A Governed Edge-of-Chaos Universe with Verified Adaptive Artificial Life under Controlled Dynamic Stress

Leroy A. Palmer
Project Borh
palmer@eormen.com eormen.com

March 2026

Abstract

Artificial-life claims are frequently undermined by poor control of the substrate on which the claimed life process operates. This paper addresses that problem by presenting an integrated result built in two bounded stages.

A companion publication (Palmer, 2026) characterised a proprietary spatiotemporal engine, treated strictly as a black box, through extensive multi-tradition diagnostic analysis. That work established, from observables alone, a reproducible dynamical phenotype: near-critical, weakly contracting dynamics close to the order-chaos boundary.

The present paper contributes two things. First, a governed universe platform was constructed on that characterised substrate: a structured, navigable environment with safe zones, corridors, terrain features, and danger boundaries, governed through closure gates, sentinel surveillance, and reproducibility controls. The platform was closed through a four-stage verification chain, all stages passing with explicit verdicts.

Second, a focused dynamic-stress programme tested whether Tovo, a proprietary population of adaptive artificial cells, could persist within the verified universe under controlled live-core perturbation. Artificial life is defined here as an adaptive population process exhibiting strictly local sensing, energy-coupled survival, death, heritable variation, conditional reproduction, lineage continuity, and environment-constrained behavioural adaptation. Tovo satisfies this criterion. The paper reports Tovo's operational interface and behavioural observables rather than its full implementation.

On the verified route, under a defined live-core pulse-train stress profile, the substrate accepted substantial perturbation (maximum local displacement of 32.4 field units, a 46-fold increase in local energy), recovered within the bounded contract horizon (52 steps against a 64-step limit), and supported Tovo adaptive persistence. All seven checks of the adaptive-persistence decision contract passed. The minimum verified claim used the pre-hospitality baseline profile and did not require enhanced environmental support.

The result is a bounded, route-specific verification of adaptive artificial-life persistence in a governed near-critical substrate-derived environment. It claims neither exact proof of ideal criticality nor broad universality of survival.

1 Introduction

Claims about artificial life are only as credible as the substrate in which the claimed life process is observed. If the substrate is poorly characterised, uncontrolled, or irreproducible, even complex behaviour remains scientifically weak. The harder task is to show that a governed universe can be built, verified, stress-tested, and then used as the setting for bounded life claims whose evidential basis is explicit.

Edge-of-chaos substrates are theoretically attractive. Langton (1990) argued that computation is maximised at the transition between ordered and chaotic regimes. Kauffman (1993) proposed that self-organisation and selection operate most effectively in near-critical regimes. Bertschinger and Natschläger (2004) demonstrated that recurrent neural networks at the edge of chaos achieve maximal memory capacity and computational expressiveness. These results suggest that a near-critical dynamical substrate could provide the right combination of structure, variation, and adaptive opportunity for artificial life.

That attraction, however, often leads to overstatement. Substrates are labelled "edge of chaos" on partial evidence. Governance is weak. Biological claims are made before the universe itself has been closed as a serious scientific system. The result is a literature in which interesting demonstrations are common but credible, governed results are rare.

This paper takes the opposite approach. It first establishes a governed universe platform, verified through explicit auditable gates, and only then presents a bounded adaptive-life result on that platform.

1.1 Operational definition of artificial life

In this paper, artificial life denotes an adaptive population process with strictly local sensing, energy-coupled survival, death from environmental or metabolic causes, heritable variation, conditional reproduction, lineage continuity, and environment-constrained behavioural adaptation. We do not claim biological life, open-ended evolution, or equivalence to natural organisms.

The life process studied here, Tovo, is a population of adaptive cells satisfying this criterion. Each cell senses only the local dynamical state at its current position, chooses actions from a heritable behavioural disposition, consumes energy, learns from outcomes, reproduces when conditions allow, and dies when they do not. Whether such a process can persist adaptively inside a governed near-critical universe under dynamic stress is the central question of this paper.

1.2 Contribution and claim structure

The paper makes three distinct claims, each bounded.

First, a governed universe platform can be constructed on the proprietary substrate characterised in the companion study (Palmer, 2026). The present paper contributes the universe construction, governance chain, and route-bounded adaptive-persistence result; it does not contribute the substrate characterisation.

Second, within that governed universe, a class of entities satisfying the operational criterion above can persist adaptively. The candidate is Tovo, a proprietary adaptive-life architecture. The claim is candidate-bounded, not family-exhaustive.

Third, this adaptive persistence survives a bounded protocol of controlled live-core dynamic stress on a verified route. The claim is route-specific, profile-specific, and contract-specific.

The engine substrate has been characterised independently in the companion publication (Palmer, 2026). The present paper builds on that characterisation but does not repeat it.

2 Background and Related Work

2.1 Edge-of-chaos dynamics and computation

The concept of computation at the edge of chaos originates with Langton (1990), who observed that cellular automata exhibit their richest computational behaviour at the transition between periodic and chaotic regimes. Kauffman (1993) extended this to biological self-organisation, arguing that natural selection operates most effectively in systems poised near the boundary between order and disorder. Packard (1988) provided early computational evidence that evolving systems can adapt toward the edge of chaos. Bertschinger and Natschläger (2004) placed these ideas on more precise footing by demonstrating that recurrent neural networks at the edge of chaos achieve the highest memory capacity and the best ability to separate distinct input histories, connecting near-critical dynamics to a concrete computational property.

These results motivate the use of near-critical substrates for artificial-life research. A system operating near the edge of chaos combines sufficient structure for persistent patterns with sufficient variation for adaptive behaviour. The challenge is to realise such a substrate in a governed, reproducible manner rather than merely invoking the concept as a metaphor.

2.2 Artificial life and substrate requirements

A recurring problem in artificial-life research is that demonstrations of interesting behaviour are rarely accompanied by rigorous characterisation of the substrate. May (1976) showed that simple mathematical models can produce complex dynamics, illustrating that apparent complexity does not by itself constitute evidence of meaningful biological organisation. Solé and Goodwin (2000) argued that the signs of life in complex systems require careful disentangling from the dynamics of the substrate itself. The implication is that artificial-life claims need to be grounded in substrates whose properties are independently established rather than inferred from the behaviour of the life process.

The present work addresses this gap. The substrate is characterised independently (Palmer, 2026). The governance layer is constructed explicitly. The life claim is made only after the substrate has passed its own verification chain.

2.3 Governance and reproducibility

Reproducibility in computational science requires more than the ability to re-run code. It requires explicit documentation of parameters, configurations, random seeds, and decision logic, together with automated verification that outputs conform to stated expectations. The governance framework used in this work includes SHA-256 hashing of all critical artefacts, a golden regression bundle with verified targets, automated quality gates, and machine-readable decision reports with explicit pass/fail verdicts.

2.4 Companion substrate result and epistemic boundary

The engine substrate has been characterised in a companion publication (Palmer, 2026) using thirty-two analytical methods drawn from dynamical systems theory, information theory, spectral and correlation analysis, and statistical physics. The companion paper treats the engine as a proprietary black box: the internal formulation, coupling topology, internal state variables, and parameter values are not disclosed. All behavioural claims are grounded in observable diagnostics alone.

The principal findings are as follows. The maximum Lyapunov exponent is weakly negative (approximately -0.0045, 95% CI $[-0.0062, -0.0022]$), indicating contracting dynamics near but not at zero. Permutation entropy is intermediate (0.440), placing the system between periodic order and maximal disorder. The power spectral density follows a $1/f^{\beta}$ scaling law with $\beta \approx 1.8$, consistent with long-range temporal correlations. The correlation dimension indicates low-dimensional attractor geometry. Perturbation recovery is rapid across all tested scales. Reference comparison with logistic-map regimes showed that the engine's diagnostic profile aligns most closely with the edge-of-chaos reference, intended as qualitative calibration rather than evidence of structural equivalence.

The established conclusion is that the engine operates in a near-critical, weakly contracting regime close to the order-chaos boundary. This is an empirical phenotypic characterisation, not a theorem-level proof of exact criticality. The present paper inherits this characterisation and the epistemic boundary that accompanies it: no claim made here depends on knowledge of the engine's internal equations.

3 System Architecture

The system comprises four distinct layers: substrate (raw dynamics), universe (spatial structure and navigability), governance (closure, surveillance, and reproducibility), and inhabitant (adaptive-life process). Evidence for each layer comes from different sources, and claims about each are supported independently.

3.1 Engine substrate

At the base of the system lies the proprietary spatiotemporal engine characterised in Palmer (2026). The engine operates on a one-dimensional spatial domain of 128 grid points evolving in continuous time via an adaptive-step numerical integrator. It is deterministic for fixed seed and initial conditions, and exposes observable quantities including the spatial mean, maximum amplitude, spatial standard deviation, energy, adaptive step size, and full spatial field snapshots.

The engine is treated throughout as an opaque deterministic substrate whose observable dynamical phenotype is established in the companion study (Section 2.4). The project does not claim access to or knowledge of the engine's internal equations, coupling topology, or parameter values. The companion study identifies two experimentally relevant controls: external perturbation injection and real-time operating-point modulation. Of these, the present paper uses external perturbation injection to deliver controlled stress to the live dynamical core.

3.2 The universe as a structured environment

The universe is not the raw numerical output of the engine. It is a governed environment constructed on top of that output.

The engine generates a spatiotemporal field in which each coordinate (x, t) carries a scalar value representing the local dynamical state. This field is not homogeneous. The project analyses it to identify spatial structure: safe zones (regions where the local amplitude remains within survivable bounds), corridors (connected paths between safer areas), terrain topology (how regions connect and constrain movement), and danger boundaries (transitions into regions classified as non-viable under the operational Tovo survival criterion).

The engine provides the raw dynamics. The universe adds the governed context: structural analysis, habitat classification, reproducibility controls, run manifests, and the conditions under which the field becomes a navigable environment for experimental artificial-life work.

A near-critical substrate sits between extremes. If too ordered, it provides insufficient variation for adaptive behaviour. If too chaotic, it is too destructive for persistent organisation. The near-critical regime offers enough structure for persistence and enough sensitivity for adaptive response (Langton, 1990; Kauffman, 1993; Bertschinger and Natschläger, 2004).

3.3 Governance layer

The governance layer provides four capabilities.

Closure logic enforces a staged verification chain in which each phase must pass explicit gates before the next begins. Decision artefacts are machine-readable JSON reports with explicit verdicts.

Sentinel surveillance provides continuous monitoring through a rolling window scheme, requiring the universe to demonstrate ongoing stability rather than passing a single check.

Reproducibility controls include SHA-256 hashing of all critical artefacts, a golden regression bundle with 26 hashed targets and 11 regression checks, deterministic random seeding, and run manifests recording all parameters and configurations. These ensure that any result can be reproduced within the governed programme environment.

Quality gates enforce automated testing. At final governance reconciliation, 164 pytest checks passed with zero failures (19 warnings, all non-blocking).

3.4 Tovo: the adaptive-life process

Tovo is the inhabitant layer of this work. It is a proprietary adaptive-life architecture. The paper reports its operational interface, evaluation criteria, and behavioural observables rather than a full implementation specification. Under the criterion defined in Section 1.1, Tovo qualifies as adaptive artificial life.

3.4.1 Operational interface

Each Tovo cell is characterised by the following publicly reported properties.

Strictly local sensing. A cell perceives only the local dynamical state at its current position. It does not receive a global map or any information about regions it has not visited. Discovery of environmental structure is experiential, not programmed.

Heritable behavioural disposition. Each cell carries a heritable genome encoding tendencies for local movement through the spatiotemporal field. Offspring

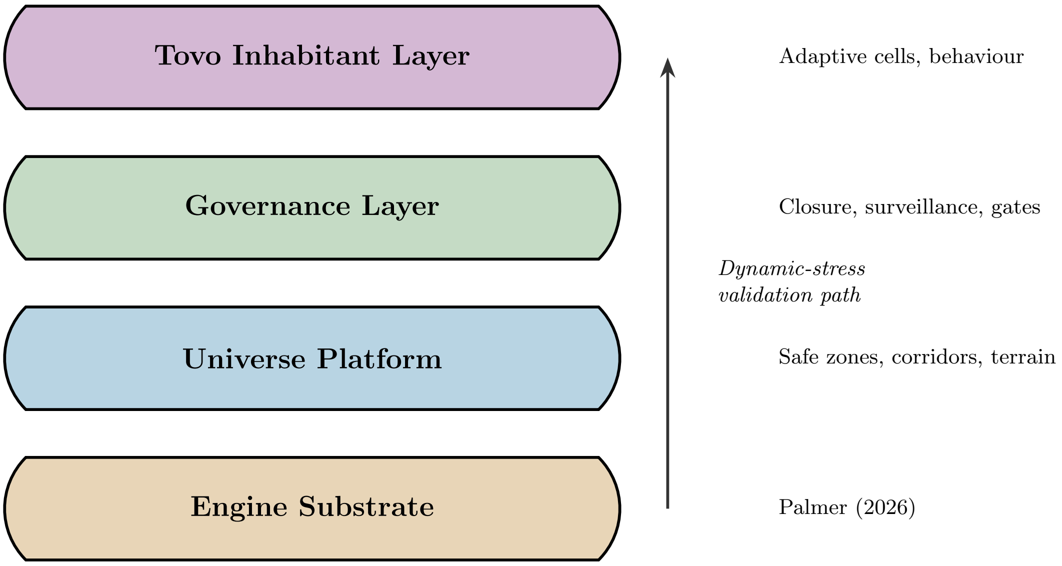

Figure 1: System architecture. The four-layer stack comprises the engine substrate (producing the spatiotemporal field), the universe platform (spatial structure including safe zones, corridors, terrain, and danger boundaries), the governance layer (closure logic, sentinel surveillance, and quality gates), and the Tovo inhabitant layer (adaptive cells executing the behavioural cycle). The dynamic-stress validation path spans all layers.

inherit a mutated copy of the parent's disposition, so that over generations lineages accumulate tendencies shaped by differential survival.

Bounded adaptive state. Each cell maintains a finite memory of recent experiences and a record of action outcomes. This accumulated state biases future decisions through a feedback loop: sense, act, survive or fail, adjust future behaviour accordingly.

Finite viability resources. Each cell has an energy budget that links action to survival. Movement consumes energy. Cells in stable regions of the universe recover energy passively. A cell that exhausts its energy dies. This creates a direct metabolic coupling between the cell and the local dynamical state of the substrate.

Conditional reproduction. Reproduction is not automatic. A cell can reproduce only if it has accumulated sufficient energy, reached a minimum age, and achieved a sufficient balance of successful actions and exploration. When reproduction occurs, the parent's energy is divided, the genome is copied with mutations, and an offspring is placed near the parent in space-time. Offspring inherit the lineage identifier.

State-dependent mortality. There are two routes to death. Energy depletion kills a cell when its budget reaches zero. Environmental lethality kills a cell when local dynamical conditions exceed a viability threshold. Mortality is discovered through consequence, not programmed as prior knowledge.

Lineage continuity. Each cell belongs to a lineage. Offspring inherit the lineage identifier of the parent, allowing tracking of whether lines of descent remain active through time.

3.4.2 Behavioural cycle

At each simulation step, every living cell executes a single cycle: it senses local conditions; selects an action from its heritable disposition under bounded exploratory variation; moves through the spatiotemporal field, consuming energy; updates its adaptive state from the outcome; checks survival; and, if alive and performance-linked conditions are met, reproduces with a mutated genome passed to the offspring. The combination of strictly local sensing, heritable disposition, adaptive state, and conditional reproduction means that Tovo behaviour is conditioned by inherited tendency, local conditions, and learned experience.

3.4.3 Population ecology

At the simulation level, Tovo operates as a distributed ecological population. The initial population is seeded across multiple regions of the universe. Each simulation step advances every living cell through the full behavioural cycle. The evolution manager removes dead cells, adds surviving offspring, and tracks population size, lineage activity, birth and death events, spatial distribution, and learning metrics over time.

3.4.4 Observable life criteria

Adaptive life is defined through deliberately minimal observable criteria: non-zero population at the end of the run, active lineages (lines of descent that have not gone extinct), observed births, observed deaths, and positive exploration or safe-zone discovery. These capture ecological turnover rather than mere transient presence.

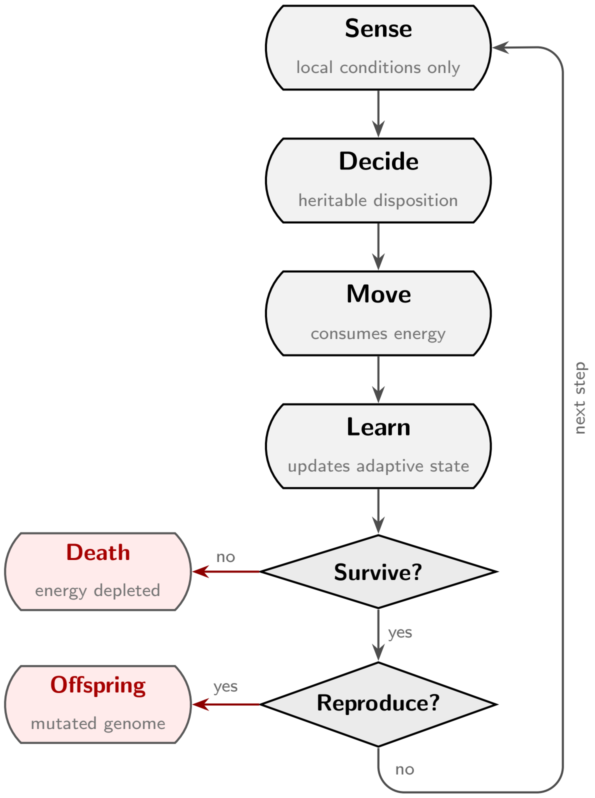

Figure 2: Tovo behavioural cycle. Each living cell executes a repeating cycle of sense, decide, move, learn, and survival check. Cells that survive and meet performance-linked conditions may reproduce, passing a mutated genome to offspring. Death occurs through energy depletion or environmental lethality.

3.4.5 Experimental profiles

Two Tovo profiles were tested. The pre-hospitality baseline profile operates with default environmental parameters. The hospitality-enabled profile adjusts these parameters to provide a more supportive local environment: reduced movement costs, enhanced passive energy recovery, increased stability bonus, and biased initial spawning toward safer regions. The minimum verified claim uses the pre-hospitality baseline, demonstrating that the governed near-critical substrate itself, without additional environmental support, provides conditions sufficient for minimal adaptive persistence.

3.5 Dynamic stress framework

The second phase replaced an earlier ordered-comparator search strategy with a dynamic-stress approach. The engine's characterised phenotype (temporally rich, spatially localised, experimentally controllable) made dynamic-stress testing the scientifically appropriate route: not whether a static comparator could be found, but whether adaptive life could persist under controlled perturbation of the substrate itself.

The stress framework uses localised live-core pulse trains injected directly into the engine's dynamical core within a defined habitat window. The profile used consists of three pulses of amplitude 8.0, spaced six steps apart with a pulse width of two steps and a spatial spread ($\sigma = 1.2$ grid points), centred on the habitat window. This perturbation modifies the engine's actual dynamical state, not merely a stored replay.

Recovery is assessed against a bounded contract: the substrate must return to within a defined tolerance of its baseline state within a specified recovery horizon (64 steps from the last active pulse).

4 Methods

4.1 Phase 1: governed universe closure

Phase 1 established the governed universe platform through a four-stage closure chain. Each stage produced a decision artefact with an explicit verdict. The stages are sequential: each must pass before the next begins.

Stage 1: Candidate screening and structural transferability (WP-30). A single promoted candidate was tested in triplicate across three conditions (holdout, stress, and support). Each condition required pipeline success rate, preflight strict gate rate, and structural quality pass rate to reach 1.0. All three passed. Verdict: `verified_transfer_evidence`.

Stage 2: Strict-causal coupling confirmation (WP-32). The same candidate was tested in triplicate to verify that Tovo coupling to the substrate was reproducible under strict-causal conditions. Each condition confirmed reproducible coupled behaviour, absence of uncontrolled failure propagation, and detection of a measurable coupling effect. All three passed. Verdict: `verified_Tovo_transfer_evidence`.

Stage 3: Sentinel surveillance block (WP-33). Three consecutive sentinel windows evaluated ongoing substrate stability after closure. All three passed, with zero failures and zero pending evaluations. Verdict: `passed`.

Stage 4: Governance reconciliation (WP-35). Final reconciliation included a fresh quality gate run (164 pytest checks passed, 19 non-blocking warnings), verification that all governance ledgers were consistent, and confirmation that no open placeholder language remained. Verdict: `verified_governance_reconciliation`.

Programme-level closure status: scientific closure verified (candidate-bounded), engineering governance verified, final sign-off verified.

4.2 Phase 2: dynamic stress protocol

Phase 2 used the verified platform as the substrate for a dynamic-stress adaptive-persistence test. The locked objective: can Tovo achieve adaptive persistence under controlled real-core dynamic stress in the verified universe?

Perturbation readiness calibration (P2-H10). The selected route was Day46, with warmup mode preserving the dynamical state from the preceding schedule. The habitat window spans spatial columns 36 to

44, classified as "edge habitat" with a safe fraction of 1.0 and an edge-habitat score of 0.78. This means the window sits at the boundary between safe and transitional dynamics. This stage confirmed operational suitability for perturbation.

Recovery baseline establishment (P2-H11). The stress profile was applied to characterise perturbation-response behaviour. Key parameters: three pulses of amplitude 8.0, spatially centred at grid point 40, spatial $\sigma = 1.2$. The recovery contract specified return to baseline within a maximum evaluation window of 512 steps, with operational recovery horizon at 64 steps from the last active pulse.

Live-core Tovo campaign (P2-H12). Both profiles (pre-hospitality baseline and hospitality-enabled) were run in triplicate on matched baseline and stressed substrate conditions using identical seed sets. Matching was intended to isolate the effect of stress from seed-level stochastic variation.

4.3 Adaptive-persistence decision contract

The contract specifies seven binary checks, all of which must pass for the verdict to be positive:

1. **Measurable stress.** The perturbation must produce measurable displacement from baseline.
2. **Bounded recovery.** The substrate must recover within the carried horizon.
3. **Non-zero population.** At least one run must end with population > 0.
4. **Active lineages.** At least one run must retain active lines of descent.
5. **Births observed.** At least one run must exhibit births under stress.
6. **Deaths observed.** At least one run must exhibit deaths under stress.
7. **Exploration or safe-zone discovery positive.** At least one run must show positive exploration or safe-zone discovery.

4.4 Justification of the decision contract

The seven checks capture the minimal observable signature of an adaptive population process as defined in Section 1.1.

Checks 1 and 2 establish that the stress is genuine and the substrate remains viable. Without these, the life claim would rest on an unperturbed or destroyed substrate.

Checks 3 and 4 establish that life persisted through the stress period. Active lineages confirm that lines of descent continued through the perturbation, not merely that individuals survived.

Checks 5 and 6 confirm ecological turnover. A population that neither reproduces nor dies is static, not adaptive.

Check 7 confirms active spatial behaviour rather than immobility.

The threshold "at least one of three runs" is appropriate because the claim is an existence result, not a statistical generalisation. The triplicate design provides context for assessing fragility without requiring universal success. For the pre-hospitality baseline, verification therefore establishes route-bounded viability rather than broad robustness.

5 Results I: Governed Universe Closure

5.1 Structural transferability

The promoted candidate exhibited reproducible structural quality under holdout, stress, and support conditions. All nine individual runs achieved pipeline success, preflight strict gate, and structural quality rates of 1.0. Verdict: `verified_transfer_evidence`.

5.2 Strict-causal coupling confirmation

Tovo interaction with the substrate operated under strict-causal conditions on the verified closure branch. All nine runs confirmed reproducible coupled behaviour, absence of uncontrolled failure propagation, and a measurable coupling effect. Verdict: `verified_Tovo_transfer_evidence`.

This result establishes that the Tovo layer is not merely conceptually attached to the universe. The coupling produces a detectable, reproducible effect under controlled conditions. This is the causal foundation on which the later adaptive-persistence claim rests.

5.3 Sentinel surveillance block

Three consecutive sentinel windows spanning the post-closure monitoring period were evaluated (Table 1).

Table 1: Sentinel surveillance block results.

Window	Pipeline success	Preflight gate	Structural quality	Pass
CB	1.0	1.0	1.0	Yes
CC	1.0	1.0	1.0	Yes
CD	1.0	1.0	1.0	Yes

All three passed with zero failures. The universe remained stable under post-closure surveillance, not only under development conditions.

5.4 Governance reconciliation

Final reconciliation passed: 164 pytest checks (zero failures, 19 non-blocking warnings, 21.0 seconds). All ledger consistency, execution log, and closure documentation checks passed. Verdict: `verified_governance_reconciliation`.

5.5 Phase 1 closure summary

Table 2 summarises the four-stage closure chain.

6 Results II: Adaptive Persistence under Dynamic Stress

6.1 Route selection

The engine characterisation (Palmer, 2026) established that the substrate is temporally rich, spatially localised, and experimentally controllable. These properties made dynamic-stress testing the appropriate route. The locked parameters were: Day46, warmup mode preserving the preceding dynamical state, habitat window spanning spatial columns 36 to 44, a defined pulse-train stress profile, recovery horizon of 64 steps, and the pre-hospitality baseline Tovo profile for the minimum claim.

6.2 Perturbation readiness

The selected habitat window was confirmed as operationally suitable: “edge habitat” with a safe fraction of 1.0, connectivity of 1.0, and an edge-habitat score of 0.78. The window sits at the boundary between safe and transitional dynamics, where adaptive persistence is neither trivially easy nor immediately impossible.

6.3 Stress and recovery

The stress produced a maximum absolute local displacement of 32.4 field units, relative to a baseline local mean absolute value of 1.12. Local mean energy rose from 1.25 to 57.1, a factor of approximately 46. The perturbation was genuine and substantial.

The three pulses produced a last active pulse index of 33 and a recovery index of 85, giving a recovery span of 52 steps. This falls within the carried horizon of 64, confirming that the substrate’s perturbation-response satisfies the bounded recovery contract.

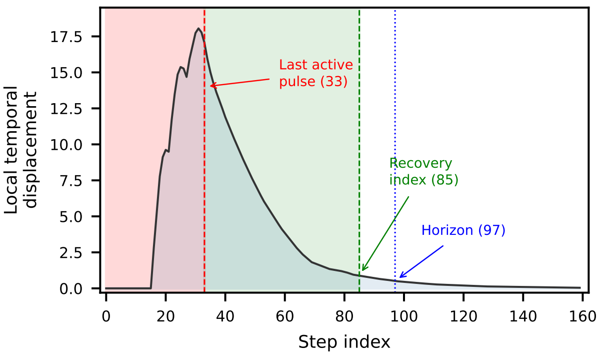

Figure 4: Perturbation and recovery time series. The three live-core pulse injections (steps 16, 22, 28) produce substantial local displacement. The last active pulse index is 33, the recovery index is 85, and the carried recovery horizon extends to step 97. The substrate returns to baseline well within the contracted horizon.

6.4 Tovo campaign results

Table 3 presents the adaptive-persistence decision contract results.

Table 3: Adaptive-persistence decision contract results.

Check	Pre-hosp. baseline	Hosp.-enabled
Measurable stress	Passed	Passed
Bounded recovery	Passed	Passed
Non-zero population	Passed	Passed
Active lineages	Passed	Passed
Births observed	Passed	Passed
Deaths observed	Passed	Passed
Exploration/safe-zone	Passed	Passed
Overall	**Passed**	**Passed**

Table 4 presents the quantitative comparison between profiles under stressed conditions.

The hospitality-enabled profile produced substantially higher population sizes, lineage counts, and exploration rates. The population delta under stress was −15.0 relative to the unstressed baseline, indicating a measurable but bounded stress effect.

The pre-hospitality baseline produced much smaller populations and fewer active lineages, but still satisfied all seven checks of the route-bounded decision contract. This profile represents persistence at the margin of viability: two of three runs retained a non-zero population with at least one active lineage. Deaths were observed in every run (rate 1.0), confirming that stress materially affected the population. Births occurred in at least one run (rate 0.33), confirming ongoing reproductive activity.

One of three pre-hospitality runs did not retain a non-zero population. The claim is a bounded existence result: verified adaptive persistence was achieved on the defined route, not in all runs or all conditions.

Table 2: Phase 1 closure chain.

Stage	Objective	Verdict	Key metric
1. Candidate screening (WP-30)	Structural transferability	`verified_transfer_evidence`	All rates 1.0
2. Coupling confirmation (WP-32)	Strict-causal Tovo coupling	`verified_Tovo_transfer_evidence`	Coupling detected
3. Sentinel surveillance (WP-33)	Post-closure stability	`passed`	Zero failures
4. Governance reconciliation (WP-35)	Final quality gate	`verified_governance_reconciliation`	164 tests passed

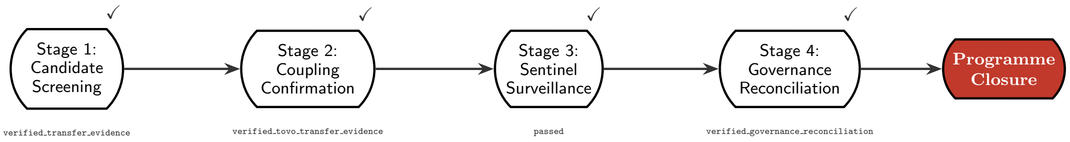

Figure 3: Phase 1 verification chain. The four sequential stages must each pass before the next begins. The chain terminates in programme-level closure.

6.5 Adaptive-persistence verdict

The decision artefact records the verdict as `verified_live_core_adaptive_persistence`, all checks passed. The selected profile for the minimum claim was the pre-hospitality baseline. Hospitality uplift was not required.

On the verified route, the governed near-critical substrate itself, without additional environmental support, provided conditions sufficient for minimal adaptive persistence.

7 Discussion

7.1 What this paper shows

The paper demonstrates three things. First, a governed universe platform was constructed on the characterised proprietary substrate (Palmer, 2026). The construction is not conceptual: it produced a four-stage closure chain with auditable verdicts, sentinel surveillance, and automated quality checks.

Second, Tovo adaptive life, satisfying the operational criterion in Section 1.1, achieved bounded persistence within that governed universe under controlled live-core dynamic stress. The minimum claim uses the pre-hospitality baseline and does not depend on enhanced environmental support.

Third, the result is governed and internally reproducible. Every parameter, configuration, seed, and decision is recorded within the governed programme environment. The result is fully auditable from the recorded artefacts, although the paper does not disclose the full proprietary inhabitant implementation required for external mechanistic reconstruction.

7.2 Interpretation and significance

Most artificial-life demonstrations proceed without a governed substrate. A simulation is configured, interesting behaviour is observed, and a claim is made. Substrate properties are either assumed or inferred from the life process itself, which is circular. The present work breaks this circularity by independently characterising the substrate and independently governing the universe platform before making any life claim.

The dynamic-stress result is stronger than a static observation of activity in undisturbed conditions. The protocol required the substrate to be measurably perturbed (32.4 field units displacement, 46-fold local energy increase), to recover within a bounded horizon, and to continue supporting adaptive life through that disruption. Tovo cells, sensing only local conditions and receiving no global information about the perturbation, retained population, lineage, reproductive activity, and spatial exploration. This is a more demanding test than observation of activity in benign conditions.

The pre-hospitality result establishes that the governed near-critical substrate itself provides conditions sufficient for minimal adaptive persistence on the verified route. The hospitality-enabled profile then shows that more supportive conditions produce substantially more robust persistence (mean final population 116.7 versus 2.0), providing information about the sensitivity of persistence to environmental conditions without undermining the minimum claim.

As with the substrate, the inhabitant process is interpreted through its operational interface and recorded behaviour rather than through full public disclosure of its internal implementation.

The connection to theoretical predictions is suggestive but bounded. Bertschinger and Natschläger (2004) predicted that systems at the edge of chaos achieve maximal computational capacity. The engine exhibits a diagnostic profile consistent with near-critical dynamics. Whether this connection reflects a deeper principle linking near-critical dynamics to conditions for adaptive life, or is specific to this engine and Tovo configuration, remains open.

Table 4: Tovo profile comparison under dynamic stress (triplicate means).

Metric	Pre-hospitality baseline	Hospitality-enabled	Unstressed baseline
Survival rate	0.67	1.00	1.00
Active lineages (mean)	1.33	35.0	35.0
Final population (mean)	2.0	116.7	131.7
Exploration (mean)	0.14	0.68	—
Safe zones discovered (mean)	1.0	8.0	—
Birth observation rate	0.33	1.00	—
Death observation rate	1.00	0.33	—

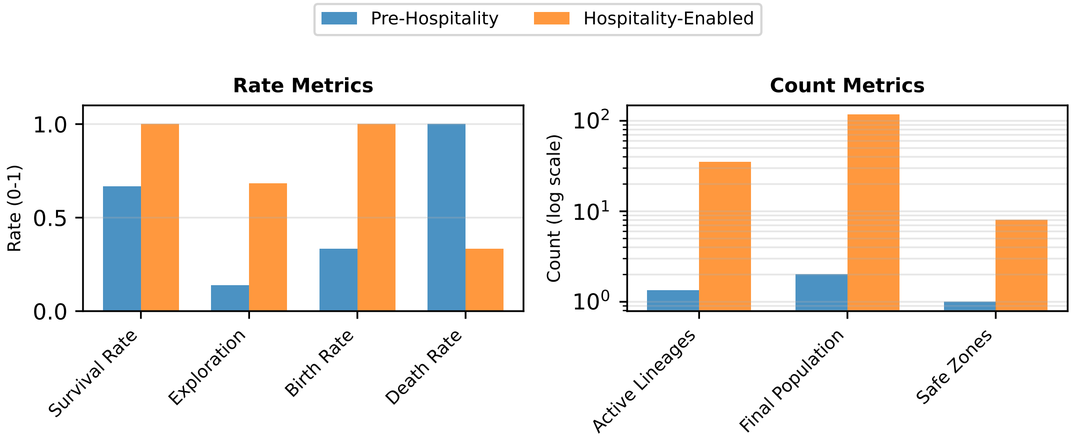

Figure 5: Tovo profile comparison under dynamic stress. Left: rate metrics (0–1 range). Right: count metrics (log scale). The hospitality-enabled profile produces substantially more robust persistence, but the pre-hospitality baseline still passes all seven decision checks.

7.3 What remains open

The result does not establish universal robustness, prove that all Tovo configurations would persist, demonstrate open-ended evolution, or confirm that the substrate occupies a mathematically ideal edge-of-chaos point. Future work can build on the verified governed foundation to investigate broader operating ranges, richer inhabitant architectures, ecological structure and lineage dynamics, and the sensitivity of persistence to stress intensity and habitat configuration.

7.4 Explicit non-claims

We do not claim biological life, open-ended evolution, universal robustness, exact proof of ideal criticality, or full public implementation-level reproducibility of the proprietary Tovo architecture from this manuscript alone. We claim a governed universe construction and a bounded, verified adaptive artificial-life persistence result within it.

8 Limitations

Candidate-bounded scope. Structural breadth is candidate-bounded, not family-exhaustive. The promoted candidate was selected through a defined promotion process; the paper does not claim that all candidates would yield the same result.

Route-specific stress result. The dynamic-stress result is bounded to specific parameters: the Day46 route, a defined warmup mode, a specific habitat window, a defined pulse-train stress profile (three pulses of amplitude 8.0), recovery horizon of 64 steps, and the pre-hospitality baseline for the minimum claim.

Empirical substrate characterisation. The engine characterisation (Palmer, 2026) is empirical, not theorem-level. The black-box constraint means no Jacobian-based or eigenvalue-based confirmation of exact criticality is available.

Black-box mechanism limitation. Because the substrate internals remain undisclosed, the paper cannot attribute the observed behaviour to a unique underlying mechanism. The result is empirical and operational, not mechanistically complete.

Proprietary inhabitant limitation. The paper does not disclose the full Tovo implementation. The manuscript supports operational and behavioural evaluation of the reported result but does not by itself enable external mechanistic reconstruction of the inhabitant process. Internal reproducibility and auditability are maintained within the governed programme environment.

Pre-hospitality fragility. The pre-hospitality baseline survival rate of 0.67 (two of three runs) means persistence was not universal even on the verified route. The decision contract requires at least one run to pass each check, not all runs.

Sample size. Triplicate runs per profile per condition are sufficient for the bounded existence claim but do not support detailed statistical inference about population-level distributions.

Unverified transformational architecture. The transformational single-cell Tovo architecture was not tested in the verified Phase 2 result. The verified claim applies to the population-level evolutionary architecture described in Section 3.4.

Operational verification contract. The adaptive-persistence contract is an operational framework for this study, not a universal definition of life. Different criteria or thresholds might yield different verdicts.

9 Conclusion

Building on prior black-box characterisation of a proprietary spatiotemporal engine as a reproducible near-critical, weakly contracting dynamical system close to the order-chaos boundary (Palmer, 2026), this paper demonstrated that a governed universe platform can be constructed on that substrate, closed through four sequential verification stages, and maintained under sentinel surveillance.

Within a verified bounded route of that governed universe, a population of adaptive artificial cells (Tovo), satisfying an explicit operational criterion for artificial life, achieved persistence under controlled live-core dynamic stress. The substrate accepted substantial perturbation (32.4 field units displacement, 46-fold local energy increase), recovered within the bounded contract horizon, and continued to support adaptive persistence. All seven checks of the decision contract passed. The minimum claim, using the pre-hospitality baseline, did not require enhanced environmental support.

The result is a bounded, governed, internally reproducible existence proof of adaptive artificial-life persistence in a near-critical substrate-derived world under controlled dynamic stress.

References

Bertschinger, N. and Natschläger, T. (2004). Real-time computation at the edge of chaos in recurrent neural networks. *Neural Computation*, 16(7):1413–1436.

Kauffman, S.A. (1993). *The Origins of Order: Self-Organization and Selection in Evolution.* Oxford University Press, New York.

Langton, C.G. (1990). Computation at the edge of chaos: phase transitions and emergent computation. *Physica D: Nonlinear Phenomena*, 42(1–3):12–37.

May, R.M. (1976). Simple mathematical models with very complicated dynamics. *Nature*, 261(5560):459–467.

Packard, N.H. (1988). Adaptation toward the edge of chaos. In J.A.S. Kelso, A.J. Mandell, and M.F. Shlesinger, editors, *Dynamic Patterns in Complex Systems*, pages 293–301. World Scientific, Singapore.

Palmer, L.A. (2026). Black-box empirical characterisation of a novel proprietary spatiotemporal engine near the order-chaos boundary. `https://doi.org/10.5281/zenodo.19066737`.

Solé, R.V. and Goodwin, B. (2000). *Signs of Life: How Complexity Pervades Biology.* Basic Books, New York.