separation are established metrology

[14]; NIST describes a four-reading drift-eliminating comparison [24]. Boundary-preserving reduction of resistive networks is well developed [13]. Reconfigurable experimental platforms, self-driving laboratories and modular science factories already combine instruments, workflow reuse and orchestration [12, 16, 25, 27]. We did not execute those external systems as comparators. This literature supports an integration-benchmark contribution, not priority for calibration, modular instruments, digital twins or autonomous laboratories.

The 27-page interval-instrument predecessor [2] and the separate 15-page launch-kit manuscript [3] remain distinct. The 21-page coalescer edition named in the assignment was not identified in the pinned tree; the current 23-page file bearing a similar general title is a different arithmetic-capability manuscript. No predecessor number is used as evidence for the present experiment.

3 The constructed foundry and its digital twin

3.1 Physical interface and conserved quantities

The twin contains two isolated through-flow thermal circuits, a shared solid cartridge and a slow thermal reservoir representing housing memory. A thermostat supplies each inlet; reservoir regulation is an external boundary service, not a free closed-system energy source. The controlled flows are q_1, q_2 , inlet temperatures T_{i1}, T_{i2} and electrical heater power P . The observable channels are inlet/outlet temperatures, wall temperature, flow and pressure; the latent reservoir temperature is withheld from the constructor's fitting input.

Let T_1, T_2 be outlet lump temperatures, T_w the cartridge wall and T_h the memory state. With $a_j = \rho c_p q_j$ and $K_j = G_j(q_j/q_*)^{0.7}$,

$$C_f \dot{T}_j = a_j(T_{ij} - T_j) + K_j(T_w - T_j) + P_{\text{pump},j}, \quad j = 1, 2, \quad (1)$$

$$C_w \dot{T}_w = P - \sum_j K_j(T_w - T_j) - U_a(T_w - T_a) - K_h(T_w - T_h) + P_{cw}, \quad (2)$$

$$C_h \dot{T}_h = K_h(T_w - T_h) + P_{ch}. \quad (3)$$

Controls use L/min, W, s and °C; q_j converts to $\text{m}^3 \text{s}^{-1}$ in the balances. Temperature differences are K. Commanded flow, pressure/flow readout and heater input are treated as exact; only temperature readout error and electronic calibration drift are stochastic. Water density is 998 kg m^{-3} , heat capacity $4180 \text{ J kg}^{-1} \text{K}^{-1}$, $C_f = 40 \text{ J K}^{-1}$, $C_h = 80 \text{ J K}^{-1}$, $T_a = 20^\circ \text{C}$, $U_a = 0.18 \text{ W K}^{-1}$, and $q_* = 0.1 \text{ L min}^{-1}$. These are declared synthetic plant assumptions, not fitted material measurements.

The energy checker recomputes

$$E(t) - E(0) = E_{\text{heater+pump}} - E_{\text{boundary},1} - E_{\text{boundary},2} - E_{\text{ambient}} + E_{\text{conditioning}}, \quad (4)$$

where $E = \sum C_i(T_i - T_a)$. Boundary heat is signed: a warm supply can add energy. Each loop's incoming and outgoing mass flow is equal; there is no fluid connection between the loops. Their thermal coupling through the wall is retained. The diagnostic energy fields go to the audit stream, not the estimator. The retained checker recalculates the residual from simulator-produced cumulative energy and stored-energy fields rather than trusting its residual column. This is a ledger-consistency check. The independent reference in Section 7 additionally reconstructs fluxes and stored energy from its own state equations and numerical solution.




Figure 2: The foundry's executable boundary. Two controlled circuits share a replaceable thermal core but keep separate fluid and controller state. The Python companion executes physics; the actual candidate SUCCESSOR worker executes a bound evidence report. Remote Fleet/Jobs execution and physical hardware remain outside the demonstrated boundary.

3.2 Hydraulics and the limits of a convenient reduction

The commanded flow requires pressure

$$\Delta p_j = R(T_j)q_j + Bq_j|q_j|, \quad R(T) = 2.6 \times 10^9 e^{-0.022(T-20)} \text{ Pa s m}^{-3}, \quad B = 4 \times 10^{13} \text{ Pa s}^2 \text{ m}^{-6}. \quad (5)$$

Pump dissipation $\Delta p_j q_j$ enters the fluid energy equation. The viscosity model is

$$\mu(T) = 0.001002 e^{-0.022(T-20)} \text{ Pa s}. \quad (6)$$

Reynolds number uses a 6 mm bore. This phenomenological viscosity law and fractional flow exponent are synthetic constitutive assumptions within the stated range, not a universal fluid correlation.

Out-Evolve's fixed-topology setting $\mathbf{q} = \mathbf{q}_0 + V\mathbf{x}$, $\mathbf{p} = R\mathbf{q}$ yields

$$D(\mathbf{x}) = (\mathbf{q}_0 + V\mathbf{x})^\top R(\mathbf{q}_0 + V\mathbf{x}), \quad (7)$$

a quadratic dissipation function when R is fixed [5, p.5]. We inspect this as a *limited near-linear diagnostic*, without a preregistered pressure-error tolerance: $q \leq 0.12 \text{ L min}^{-1}$ and $|T - 20| \leq 0.5 \text{ K}$. The full experiments retain the nonlinear pressure term and thermal state. Changing viscosity, turbulent operation, a new topology, uncontrolled valves or an altered thermal path requires new qualification. We claim no exponential measurement advantage over specialist hydraulics; the source itself admits efficient specialist baselines.

Quantity	Range or fixed value	Interpretation
Controlled flow, each loop	0.04–0.20 L/min	Positive flow only; no fluid reversal assumption.
Inlet setpoint	15–30°C	Input check before simulation.
All fluid/wall temperatures	15–45°C	Runtime trip on violation.
Heater	0–80 W	Bounded schedule, no negative heating.
Pressure / Reynolds limits	20 kPa / 1800	Safety envelope in this single-phase model.
Fitted G_1, G_2	[6,14], [4,11] W/K	Conductance at reference flow.
Fitted C_w, K_h	[140,320] J/K, [0.08,0.55] W/K	Positive bounded identification.
Hidden initial history	$T_h - T_a \in [-10, 10]$ K	Independent repeated-use stress.
Numerical step	1 s; 0.25 s audit	Solver steps are not physical measurements.

Table 2: Declared plant and acceptance envelope. A physical bill of materials, controller or sensor certificate has not verified these values.

3.3 Constructing a repeatable world

The reset is an actuator and a time interval. Both flows are set to 0.15 L min^{-1} , both inlets to 20°C , and the heater to zero for 180 seconds. A modeled wall clamp contributes $P_{cw} = -12(T_w - 20) \text{ W}$ and a housing contact contributes $P_{ch} = -8(T_h - 20) \text{ W}$. These terms extract or supply signed energy; the ledger retains them. No simulator assignment sets the hidden state directly to ambient.

The declared residual target is 0.005 K. The independent repeated-history check on the selected simulated instrument reaches at most 0.001219 K after conditioning for histories $-10, 0, +10 \text{ K}$. A short passive flush cannot establish this reset contract. A separately derived weighted-norm certificate for the declared Euler model bounds the residual by 0.002506 K over the stated parameter box, conditional on the modeled contacts and initial observable nodes at 20 degrees C. Removing reset increases history sensitivity but still meets the coarse initial cooling RMSE limits; it fails the separate repeatability contract, not that cooling criterion. The clamp’s ability to contact the housing, its integrity, the initial history bound and the absence of a new hidden reservoir remain modeling premises. A blocked clamp and a different topology are new worlds requiring rejection or renewed qualification. Finite trials do not prove physical isolation zeros [11, pp.14–15].

The selected response keeps four dynamic states and the full ordered input schedule. A versioned parameter record, units, initial-state contract, raw observations, input identities and dependency hashes accompany it. The underlying sensor observations remain available; retiring a detailed working representation does not mean deleting evidence. A deliberately lossy steady-state summary incurs 3.0993 K wall-temperature RMSE on a pulse continuation. This is the concrete failure that motivates an executable dynamic representation, consistent with Out-Think’s update requirement [9, Proposition3.2].

4 Construction, calibration and qualification

4.1 A generator with alternatives and rejection

The constructor explores a finite grammar: two fluid topologies, three sensor layouts, two conditioning protocols, three calibration programs and three control/reference arrangements, giving $2 \times 3 \times 2 \times 3 \times 3 = 108$ configurations. Constraints require independent fluids, required boundary

and wall observations, the finite reset interface, separate controller failure domains and the reference-measurement threshold. Structural exclusions precede expensive trials. Surviving calibration alternatives face simulated reference experiments; the lowest modeled parts cost among qualified candidates is selected.

This is bounded combinatorial construction. The grammar, allowable components and relevant physics are supplied by the researcher; the program did not invent a new pump or discover an open-ended topology language. Nevertheless, a winning design is not simply returned as a fixture: all candidates, failure reasons, competing prices and simulated calibration observations are retained. Of 108 candidates, 106 are rejected and 2 qualify. Candidate106 uses two isolated loops, port and wall sensing, a clamp, ABBA calibration, independent controllers and a shared reference fixture; its assumed parts cost is 4,530 CAD. Candidate107 also qualifies but uses two reference fixtures and costs 260 CAD more. A single controller is rejected despite its apparent saving because one failure would stop both services.

The reference fixture is scheduled in two four-second reservations, completing by the unchanged ten-second calibration deadline. Independent local controller state and already calibrated operation remain separate. The second fixture dock is released specifically for the conditioning clamp. Reference failure blocks *new* calibration; it does not retroactively prove existing sensors valid. Simultaneous urgent recalibration and common reference corruption are not covered by this availability contract. This is a narrow Out-Design adaptation, not a general certification of fault isolation [4, p.5].

4.2 Sampling, polarity, weights and observed gain

Each measurement first captures the differential signal into a modeled sample-and-hold. Electronic input reversal changes its sign without reversing a fluid process. The frozen signal assumption is necessary: the thermal state can change while the held electrical sample is read. A physical implementation must establish hold droop, settling, sensor linearity and reference traceability. Here hold droop at most 0.001 K is an explicit unverified interface assumption.

For one block,

$$y_i = g\sigma_i\theta + b_0 + b_1t_i + b_2t_i^2 + \epsilon_i. \quad (8)$$

The simulator draws noise with standard deviation 0.002 K. Qualification uses a finite Cartesian set of positive gains 0.96, 1.04, offsets $-0.3, 0.4$ K, slopes ± 0.09 K/s, quadratic coefficients ± 0.002 K/s², and known differences 2, 5, 11 K. These electronic-nuisance stresses are supplied scenarios. Gaussian noise is not bounded, so finite maximum errors are not universal confidence bounds.

A program constructs $\mathbf{t}$, σ and $\mathbf{w}$ subject to

$$\mathbf{w}^\top \sigma = 1, \quad \mathbf{w}^\top \mathbf{t}^k = 0 \quad (k = 0, \dots, d). \quad (9)$$

An observed 8 K reference estimates $\hat{g} = \mathbf{w}^\top \mathbf{y}_{\text{ref}}/8$, and $\hat{\theta} = \mathbf{w}^\top \mathbf{y}/\hat{g}$. References are per channel and block; unrelated sensors inherit no gain. Readings outside the admissible gain interval are rejected. For affine drift, ABBA uses times (0, 1, 2, 3) s, polarities (+, −, −, +) and weights (1, −1, −1, 1)/4. The initial qualifier accepts maximum observed error at most 0.025 K.

This is an adaptation of Out-Learn’s observation-program and unknown-gain mechanisms [7, Theorems 3.1/4.1 and §5], with conventional reversal/drift cancellation explicitly credited [14, 24]. The same policy executes across the nuisance worlds; hidden gain or drift never selects a different program. That uniform policy implements the separate-or-neutralize discipline of Out-Strategise within the supplied nuisance class [8, p.4].




Figure 3: Constructed calibration and its limits. Initial one-reading and simple-reversal alternatives fail the unchanged target. A degree-two procedure later restores margin under stronger curvature; the conventional finite-difference improver ties it. Errors are maxima over retained simulated reference cases, not physical calibration certificates.

4.3 Response construction without latent-state fitting inputs

The station applies a six-segment 1,260-second training schedule, after 180 seconds of conditioning. Heater levels and both flows vary within the original envelope. Every ten seconds it records three calibrated temperatures. The resulting 126 observations per channel require 3,024 signal/reference readouts; the initial alternative-program qualification costs another 672 readouts, including the rejected programs. The estimator receives an explicit observable table containing time and the three temperatures. Latent parameters, internal memory temperature and conservation-audit fields do not enter its objective.

Bounded nonlinear least squares identifies (G_1, G_2, C_w, K_h) in the known four-state grammar. It uses 30 forward-model calls and reaches 0.00155 K training RMSE. The plant, fitter, predictor and specialist all use [foundry/twin.py](#); this millikelvin agreement is internal identification evidence. The retained public-cooling refinement comparison finds maximum differences of 0.02538 K at primary outlet, 0.01979 K at secondary outlet and 0.03495 K at the wall between one-second and quarter-second investigation steps. These exceed the internal agreement but concern different observables and tests. They do not alone violate the coarser acceptance criteria. Moreover, that original refinement leaves reset integration at one second. The new end-to-end reference explicitly refines reset as well (Section 7).

The selected manifest, model/procedure source, training configuration and observations are hashed before disclosure of the downstream requirement. The constructor and evaluator are separated by code and process discipline, not by a secure enclave: the project author can inspect the simulator. The freeze is an in-session chronological commitment, not an externally registered blind trial. The retained first-freeze record provides the original timestamp and identities even when a later reproduction replays the sequence with a now-public mission.

5 One three-minute journey, with downloadable evidence

The offline companion opens on a single request: help construct reusable thermal-fluid research capability for investigations that cannot economically be conducted separately. Five screens follow the transformation:

1. **Need and bottleneck.** The cooling and heat-recovery questions, budgets, quality thresholds and shared experimental bottleneck are visible before results.
2. **Construction.** A physical assembly drawing shows the two loops, replaceable core, sensing, reference and clamp. The selected and rejected candidates remain inspectable.
3. **Qualification.** Experiment traces, prediction errors, the reset evidence and refusal at high temperature distinguish qualification from a visual demonstration.
4. **Successor.** The held-out pulse-calorimetry requirement, frozen-parent response and subsequent calibration improvement show what changes and what remains inherited.
5. **Retained value.** The user can download the capability and evidence snapshot, inspect allowed scope and modeled opportunities, and reproduce the calculations.

The three minutes refer to the guided explanation. The browser replays stored, hashed research evidence, plots it and creates JSON downloads; it does not rerun the Python simulator or invoke the harness. Advanced details expand on demand. A local Python command performs the computation. The four evidence labels remain visible, and the interface links its source records with paths checked against the final archive.

6 A fresh mission: from cooling response to pulse calorimetry

After the first capability freeze, a separately authored requirement is revealed: construct a dynamic receiving-port calorimeter for programmed heat-delivery pulses. It requests the full pulse and recovery windows, receiving-port RMSE at most 0.35 K, cumulative secondary heat error at most 5%, three independent validation realizations and dependency rejection. The protocol is 720 seconds long, with 42 W and 64 W pulses separated by recovery, and changed flow settings. The requirement is different from predicting a cooling device's wall limit or a steady useful return temperature: it asks for time-resolved delivered heat and its integral.

The successor method is a hand-written composition of the frozen response with a segment-aware sampling and integration protocol. Its corrected incremental interface verifies the initial sample as well as every later sample: schedule identity, finite values, unit convention, time, calibration identity, temperature/flow scope, acquisition gaps and command-boundary coverage. The v2 regression suite passes 15 positive and 116 rejection cases, including unchanged-state checks for failed updates and replay recovery. A new manifest retains the parent hash, ten-second acquisition schedule, boundary-temperature/flow dependencies, unit conversion, calibration provenance and acceptance conditions. It is an executable *software calorimeter built from the station*, not a fabricated new physical apparatus. The first frozen model is not refitted. Predictions are written and hashed before the evaluator generates new measurements. Three separate sensor/nuisance realizations vary history and readout noise on the *same* physical parameter instance; they are not three different manufactured instruments.

All three realizations meet the unchanged criteria. The largest receiving-port RMSE is 0.001812 K and the largest prediction-to-measured heat-integral error is 0.1711%. The underlying boundary heat is approximately 5.176 kJ. Ledger consistency, the calorimeter's finite-sampling error and the new independent numerical-reference discrepancies are reported separately. A 55°C inlet request is refused before any programmed operation, and both removed and modified dependencies invalidate reuse.

Nuisance seed	History / K	Port RMSE / K	Heat error / %	Result
7319	-10	0.001203	0.1048	Pass
2791	+0	0.001812	0.1711	Pass
4421	+10	0.001453	0.1203	Pass

Table 3: Held-out pulse calorimetry, with unchanged thresholds 0.35 K and 5%. These are sensor/history realizations of one simulated instrument; the parent response was not refitted.




Figure 4: A new measurement requirement after freeze. The first panel shows the receiving-port response and its frozen prediction; the second shows heat delivered over the complete pulse/recovery schedule. Calibration and quadrature produce the measured estimate. Agreement concerns the supplied simulator; physical transfer is NOT RUN.

Ordinary reuse is allowed to succeed where justified. The full dynamic parent passes this fresh requirement; it need not be made to fail to dramatize a successor. The construction lies in the new scoped measurement method and its checked composition. A static wall-temperature summary, by contrast, cannot inherit transient calorimetry acceptance. This result supports transfer within a known instrument grammar, not an unbounded expansion into unknown physics.

6.1 Correct continuation at the first sample and every command boundary

The first edition checked later samples against the schedule but admitted an in-range, inconsistent initial inlet or flow. It also had no explicit one-sided convention for changing inlet setpoints. These were executable interface defects. They did not alter the retained constant-20°C mission; its predictions, observations, traces and case metrics reproduce unchanged after correction ([evaluation/calorimeter_revision_comparison.json](#)).

The corrected method assigns command metadata at an exact boundary to the interval that just ended. For an integration interval $[t_a, t_b]$ lying within committed segment j , both endpoint heat fluxes use that segment's flow and inlet:

$$\Delta Q_j = \frac{\rho c_p q_j (t_b - t_a)}{2} [(T_{o,a} - T_{i,j}) + (T_{o,b} - T_{i,j})]. \quad (10)$$

Outlet temperature is continuous in this lumped model; commanded inlet and flow may jump. A

sample exactly at a boundary closes the old interval. The next interval uses the new committed values at both ends, without blending inlet setpoints. A missing boundary sample, malformed schedule or invalid initial value is rejected before mutating accepted history. Method identity is versioned so an old checkpoint cannot silently inherit the new semantics.

A two-segment, hand-calculable regression gives 1126.3428 J. The former blended-inlet expression gives 1230.6338 J on the same samples, exposing the defect directly. The full test inputs, expected arithmetic, unit conventions and recovery rules are in [evaluation/CALORIMETER_METHOD.md](#) and [evaluation/calorimeter_checks.json](#).

7 Numerical accuracy and the failure an acceptance policy misses

7.1 A separate implementation, with reset inside the comparison

The decisive extension compares the receiving-port temperature and complete investigation heat with independently transcribed differential equations, hydraulic/thermal fluxes, state-energy reconstruction and discontinuity handling. The reference module imports no production forward simulator. It uses piecewise DOP853, an eighth-order method with seventh-order dense output, through SciPy's documented interface [26]; a separately written classical RK4 implementation checks refinement. This is independence of numerical implementation within the same declared physical model. Water properties, lumped states, contacts and conditioning assumptions remain shared modeling assumptions.

The unchanged K3 hash was recorded at 20:00:52 UTC on 22 October. Numerical tolerances, six discrepancy cases and a separate downstream catalog selection were committed locally at 20:02 UTC; the implementation and an observable-only admission clarification were recorded before the first new integrations at 20:09 UTC. These timestamps document within-project chronology, not external preregistration or laboratory blinding. The parent, useful requirements and all failed cases are retained. See [evaluation/reference_protocol_commitment.json](#), [evaluation/reference_execution_commitment.json](#) and [evaluation/challenge_admission_clarification.json](#).

For each initial memory offset $-10, 0, +10$ K, the reference integrates the 180-second reset followed by the complete 720-second investigation. The production comparison explicitly refines *both* phases at $h = 1, 0.5, 0.25, 0.125$ s, starting each mission from its own computed reset endpoint. Temperature comparisons use the same receiving-port times $10, 20, \dots, 720$ s; an all-integer-second diagnostic is reported separately. Continuous signed secondary boundary heat is integrated over the whole mission. The preregistered local numerical limits are 0.05 K RMSE, 0.10 K maximum error and 1% heat error for the production $h = 1$ arm. They supplement, and do not replace, the original useful limits of 0.35 K and 5%.

The primary reference uses relative tolerance 10^{-10} , temperature-state absolute tolerance 10^{-12} , integral tolerance 10^{-9} and a two-second maximum step. A confirmation uses 10^{-12} , 10^{-13} , 10^{-11} and 0.5 s, respectively. Solvers restart at every command or contact discontinuity. Their maximum receiving-port difference is 7.90×10^{-10} K, below the fixed 10^{-7} K confirmation limit. Independent RK4 at $h = 0.125$ s also meets its fixed 10^{-6} K and 10^{-7} relative-heat limits. These checks support numerical convergence of this model; they do not assign a physical uncertainty to the apparatus.

At $h = 1$ s the worst matched-grid receiving-port RMSE is 0.010130 K and the maximum error is 0.025470 K; the all-integer-second maximum is 0.027816 K. The worst mission heat error is 0.00009657% of reference heat. All three histories meet both the added numerical criteria and the unchanged useful mission requirements. The very small heat discrepancy reflects cancellation over this complete pulse/recovery protocol and is not a bound for arbitrary time windows. The millikelvin agreement in

Step (s)	Euler RMSE (K)	Euler max (K)	Heat error (%)	RK4 max (K)
1	0.010130	0.025470	0.00009657	6.93×10^{-6}
0.5	0.005017	0.012593	0.00005380	3.48×10^{-7}
0.25	0.002497	0.006261	0.00002831	1.95×10^{-8}
0.125	0.001245	0.003122	0.00001451	1.16×10^{-9}

Table 4: End-to-end refinement including reset. Each entry is the worst of three initial memory histories, compared with the tight DOP853 reference. Temperature uses the matched ten-second receiving-port grid; heat error divides by the magnitude of reference mission heat. The RK4 column uses an independent fixed-step implementation and is a numerical cross-check, not a competing physical instrument. Data: [results/reference_validation/report.json](#).

Section 6 is therefore an internal consistency result, not the end-to-end numerical accuracy claim.

7.2 Reconstruct energy rather than trust its ledger

The retained checker verifies the production simulator’s energy ledger. The new checker additionally reconstructs stored energy from state and independently integrates boundary, input, ambient and conditioning fluxes over each constant-command piece of the reference trajectory. It does not consume production energy fields. The production-side reconstruction uses the declared old-state Euler convention, recorded simulated temperatures and an independently advanced memory balance; ledger fields enter only a subsequent diagnostic comparison. A mutation adding 1 J to a production energy field changes the diagnostic by 1 J without changing reconstructed state or flux integral. Raw state/flux records and the reconstruction residuals are retained in [results/reference_validation/report.json](#) and its hashed run directory.

Discretization, ten-second sampling, readout noise, fitted-parameter differences and model discrepancy are evaluated separately. The calibration assumes an exact 8 K reference and ideal polarity/sample hold with 0.002 K raw temperature noise. An eight-case, one-at-a-time sensitivity set perturbs the two conductances by $\pm 5\%$, wall capacity by $\pm 10\%$ and memory conductance by $\pm 20\%$. Its largest receiving-port RMSE is 0.040109 K and largest heat change is 1.643236%. These are declared stress magnitudes, not an estimated posterior, confidence interval or measured uncertainty budget. The components need not be independent and are not added in quadrature. Appendix D gives the separated quantities.

7.3 One admission policy across the declared challenge worlds

The common policy receives only the request, dependency identities, frozen prediction and observed sample stream. It refuses unsupported controls before operation and requires renewed qualification when observed temperature or heat residuals exceed 0.35 K or 5%. It receives no case name, latent state, true flow, power multiplier or reference heat. The independent audit then compares accepted results against reference truth. This separation matters: a failure invisible to the policy must remain an *undetected failure*.

The flow/power case provides that failure. Actual primary flow is 8% high, secondary flow 8% low and heater power 5% high, while command telemetry stays nominal. Observed temperature RMSE is 0.064854 K and the prediction-to-observed-heat residual is 4.947712%, so the runtime accepts. Predicted heat is within 3.318417% of true heat, but the calorimeter’s use of commanded secondary flow produces a 8.696402% error relative to true delivered heat. The measurement method fails its 5% useful heat

Case	RMSE (K)	e_{obs} (%)	e_{true} (%)	Outcome
Nominal control	0.0103	0.0005	0.0007	Pass tested case
8 s sensor lag	0.1909	0.0123	0.0121	Pass tested case
Flow / power mismatch	0.0649	4.9477	8.6964	Undetected failure
Contact –40%	0.1803	12.1465	0.0040	Requalify
45 s sensor lag	0.6473	0.4781	0.4755	Requalify
55°C inlet	—	—	—	Refuse before run
Fresh varying inlet	0.0130	0.5602	0.5475	Pass tested case
Fresh + 8 s lag	0.2724	2.1159	2.0627	Pass tested case

Table 5: Every declared challenge, including the separately selected continuation. $e_{\text{obs}} = |Q_{\text{pred}} - Q_{\text{cal}}| / \max(|Q_{\text{cal}}|, 1\text{J})$ is visible to runtime. $e_{\text{true}} = |Q_{\text{cal}} - Q_{\text{true}}| / \max(|Q_{\text{true}}|, 1\text{J})$ is available only to the independent reference audit. RMSE compares frozen prediction with the sensed receiving port. Admission uses 0.35 K and 5%; “requalify” does not claim recalibration was performed. Data: [results/model_challenges/report.json](#).

requirement without detecting the failure. The required physical improvement is independently qualified flow and power metrology, with uncertainty propagated into heat acceptance; a passing receipt cannot supply it.

A 40% contact-conductance drop at 360 s and a 45-second sensor lag exceed observable residual limits and trigger renewed qualification. An eight-second lag passes this finite test. The 55°C inlet request is refused before simulation. The table retains every case; no refit, tolerance change or privileged selector rescues an unfavorable world.




Figure 5: Numerical refinement and an undetected measurement failure. Left: maximum receiving-port error on the ten-second observation grid, including refined reset, decreases as the production step is halved; curves summarize the worst of three histories. Right: the same acceptance policy compares predicted and observed heat, while the independent audit compares calorimeter and true heat. The flow/power case falls below the runtime’s 5% line but exceeds the true-heat requirement. Contact change is detected through the observable residual. All curves and bars come from retained simulated results; no laboratory measurements are shown.

7.4 A fresh varying-inlet continuation

After recording the unchanged parent, a deterministic selection from a declared downstream catalog chooses staged delivery with receiving inlet temperatures 19, 22 and 20°C, flows 0.06, 0.16 and 0.08 L/min, and segment durations 120, 180 and 180 s. This challenges the corrected command-boundary semantics rather than replaying the original constant-inlet case. The fixed parent and corrected method pass both nominal and eight-second-lag versions: temperature RMSE is 0.013025 and 0.272425 K, and calorimeter-to-true-heat error is 0.547453% and 2.062660%, respectively. There is no refit. This is another finite numerical continuation test, not a third demonstrated venture family or physical instrument. The catalog, selection rule, request, pre-observation prediction commitments and all samples are downloadable.

8 Qualified improvement of the construction procedure

The calibration records expose a limitation: curvature in electronic drift consumes the affine procedure's margin. At the predeclared training curvature, the old procedure still passes, with maximum error 0.01804 K; this is not reported as an initial failure. The Architect's bounded proposal asks for cancellation of quadratic drift within the same maximum block size and unchanged 0.025 K criterion.

The constructor searches 242 feasible timing/polarity programs and minimizes reading count plus a correction-weight variance term. It freezes the chosen program before the stress evaluation. The result uses the same four times, alternating polarities (+, −, +, −) and weights

$$w = \frac{1}{8}(1, -3, 3, -1), \quad (11)$$

which annihilate constant, linear and quadratic drift. The old program fails the predeclared stronger validation curvature at 0.03256 K; the revised one passes at 0.00355 K. A separately implemented conventional third-finite-difference construction produces the same program and ties the result. The improvement is a known mathematical technique selected and installed by a bounded procedure, not general recursive self-improvement.

Three additional sensor-suite realizations retain the improvement's ancestor cost and cumulative calibration operations. More substantially, one distinct simulated module is commissioned from fresh sensor observations using the inherited fixture, reset and model grammar. Its parameters are re-estimated rather than copied. The old calibration fails its fresh standard. The revised procedure qualifies the new module against three calibrated validation realizations. The conventional arm repeats calibration with the same correction program and reuses the same observations, response fit and validation result: it is an *equivalent-procedure reuse comparison*, not a third independent commissioning. The increment is 1,723.73 CAD under declared component/labor assumptions; cumulative construction cost including the original station, improvement labor and all compared calibration/validation arms is 10,861.73 CAD. Costs do not disappear when a parent identifier changes.

The generated lineage manifest (Figure 1) distinguishes K1 fixtures, K2 station and K3 research instrument from the subsequent METHOD and MODULE objects. The child reuses topology, reset and model grammar; it reconstructs calibration, observations, parameters and acceptance. Its active raw history points to its own observations while the parent's history and costs remain explicit ancestors.

The procedure remains limited. The independent reviewer adds a post-freeze cubic-drift stress; the revised procedure fails at 0.06083 K and must be requalified for that class. A gain that looks plausible does not establish an adequate nuisance model. No acceptance rule, budget, permission or valid operating envelope is expanded automatically. Parameter fitting, procedure search, capability construction




Figure 6: A strong comparator and a meaningful failure. Left: the conventional specialist meets the same useful quality with 384 initial qualification readouts, versus 672 for finite-grammar search including rejected programs. Right: quadratic-drift correction succeeds and ties a conventional finite-difference construction; cubic drift breaks the unchanged 0.025 K criterion. These are calibration errors, not held-out receiving-port RMSE. Evidence: [evaluation/specialist_results/report.json](#) and [results/rsi.json](#).

and improvement of the improvement process are distinct operations. This study executes the first three in a supplied grammar. It does not change the general improvement algorithm. The v2 challenge set tests limited model discrepancy, while physical fault tolerance remains a bench question. The old procedure remains available for rollback within its qualified scope; adoption never relaxes acceptance or expands authority.

9 Feasibility, voluntary commissioning and venture families

9.1 Define viability before evaluating it

The prospective buyer receives a designed experimental interface, an executable response, a checked measurement procedure and the evidence needed to reuse them. Participants share fixtures, calibration and instrument time; each pays for its own reserved investigation and receives its data and reusable software. A modeled offering is newly viable when it meets its declared useful-performance requirements within the total cost and deadline, and every participant voluntarily accepts its own assigned share, access rights and evidence terms. The two primary participants have hypothetical caps of 7,000 CAD for cooling and 6,800 CAD for heat recovery; each receives eight reserved, nonexclusive instrument hours, evidence export and no equity or exclusive IP. The shared commissioning deadline is 80 hours. A participant opts in only if quality passes and its share is no greater than both willingness to pay and budget. The campaign commits only if all required participants accept.

These are assumptions, not surveyed demand or real transfers. The cost model assigns 100 CAD/h to human work and 80 CAD/h to lab occupancy. Construction includes the 4,530 CAD generated parts list, 18 assembly hours, 16 engineering hours, eight qualification hours and simulated commissioning time. Its fixed total is 8,844.13 CAD. Each investigation costs 1,040 CAD for four human hours and eight

reserved lab hours. Shared coordination, financing administration and maintenance add 600, 240 and 400 CAD, respectively. Human hours are hypothetical deployment accounting, not measured authoring time. Actual research authoring/review time was not instrumented and is not asserted to be zero. Lab allowances conservatively add one billable second per raw readout to thermal-protocol duration; this is a modeled serial allowance, not measured physical wall time, and real sampling could overlap the thermal run.

The counterfactual existing lab is competent: it can meet both requirements with conventional engineering and refixturing. We do not deny it the ability to redesign. Its assumed shared base and per-mission reconfiguration costs are 4,600 and 4,800 CAD, with 600 CAD coordination. Independent existing construction costs 8,000 CAD per mission. These comparisons concern affordability under explicit assumptions, not physical impossibility or a universal advantage over established laboratories. Legacy-arm quality is assumed achievable by competent refixturing; a separate legacy apparatus was not simulated. They instantiate Invent the Economy's distinction between capability, acceptable terms, actual consent, cash and repeated demand [1, §§2–3].

Capability / arrangement	Total CAD	Share CAD	All opt in?
Existing / separate	16,000.00	8,000.00	No
New / separate	20,128.27	10,064.13	No
Existing / shared	14,800.00	7,400.00	No
New / shared	12,164.13	6,082.07	Yes

Table 6: Four-way economic counterfactual under fixed requirements and individual caps. Values are hypothetical assumptions plus derived operation accounting. Legacy-arm quality is assumed competent engineering, not a separately executed physical apparatus.

9.2 What the four-way result does and does not show

At the same requirements, the proposed new-capability/shared-arrangement combination costs 12,164.13 CAD, or 6,082.07 CAD per participant. Both primary families cross their individual thresholds. Neither change alone suffices in this declared scenario. The finding is conditional complementarity: a new shared experimental means and acceptable terms jointly resolve a budget obstruction. It does not infer revenue from willingness to pay, nor product-market fit from a funded campaign.

A competent centralized lab using the *same new capability* costs 11,704.13 CAD, 460 CAD less, because its modeled administration is lower, holding the construction ledger fixed. The specialist's 288 fewer initial calibration readings would save another 6.40 CAD at the assumed lab rate. Conventional engineering can therefore reach the same transformation; Ascension is an accountability/funding option, not a demonstrated necessity or cost advantage. This loss remains in the results.

If the heat-recovery participant withdraws before commitment, the remaining cooling participant faces 11,124.13 CAD, exceeds its cap and does not commit. At 50% utilization, cost per realized mission becomes 12,164.13 CAD; at 25%, it doubles again. This is sunk-cost allocation per realized use, not an automatic extra payment, contractual default or promised refund. These are financial feasibility failures, not silently guaranteed recurring revenue. The first feasible beneficiary count in the modeled schedule is two; counts two through four meet both the cost cap and deadline, whereas six or more require more capacity or an explicitly changed offering. The proposed campaign has only 1,435.87 CAD of total cost headroom at the tighter participant cap. At least 89.44% effective utilization is required under this allocation. Reducing the assumed legacy refixturing cost by 12.5%, from 4,800 to 4,200 CAD per mission, makes the existing shared alternative viable as well; the joint-necessity conclusion then




Figure 7: Conditional economic feasibility. Left: cost per participant in the four counterfactuals; the dashed line is the more restrictive primary cap. Right: sharing reduces unit cost, but serial capacity violates the original 80-hour deadline beyond four beneficiaries in this schedule. More participants are not automatically feasible scale. All amounts are hypothetical CAD.

disappears. Price/WTP sensitivity and ordinary lab alternatives must be reconsidered for any real proposal.

The retained four-way table prices the original declared construction/commissioning scenario. V2's additional software qualification is not silently free. Its executed compute and operation counts are recorded separately in [results/revision_resource_ledger.json](#); authoring and review labor were not measured. Let $C_{rev} \geq 0$ denote a quoted incremental qualification charge. The two-party shared offer remains within both equal-share caps only for $C_{rev} \leq 1,435.87$ CAD. An independently reconstructed sensitivity retains that boundary rather than inserting an invented zero-cost estimate. The centralized comparison retains its 460 CAD advantage if both offers bear the same new qualification charge.

The withheld calorimeter has a distinct hypothetical 6,500 CAD cap and a twenty-working-day deadline. Its technical acceptance is executed; actual commissioning consent is not. A same-policy extension can price it from the retained fixed/variable ledger, but this third offering is not needed to claim the primary two-family threshold result. The out-of-scope high-temperature family fails. A Cambrian expansion of enterprises is the ambition; evidence-bearing qualification and demand remain the filter.

10 Ascension accountability and the operational token path

10.1 Use the owner-designated suite and its actual boundaries

The owner-designated five-contract inventory is Ethereum chain ID1, dated 18 October 2026. We verify the supplied addresses against the pinned inventory and manifest, then derive calls and guards from deployed-source revision `c0fc5217e364a39a4cd123d17a8317c06b80f820` [17, 19]. This is a source/inventory reconciliation. No new live configuration query or deployed-bytecode equivalence test ran, so the paper does not claim current readiness.

The dated configuration specifies 10-token minimum agent and validator stakes, oracle quorum two

Inventory item	Dated Ethereum chain-1 address
AscensionAccess	0x9eE41dF6d076Fa4F89b102b76DEf1c503a1e7C3A
NovaSeed	0x31BD0736DcD55DA2f9a2BB586C8a98695F86f19E
AscensionRiskOracle	0x198FC6EBee21c6A0B34E4d055AFD6D811AA6dE58
AscensionJobMarket	0xc1d3d7fb064C23c46C4EB20d24ab6EA650FF0ceE
AscensionMark	0x6D5323b2609E36fDc4dB3B2501B7BcF0E131033c
AGIALPHA token	0xA61a3B3a130a9c20768EEBF97E21515A6046a1fA

Table 7: Owner-designated inventory dated 18 October 2026, reconciled with the pinned public manifest. The token has 18 decimals. Historical inventory is not present-state, activation or operational-readiness evidence [19].

and maximum risk 2,000 basis points. It also records activation/public-funding/production readiness as false. The Solidity suite has no global pause function; off-chain readiness flags are not pause controls. A production operator must inspect live state and satisfy current readiness before commissioning. That is a future authorized step, not hidden inside this local study.

10.2 Roles become actions in the same foundry

Insight forms the frozen need/capability hypotheses. Nova-Seeds commits a capsule, FusionPlan and evidence. The Risk Oracle answers the bounded risk-policy question about scope, exposure, participation and funding; it does not replace scientific acceptance. MARK can conditionally finance committed jobs. Sovereign/Business coordinates that bounded work, agents construct and execute, Jobs records assignment/submission/settlement, and the Validator Council reviews the stated acceptance. Architect proposes procedure changes; SUCCESSOR Ω retains source-bound qualified artifacts. These are workflow responsibilities; Sovereign/Business and Architect are not extra deployed contracts in the five-contract inventory.

Role eligibility requires a current canonical ENS ownership check and applicable unexpired governor-issued admission. The source roots are *.alpha.agi.eth, *.alpha.agent.agi.eth and *.alpha.club.agi.eth for business, agent and validator. A resolver address or token balance is not ownership, and tokens alone confer no such role. Admission is checked again at relevant execution/settlement boundaries.

NovaSeed.seal(capsuleHash,planRoot,jobs,metadataURI,parent,evidenceHash) commits an immutable root, nonempty capsule/evidence, 1–128 jobs and bounded metadata. A parent requires current NFT ownership. That lineage is not qualified capability inheritance, equity or intellectual-property title. The companion’s full source/data/calibration dependencies carry qualification separately.

10.3 Funding and work must remain inside the committed plan

An eligible current seed owner opens a risk round; independent staked validators attest a risk number, policy decision and nonzero evidence. Green status depends on quorum, expiry, current seed ownership and eligibility. The review round is seven days. It is neither investment suitability nor a scientific truth oracle.

MARK’s supported campaign route uses open, voluntarily purchased lots through buy, a funded bloom, and route of a specific committed leaf. Funding ends within thirty days; execution duration is greater than three and at most ninety days. Routing requires reserve, current authority, valid proof, an unposted job index and a horizon covering auction, job and three-day review. Exact bounty allowance is set for postFromSeed and then cleared to zero. Uncommitted successor work needs a fresh

valid plan/review/funding boundary. Leftover reserve is not unrestricted reinvestment permission. Funding-phase withdrawal through `sell` is a separate path; expiry or revoked admission does not trap unused capital in the source's withdrawal path.

A directly funded `AscensionJobMarket.post` is the simpler alternative when the business can fund the bounded job without a seed campaign. Agents bid with collateral at least $\max(\lfloor B/10 \rfloor, \text{minimumStake})$ in token base units. The source caps auctions at 32 distinct bids; a job's committee has three distinct, eligible, staked validators with explicitly reserved review availability. An agent cannot also be a committee validator. An assigned eligible worker submits a nonzero result commitment by its due time. The three selected validators may review; the supported payout path needs two approvals bound to the same submitted result, and rechecks their live eligibility. Two rejections fail the job. Timeout distinguishes missed delivery from absent review quorum.

10.4 Exact allocation, burning and refunds

All amounts use token base units, with 18 decimals. For settled price p , each approving validator's gross share is

$$v = \lfloor p/40 \rfloor, \quad w = p - 2v. \quad (12)$$

For each gross allocation $g \in \{w, v, v\}$, burning is $\lfloor g/100 \rfloor$ and the recipient obtains $g - \lfloor g/100 \rfloor$. This is nominally 95/5 allocation and one-percent *payout* burning, with exact integer floors. It is not burning on every transfer. Unused bounty $B - p$ is a separate refund and is not a validator fee or a return on investment.

For the retained synthetic example, $p = 100000000000000000103$ and $B = 120000000000000000123$. Each validator gross is 25000000000000000002 ; worker gross is 950000000000000000099 ; total burning is 10000000000000000000 ; unused bounty refund is 20000000000000000020 . These are arithmetic/model outputs only. No token appreciation, passive return or fiat exchange rate is assumed. The Manager USDC route remains separate from this AGIALPHA path.

10.5 Replayable model, honest scope

The companion implements an explicitly identified Python contract model, not a local EVM. Fourteen scenarios execute: ordinary success, rejected evidence, expired admission/review, insufficient funding/reserve, lost ownership, unavailable review capacity and other boundary failures. A separate checker tests 162 price edge cases, exact floors, conservation and refund separation. A separately written three-leaf Merkle construction reconciles with the actual pinned native FusionPlan compiler, including the odd-leaf rule. Synthetic receipts use an explicit model namespace and never resemble mainnet transaction evidence.

These checks do not establish gas behavior, external token/ENS behavior, reentrancy safety, race equivalence or deployed-bytecode correctness. Source-derived canonical signatures/selectors and the full caller/event/guard matrix are retained. Even a genuine settlement receipt would not prove physical correctness: scientific acceptance must bind to the underlying instrument evidence and scope. Appendix B gives the operational call map.

11 Extend the real harness, with a narrow demonstrated handoff

Read-only verification on 22 October 2026 at 16:01 EDT found PR4863 open/unmerged at c455ad4ce109c86df2fef09d19674810290b5185 and bridge PR4865 open/unmerged at 6bb361d522d80dc900cc2a480014531c6a44b66e. The newly observed main revision was 61532fb82621a065e668fd11e57289dcfb65c67b; the requested enterprise module and harness guide were absent there as well as at the original paper-source pin. The retained native execution uses the unchanged candidate sources, with original and refreshed provenance kept separately. It is not evidence of released-main thermal functionality [18, 21]. The read-only responses are retained in [sources/status_refresh/2026-10-22.json](#).

The unchanged candidate provides fixed reliability, performance, documentation, public-data and numerical work families, immutable configuration identities, an existing signed journal, bounded attempts, exact-result review, explicit resume and terminal stop. Our thin bridge sends the foundry's integer operation ledger through its maintained public-data Node worker. A separate Python checker recomputes every group result before accepting the exact digest. The report's grand total mixes categories and has no physical-resource meaning; only the category-specific counts are scientifically interpreted. The bridge then reloads retained records, exports public evidence, derives an *unaccepted* child, injects interruption and resumes, tests stop, compiles a native FusionPlan, validates the existing typed Studio handoff and rejects caller-identity substitution/provisioning requests. Thirteen checks pass on the actual foundry input, not a substitute fixture.

This proves that the candidate's report execution/review/recovery and typed handoff work locally on bound foundry evidence. It does not prove that the maintained executor ran arbitrary thermal code. Physics runs in this research companion. A reviewed fixed thermal-fluid adapter, bounded schema, implementation, independent checker and corresponding Studio allowlist entry are narrow proposed additions. The precise input/output, units, dependency requirements, allowed operations, errors, output limits and recovery contract are in [integration/thermal_adapter_proposal.json](#). They are **NOT RUN** as installed functionality. The bridge introduces no token, chain deployment, parallel database or external service. Free local planning requires neither campaign funding nor a wallet, although compute and human time remain resources.

Repository	Actual responsibility	Demonstrated boundary
AGI-Alpha-Agent-v0	Discovery, retained work/review/recovery; Ascension commitments	Candidate report worker and typed handoff ran.
AGI-Fleet-Linux	Qualified installed execution and authorized independent review	Source inspected; installation/remote execution NOT RUN .
AGIJobsv0	Its own JobRegistry v2 / AGIALPHA work lifecycle	Separate protocol; no remote receipt translation.
AGIJobManager	Separate USDC escrow and lifecycle	No USDC call or transfer.

Table 8: Four repositories, with no invented round trip. Companion source pins are retained in the evidence manifest [20, 22, 23]. Alpha's AscensionJobMarket is not Jobs v0's JobRegistry.

12 Matched comparisons, resource accounting and negative results

The strong conventional comparator uses the same instrument grammar, sensing, redesign opportunities, training observations and acceptance. It constructs standard ABBA analytically and fits the

response with a separately initialized least-squares procedure. It meets the same held-out criteria. It uses 384 initial qualification readouts versus the proposed search's 672, avoiding failed calibration candidates; its optimizer makes 35 calls versus the proposed 30. Training readouts and parts are the same. Thus useful quality ties, and the proposed constructor loses on initial qualification measurements. No unavailable external system is claimed inferior.

Independent mission construction is represented with the same useful requirements and full permission construction cost. Ordinary reuse of the *qualified full* instrument succeeds on the held-out mission. Reuse of an underspecified static summary fails the pulse continuation. The constructed successor method adds a specific measurement operation and retained acceptance; the revised procedure helps commission a fresh module, while the conventional improver ties. These four comparisons distinguish capability construction from weak baselines and from mere copying.

Resource	Count	Meaning
Candidate configurations	108	106 rejected; two satisfy constraints
Initial calibration readouts	672	Includes rejected alternatives; conventional specialist 384
Initial training readouts	3,024	Three channels, 126 time points, signal + reference
Initial simulated protocol duration	1,440	Seconds: 180 conditioning + 1,260 investigation
Initial forward-model evaluations	30	Numerical fitting calls; conventional specialist 35
Held-out readouts	1,752	219 secondary-temperature acquisitions across three traces
Held-out protocol duration	2,700	Seconds: 540 conditioning + 2,160 investigation
Held-out control segments	18	15 investigation segments + three conditioning commands
Procedure-search alternatives	242	Mathematical candidate programs, not sensor operations
Distinct-module training readouts	3,024	Fresh observations; original ancestor costs retained
Actual physical hardware operations	0	Physical construction and deployment NOT RUN

Table 9: A disaggregated resource ledger. Seconds, control segments, sensor readouts and model evaluations are not added into one performance score. Full per-stage raw evidence and timing are included.

Counts are kept in their natural units. A control-segment command, a sensor readout, a numerical time step, a forward-model call, an assumed human hour, a CAD allowance and a token base unit are different quantities. The suite uses no foundation-model calls during reproduction. Authoring and reviewing this research did use human/assistant effort that was not measured by a stopwatch. Local timings are retained with environment identity and may vary across runs; no timing or local receipt is advertised as an invariant scientific result. Protocol counts refer to simulated commands; actual physical operations are zero.

The retained negative results matter: 106 initial candidate rejections; failed simple calibration methods; invalid steady-only dynamic reuse; loss of original reset guarantee under passive/blocked conditioning; refusal outside the thermal envelope; invalidation on dependency loss/mutation; the frozen calibration's curvature failure; the revised calibration's cubic-drift failure; the conventional specialist's reading-count advantage; the centralized lab's lower modeled cost; and the withdrawal/underutilization failures. No subsequent success erases them.

13 Limits and a practical deployment path

The digital twin is deliberately small. It assumes a single-phase fluid, known flow exponent, lumped thermal states, externally regulated inlets, fixed heat capacities, ideal control-setpoint tracking and a specified calibration/noise family. It excludes leakage, fouling, cavitation, boiling, distributed transport delay, changing contact geometry, environmental disturbances beyond those represented and common-cause reference corruption. Conservation consistency is necessary but cannot validate those constitutive choices. The production plant, fitter, predictor and specialist share one forward implementation. The independently implemented high-order reference reduces that implementation coupling for numerical checks, while the declared discrepancy cases expose selected departures. All still depend on synthetic constitutive assumptions; physical validation is a distinct task.

The search does not autonomously invent its grammar or perform unrestricted scientific discovery. The held-out requirement is unknown to the constructor until freeze, but the evaluator knows the available primitives and model class. The three realizations share one plant parameter instance; the later module-commissioning experiment changes the plant separately. No large-population inference follows from either. Scope hashes and local signatures bind records, not truth or legal ownership. The separately written checks share this project and its assumptions. The fresh challenge was selected and committed after the parent capability freeze and before its new outcomes. Local timestamps and agent separation document chronology; they do not prove external blinding.

Economic viability is conditional on hypothetical prices, labor, useful heat thresholds, budgets, reservations and utilization. No customer interview, invoice, supplier quotation, tax analysis, certification or purchase commitment occurred. Conventional shared laboratories may already be the best implementation. Any physical offer needs actual pricing, validated calibration, maintenance capacity, terms allocating responsibility and evidence of willingness to commission. Funding and token transfers cannot supply missing demand or scientific evidence.

The next useful deployment unit is small: deliver one reproducible design/qualification study for a specified beneficiary question, with its command envelope, candidate bill, uncertainty assumptions, response and checker. Free local work needs no fundraising. A bench pilot then measures one two-loop module: calibrated inlet/outlet and wall temperatures, independently measured flow and electrical power, pressure, reset carryover, isolation and repeated pulse response. Obtain component quotations and participant acceptance before any physical commissioning commitment. Compare measured fluxes and stored energy with a predeclared uncertainty budget; renew the scope when a criterion fails.

Only after that pilot should the fixed thermal worker be reviewed and installed in the existing harness, followed by a bounded commissioning arrangement and qualified remote execution. A second bench or venture family repeats the relevant qualification; it does not inherit unlimited capacity by copying software. The declared serial schedule already fails beyond four beneficiaries. Broader adoption must price model diversity, calibration renewal, shared-fixture contention, failed launches and marginal demand. Certification, warranties, customer contracts and live-chain readiness are future evidence-bearing milestones.

14 Conclusion

The demonstrated transformation is a reusable means of investigation: K1 configurable fixtures, K2 assembly/calibration station and K3 qualified research instrument; a separately constructed calorimeter software method; and a bounded procedure revision that helps commission one distinct simulated module. Corrected continuation checks, independent numerical reference and explicit discrepancy

challenges make the retained operating conditions inspectable. The evidence also establishes limits. Conventional methods match the scientific result, a specialist saves calibration readings, the centralized shared-lab alternative is cheaper, and unfavorable participation or utilization defeats the modeled market. Those findings make the construction inspectable and constrain what should be built next.

An enterprise builds an instrument.
The instrument makes new discoveries possible.
Those discoveries help construct its successor.

ACTIVATE available capability.
EVOLVE the means of discovery and production.
SCALE what evidence shows can create value.

FIND YOUR ALPHA. BUILD YOUR SUCCESSOR Ω .

A Reproduction, identity and the evidence docket

From the extracted companion root, the single entry point is:

```
python reproduce.py
```

It checks the pinned runtime, executes construction, public qualification, ablations, economics, procedure improvement, the separately authored held-out/checking modules, fresh-module commissioning, the contract model, the actual candidate harness report path, figure/table generators and the offline experience. The paper can then be compiled with the retained LaTeX/BibTeX sources; `python reproduce.py --paper` includes that step when the TeX toolchain is installed. The bootstrap command documented in the README creates a local environment from exact dependency pins. No account, wallet, external service or foundation model is needed for the scientific reproduction.

The retained execution uses Python 3.12.14, NumPy 2.3.5, SciPy 1.17.0 and Matplotlib 3.10.8; the complete Python import closure and Node version are in the environment docket. Seeds, configuration, sources and raw logs are included. Initial and held-out commitments are retained separately from rerun outputs. Numerical artifacts reproduce within declared tolerances with the pinned environment; wall times, local signing keys/signatures and timestamp-bearing receipts intentionally vary. A replay cannot recreate the original fact that a now-public mission was unseen; the chronological original freeze/reveal record documents that limited claim.

Table 10: A compact evidence route. HTTPS links retrieve exact archived files; labels show package paths.

Claim or operation	Primary record
Generated alternatives	results/candidate_ledger.json
Frozen K3 and lineage	results/capability_v1.json ; results/foundry_lineage.json
Original transfer	evaluation/heldout_results/report.json
Corrected continuation	evaluation/calorimeter_checks.json ; evaluation/calorimeter_revision_comparison.json
Reference and convergence	results/reference_validation/report.json
Declared discrepancy failures	results/model_challenges/report.json
Successor and its costs	results/successor_build.json

Claim or operation	Primary record
Economic model and new charge	results/economy.json ; results/revision_resource_ledger.json
Python contract model	results/contracts/report.json
Actual candidate harness	results/harness/report.json
Proposed thermal adapter	integration/thermal_adapter_proposal.json
Reproduction and archive replay	results/reproduction_report.json ; results/archive_replay/report.json
Visual/browser checks	evidence/visual_qa.json ; interface/qa/report.json
Revision and focused review	REVISION_NOTE.txt ; notes/v2_numerical_review.md

B Ascension call and failure map

The complete declarations, source lines, guard strings and event names are in [contracts/source_call_matrix.json](#); 41 canonical ABI signatures and selectors are in [contracts/source_abi_signatures.json](#). They are source-derived, not a freshly compiled or live-verified ABI. Here Spec is (string goal, string successMetric, uint96 bounty, uint32 duration, uint16 priceWeight) and Role is a uint8 enum, Business0, Agent1, Validator2. Governor-only, token allowance and current eligibility requirements must not be replaced by UI role labels.

Table 11: Operational call map derived from pinned Ascension source. Spec expands to the exact tuple defined above. Full canonical signatures and selectors accompany the code; this table does not authorize calls.

Call	Caller and acceptance boundary	Events, deadline or refusal
Access <code>register(uint8,string)</code> <code>setAdmission(address,bytes32,uint64)</code>	Register: caller owns the canonical ENS child. Admission: governor only; nonzero evidence for an active admission, at most 365 days.	Registered; AdmissionSet. Reject wrong ownership, noncanonical label or invalid expiry. Token holdings confer no role.
NovaSeed <code>seal(bytes32,bytes32,uint32,string,uint256,bytes32)</code>	Eligible business; nonzero capsule, root and evidence; 1–128 jobs; metadata 1–512 bytes. Optional parent NFT owned by caller.	SeedSealed. Reject empty commitment, job bound, metadata or parent ownership. Root immutable; lineage conveys no scientific acceptance.
Oracle <code>deposit(uint256)</code> <code>setJobAvailability(address,uint64,uint32)</code>	Positive validator deposit requires live eligibility. Explicit business availability up to 120 days; capacity at most 128.	StakeChanged; JobAvailability. Job locks require staked eligibility, sufficient reserved horizon and free capacity.
Oracle <code>openReview(uint256)</code> <code>attest(uint256,uint16,bool,bytes32)</code>	Current eligible seed owner opens after prior review expiry. Staked validators differ from owner and creator; unique vote, nonzero evidence.	ReviewOpened; RiskAttested. Seven-day round; risk 0–10000. Reject expired/transferred/decided round or duplicate vote.
Oracle <code>green(uint256)</code> <code>withdraw(uint256)</code>	Green requires quorum, unexpired round, same eligible owner and continuing eligibility of all supporters. Withdrawal requires time strictly after stake lock.	Green is a view, not a science verdict. StakeChanged on withdrawal; insufficient or locked stake rejects.
MARK <code>open(uint256,uint96,uint96,uint32,uint96,uint64,uint32)</code>	Current green seed owner. Positive base; at most one million lots; reachable positive target. Funding end in next 30 days; duration greater than 3, at most 90 days.	MarketOpened. Reject non-green owner, existing campaign, bounds or unreachable target. Linear-curve terms are committed before buying.

Continued on next page

Ascension operational call map (continued)

Call	Caller and acceptance boundary	Events, deadline or refusal
MARK buy(uint256,uint32,uint256) sell(uint256,uint32,uint256)	Buy: admitted contributor, green current owner, Funding phase before deadline, lot and max-cost bounds. Sell: owned lots in Funding, minimum return.	Traded. Buy rejects slippage or insufficient funding balance. Sell does not require current admission or green review; unused-funding withdrawal survives their expiry.
MARK transferLots(uint256,address,uint32) bloom(uint256)	Lot transfer: admitted sender/recipient, open funding and green review. Bloom: same business/seed owner, live business ENS, target met before funding end.	LotsTransferred; Bloomed. Reject authority, insufficient lots or reserve. Bloom fixes campaign executionEnd.
MARK route(uint256,uint32,Spec,bytes32[],uint32,address[3])	Bloomed business with live node, current seed owner and green review. Unposted committed index, valid Merkle proof, enough reserve.	MissionRouted. Require now + auction + job duration + 3 days no later than executionEnd. Exact bounty allowance is cleared after postFromSeed.
Jobs post(Spec,uint32,address[3]) postFromSeed(address,Spec,uint32,address[3])	Direct caller or routed business must be eligible. postFromSeed is MARK-only. Client supplies exact bounty. Committee: three distinct eligible staked validators, independent of client/business.	JobPosted. Reject missing reserved capacity. Bounty at least 100 base units; job duration 1 second–90 days, auction 1 second–7 days, weighting 0–10000.
Jobs deposit(uint256) withdraw(uint256) bid(uint256,uint96,uint32)	Eligible agent deposits; withdrawal cannot consume locked stake. Bidder differs from business and committee; free stake covers max(floor(B/10), minimumStake).	StakeChanged; BidPlaced. Before auctionEnd; at most 32 distinct bids. Price 100 base units–bounty, duration positive and within specification.
Jobs award(uint256)	Callable after auctionEnd. Only currently eligible bids still fitting executionEnd compete. Lowest source score; ties by price, duration, numeric address.	Assigned or Closed(Refunded). Losing locks released. Due is award time plus winning duration; no feasible bid refunds bounty.
Jobs submit(uint256,bytes32,string)	Assigned worker, still eligible, at or before due. Nonzero result commitment; URI 1–512 bytes.	Delivered. Reject wrong worker/state, lateness, expired eligibility or missing result. Result commitment alone proves no physical property.
Jobs validate(uint256,bytes32,bool,bytes32)	Selected currently eligible validator; unique vote; exact submitted hash and nonzero evidence; at or before reviewEnd.	Validated; Payout; Closed. Two approvals pay; two rejections fail and slash. Both deciding quorums are rechecked for current eligibility.
Jobs timeout(uint256) claimRefund(uint256)	Timeout after due for Assigned, or after reviewEnd for Submitted. Refund claimed once by client in terminal state; MARK is client for routed jobs.	Closed; RefundClaimed. Missed delivery slashes; absent reviewer quorum refunds without slashing. Unused bounty remains separate from payouts.
MARK collect(uint256) close(uint256) reclaim(uint256)	Collect terminal routed job once. Close after funding deadline, or with zero active jobs and all committed jobs posted / execution expired. Closed-campaign lot holders reclaim unused reserve.	MissionCollected; CampaignClosed; CapitalReturned. Reject duplicate collection, live missions or no capital claim. Reclaim is pro rata unused capital, not promised yield.
Oracle bindMarketplace(address) Jobs: bindMark(address)	Governor-only one-time graph bindings; deployed target and reciprocal dependency identity checks.	MarketplaceBound; MarkBound. Reject rebinding or wrong dependency graph. No global Solidity pause exists in this five-contract suite.

C Mechanisms, inputs and checked consequences

Source and checked location	Distinct construction here	Consequence of removing it
Transformational Intelligence, §§2.1–2.3, pp.3–4	Fixed mission charter, thresholds, scope and authority	Changed goals can masquerade as a discovery.
Invent the Economy, §§2–3, pp.3–5	Capability and individual commissioning choices evaluated together	Aggregate surplus is incorrectly counted as consent or revenue.
Out-Learn, pp.1,14; Thms. 3.1, 4.1, pp.5–6; §5	Times, polarity, correction weights and observed gain reference	Drift and unknown gain remain aliased with the quantity measured.
Out-Think, p.2; Def. 2.1; Prop. 3.2	Executable dynamic port response with retained updates and history identity	A steady summary gives the wrong answer to a pulse continuation.
Out-Design, p.5; Def. 3.1; Prop. 3.2	Shared reference fixture, separate controller state, checked reservation window	A released fixture at the wrong location is useless; consolidation can destroy availability.
Out-Strategise, p.4, §§2.2–2.3	One calibration/reset policy across nuisance worlds	Hidden ground truth chooses a different policy in each world.
Out-Execute, p.4; Def. 3.1; Thm. 3.2	Boundary response, full input identity, evidence and dependency checks	A missing sensor, source or witness invalidates composition.
Out-Evolve, pp.9–11; §§4.1–4.5	Qualified station helps construct a calorimetric method and commission a new module	Parent identifiers or renamed products are mistaken for inheritance.
The World as an Output, pp.1,14–15; Theorems 3.1 and 4.1	Finite clamp/flush and an explicit nuisance-state contract	Prior thermal history contaminates subsequent measurements.

Table 12: Source mechanisms and new thermal adaptations. The cited mathematical results retain their own assumptions; this table does not transfer their guarantees to physical hardware. References: [1, 4–11].

Table 13: Causal operations in the companion. These are new thermal adaptations of the source mechanisms, with explicit acceptance and downstream consequences.

Mechanism	Input and constructed output	Acceptance test	Downstream consequence
Out-Learn	Unknown gain/drift; choose sample times, signs and weights with a reference.	Calibration max error ≤ 0.025 K; retain rejected programs and readings.	Calibrated observations can identify and validate a response. Cubic drift later fails.
Out-Think	Ordered commands and histories; construct dynamic port response and checked calorimeter continuation.	Pulse prediction; initial/boundary checks; rejected updates preserve state; recovery replays history.	A new heat-integral query is executable; a steady summary loses necessary dynamics.

Mechanism	Input and constructed output	Acceptance test	Downstream consequence
Out-Design	Two independent services and one reference fixture; reserve shared calibration while retaining two controllers.	Isolation and reservation-window checks; retain consolidation failure.	A fixture is available at the required place/time without merging control state.
Out-Strategise	Remaining nuisance worlds; use one reset/calibration and observation-based admission policy.	Same protocol for every world; fixed useful thresholds.	Residual failures require requalification; hidden flow error remains an undetected failure.
Out-Execute	Boundary response, units, input identity, calibration and dependencies; construct a composable manifest.	Remove or alter a dependency and reuse is refused; keep attaining sample history.	Successors can check whether a retained capability applies to their exact request.
Out-Evolve	Qualified station and a revealed limitation; revise calibration and commission a distinct module.	Fresh calibration, refit and three validation traces; retain inherited costs.	One new module qualifies; conventional equivalent-procedure reuse ties. No extra authority follows.

D Revision history and numerical evidence

The supplied review reported reconciliation of 465 archive-manifest entries and 21 HTML receipts, manuscript/interface inspection and independent recomputation of central calculations. It explicitly did not run the complete pipeline. We preserve that supplied report as review context; the original edition's own 16-stage execution and extracted-archive replay remain separate evidence. The previous PDF, source, result summaries and freeze records are retained under [history/v1/paper/main.pdf](#).

This edition corrects the initial calorimeter check and integration semantics; versioned recovery rejects incompatible checkpoints. It fixes the contract receipt path, standardizes K1/K2/K3 lineage, distinguishes equivalent-procedure reuse, adds the independent reference and challenges, and regenerates affected manifests, hashes, figures and interface. Original constant-inlet transfer outcomes remain unchanged. The short revision note records changed results and the exact new execution scope. It never retroactively relabels the first edition as having passed the new checks.

Table 14: Separated error sources. Quantities are maxima over the stated finite tests and have different conditioning; they must not be combined as independent uncertainties.

Component	Retained magnitude	Meaning and limit
Production discretization	0.010130 K RMSE; 0.00009657% heat.	Same exact parameters and history; reset and investigation refined together.
Reference confirmation	7.90×10^{-10} K maximum.	Tighter tolerance and maximum step; same declared equations.
Fitted parameters	0.000221 K RMSE; 0.002303% heat.	Fitted versus nominal parameters under independent reference; not a confidence interval.

Component	Retained magnitude	Meaning and limit
Post-reset approximation	0.000057 K RMSE; 0.002123% heat.	All-20 C prediction start versus integrated reset, fitted parameters held fixed.
Ten-second sampling	0.000732% heat.	Corrected noiseless calorimeter versus continuous reference integral.
Readout nuisance/noise	0.001795 K RMSE; 0.170754% total heat.	Three seeded sensor realizations; includes sampling, not an independent error term.
Parameter stress	0.040109 K RMSE; 1.643236% heat.	Eight one-at-a-time perturbations; not an empirical uncertainty distribution.
Calibration and reference	Exact 8 K reference; 0.002 K raw noise; ideal polarity/sample hold.	Not measured, traceable or physically certified.
Model discrepancy	8.696402% true-heat error accepted by runtime.	Flow/power mismatch exposes missing metrology; not bounded by internal residual.

E Interpretation and attribution

The ordinary numerical tools and conventional physical methods are reusable commercial building blocks subject to their own licenses; this paper grants no license to someone else's patents, data or hardware. The companion's dependency notices retain the upstream software attribution. The thermal-fluid benchmark, additional checker and manuscript are new research artifacts in this edition. Original repository files are retained at their pinned revisions; their local execution is not a merge, release or deployment.

A genuinely useful instrument-design service can begin with reproducible software and a precise construction/qualification plan. Whether it becomes a physical foundry, a successful commercial offer or an expanding family of enterprises is a question for the subsequent evidence, not a result supplied by the name SUCCESSOR Ω .

References

- [1] Vincent Boucher. Invent the Economy: Transformational Intelligence for Creating Viable Markets. MONTREAL.AI research manuscript, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/2b92933a5e3955440be770362acc18f06abf1708/docs/Invent_the_Economy.pdf. Edition 2026-10; repository revision 2b92933a5e39; retrieved 22 October 2026.
- [2] Vincent Boucher. Autonomous Enterprises: Activate. Evolve. Scale.—SUCCESSOR Omega Harness: Constructing and Financing Reusable Productive Capability. MONTREAL.AI research manuscript, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/2b92933a5e3955440be770362acc18f06abf1708/docs/Autonomous_Enterprises_SUCCESSOR_Omega.pdf. Edition 2026-10-22; repository revision 2b92933a5e39; retrieved 22 October 2026.
- [3] Vincent Boucher. Autonomous Enterprises: Activate. Evolve. Scale.—Reusable Launch Kits with Rejection-Only Improvement Transfer. MONTREAL.AI research manuscript, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/2b92933a5e3955440be770362acc18f06abf1708/docs/Autonomous_Enterprises_Activate_Evolve_Scale.pdf. Edition 2026-10-22; repository revision 2b92933a5e39; retrieved 22 October 2026.

- [4] Vincent Boucher. Out-Design: Creating the Freedom to Transform. MONTREAL.AI research manuscript, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/2b92933a5e3955440be770362acc18f06abf1708/docs/Out_Design.pdf. Edition 2026-10-21; repository revision 2b92933a5e39; retrieved 22 October 2026.
- [5] Vincent Boucher. Out-Evolve: Building the Means of Discovery and Production. MONTREAL.AI research manuscript, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/2b92933a5e3955440be770362acc18f06abf1708/docs/Out_Evolve.pdf. Edition 2026-10-20; repository revision 2b92933a5e39; retrieved 22 October 2026.
- [6] Vincent Boucher. Out-Execute: Constructing the Conditions of Completion. MONTREAL.AI research manuscript, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/2b92933a5e3955440be770362acc18f06abf1708/docs/Out_Execute.pdf. Edition 2026-10-21; repository revision 2b92933a5e39; retrieved 22 October 2026.
- [7] Vincent Boucher. Out-Learn: Change What You Need to Know—Constructed Experiments and Sharp Calibration Guarantees for Reusable Learning. MONTREAL.AI research manuscript, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/2b92933a5e3955440be770362acc18f06abf1708/docs/Out_Learn.pdf. Edition 2026-10-21; repository revision 2b92933a5e39; retrieved 22 October 2026.
- [8] Vincent Boucher. Out-Strategise: Transformational Intelligence for Constructing the Conditions of Success. MONTREAL.AI research manuscript, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/2b92933a5e3955440be770362acc18f06abf1708/docs/Out_Strategise.pdf. Edition 2026-10-20; repository revision 2b92933a5e39; retrieved 22 October 2026.
- [9] Vincent Boucher. Out-Think: Transformational Intelligence for Inventing the Space of Reasoning. MONTREAL.AI research manuscript, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/2b92933a5e3955440be770362acc18f06abf1708/docs/Out_Think.pdf. Edition 2026-10-21; repository revision 2b92933a5e39; retrieved 22 October 2026.
- [10] Vincent Boucher. Transformational Intelligence: Learning to Change the Game. MONTREAL.AI research manuscript, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/2b92933a5e3955440be770362acc18f06abf1708/docs/Transformational_Intelligence.pdf. Edition 2026-10; repository revision 2b92933a5e39; retrieved 22 October 2026.
- [11] Vincent Boucher. The World as an Output: Constructive Sufficiency for Reusable Discovery. MONTREAL.AI research manuscript, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/2b92933a5e3955440be770362acc18f06abf1708/docs/The_World_as_an_Output.pdf. Edition 2026-10-20; repository revision 2b92933a5e39; retrieved 22 October 2026.
- [12] Anne-Catherine Bédard, Andrea Adamo, Kosi C. Aroh, M. Grace Russell, Aaron A. Bedermann, Jeremy Torosian, Brian Yue, Klavs F. Jensen, and Timothy F. Jamison. Reconfigurable system for automated optimization of diverse chemical reactions. *Science*, 361(6408):1220–1225, 2018. doi: 10.1126/science.aat0650. URL <https://doi.org/10.1126/science.aat0650>.
- [13] Florian Dorfler and Francesco Bullo. Kron Reduction of Graphs With Applications to Electrical Networks. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 60(1):150–163, 2013. doi: 10.1109/TCSI.2012.2215780. URL <https://doi.org/10.1109/TCSI.2012.2215780>.
- [14] Chris J. Evans, Robert J. Hocken, and W. Tyler Estler. Self-Calibration: Reversal, Redundancy, Error Separation, and ‘Absolute Testing’. *CIRP Annals*, 45(2):617–634, 1996. doi: 10.1016/S0007-8506(07)60515-0. URL [https://doi.org/10.1016/S0007-8506\(07\)60515-0](https://doi.org/10.1016/S0007-8506(07)60515-0).

- [15] Rüdiger Franke, Francesco Casella, Michael Sielemann, Katrin Proelss, Martin Otter, and Michael Wetter. Standardization of Thermo-Fluid Modeling in Modelica.Fluid. In *Proceedings of the 7th International Modelica Conference*, pages 122–131, 2009. doi: 10.3384/ecp09430077. URL <https://simulationresearch.lbl.gov/wetter/download/2009-modelicaFluid.pdf>.
- [16] B. P. MacLeod, F. G. L. Parlane, T. D. Morrissey, F. Häse, L. M. Roch, K. E. Dettelbach, R. Moreira, L. P. E. Yunker, M. B. Rooney, J. R. Deeth, V. Lai, G. J. Ng, H. Situ, R. H. Zhang, M. S. Elliott, T. H. Haley, D. J. Dvorak, A. Aspuru-Guzik, J. E. Hein, and C. P. Berlinguette. Self-driving laboratory for accelerated discovery of thin-film materials. *Science Advances*, 6(20), 2020. doi: 10.1126/sciadv.aaz8867. URL <https://doi.org/10.1126/sciadv.aaz8867>.
- [17] MONTREAL.AI. Ascension five-contract suite: deployed source. Public repository source, 2026. URL <https://github.com/MontrealAI/AGI-Alpha-Agent-v0/tree/c0fc5217e364a39a4cd123d17a8317c06b80f820/contracts/ascension>. Deployed-source pin c0fc5217e364a39a4cd123d17a8317c06b80f820; retrieved 22 October 2026.
- [18] MONTREAL.AI. SUCCESSOR enterprise to Ascension journey: PR 4865 candidate. Public repository source, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/6bb361d522d80dc900cc2a480014531c6a44b66e/docs/agent/SUCCESSOR_ASCENSION_JOURNEY.md. Open and unmerged at retrieval; candidate pin 6bb361d522d80dc900cc2a480014531c6a44b66e; retrieved 22 October 2026.
- [19] MONTREAL.AI. Owner-designated current mainnet deployment and sanitized 18 October 2026 manifest. Public repository source, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/2b92933a5e3955440be770362acc18f06abf1708/docs/agent/CURRENT_MAINNET_DEPLOYMENT.md. Historical inventory observed 18 October 2026; documentation pin 2b92933a5e3955440be770362acc18f06abf1708; retrieved 22 October 2026; not current readiness.
- [20] MONTREAL.AI. AGI Fleet Linux: qualified execution, review and project continuity source. Public repository source, 2026. URL <https://github.com/MontrealAI/AGI-Fleet-Linux/blob/45759fb096faf8988f58ae5ea71248d0f35a97c7/README.md>. Source pin 45759fb096faf8988f58ae5ea71248d0f35a97c7; retrieved 22 October 2026; no installed or remote execution in this study.
- [21] MONTREAL.AI. SUCCESSOR enterprise harness candidate: PR 4863 source and runbook. Public repository source, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/c455ad4ce109c86df2fef09d19674810290b5185/docs/agent/SUCCESSOR_HARNESS.md. Open and unmerged at retrieval; candidate pin c455ad4ce109c86df2fef09d19674810290b5185; retrieved 22 October 2026.
- [22] MONTREAL.AI. AGI Jobs v0: JobRegistry protocol and bounded Alpha interoperability. Public repository source, 2026. URL <https://github.com/MontrealAI/AGIJobsv0/blob/ba67d5d19ef7d9ac643f99f2b8f49c160714fdf6/README.md>. Source pin ba67d5d19ef7d9ac643f99f2b8f49c160714fdf6; retrieved 22 October 2026; distinct from Ascension-JobMarket.
- [23] MONTREAL.AI. AGIJobManager: separate USDC work lifecycle and SUCCESSOR integration. Public repository source, 2026. URL <https://github.com/MontrealAI/AGIJobManager/blob/900bb4b46cb7de7f6f69f553a5687866f978a413/README.md>. Source pin 900bb4b46cb7de7f6f69f553a5687866f978a413; retrieved 22 October 2026; separate USDC protocol.

- [24] NIST/SEMATECH. e-Handbook of Statistical Methods, Section 2.3.3.1.2: Bias Caused by Instrument Drift, 2026. URL <https://www.itl.nist.gov/div898/handbook/mpc/section3/mpc3312.htm>. Retrieved 22 October 2026.
- [25] Malcolm Sim, Mohammad Ghazi Vakili, Felix Strieth-Kalthoff, Han Hao, Riley J. Hickman, Santiago Miret, Sergio Pablo-García, and Alán Aspuru-Guzik. ChemOS 2.0: An orchestration architecture for chemical self-driving laboratories. *Matter*, 7(9):2959–2977, 2024. doi: 10.1016/j.matt.2024.04.022. URL <https://doi.org/10.1016/j.matt.2024.04.022>.
- [26] The SciPy Community. scipy.integrate.solve_ivp: SciPy v1.17.0 manual. https://docs.scipy.org/doc/scipy-1.17.0/reference/generated/scipy.integrate.solve_ivp.html, 2026. Accessed 22 October 2026; DOP853 method and tolerance definitions.
- [27] Rafael Vescovi, Tobias Ginsburg, Kyle Hippe, Doga Ozgulbas, Casey Stone, Abraham Stroka, Rory Butler, Ben Blaiszik, Tom Brettin, Kyle Chard, Mark Hereld, Arvind Ramanathan, Rick Stevens, Aikaterini Vriza, Jie Xu, Qingteng Zhang, and Ian Foster. Towards a modular architecture for science factories. *Digital Discovery*, 2(6):1980–1998, 2023. doi: 10.1039/D3DD00142C. URL <https://doi.org/10.1039/D3DD00142C>.