FORKED

Observer-Indexed Reality, Suspension Lattices, and Mission-State Integrity

Breyden E. Taylor - Inventor and Architect, Prompted LLC

Version 2.0 Foundational Whitepaper - July 2026

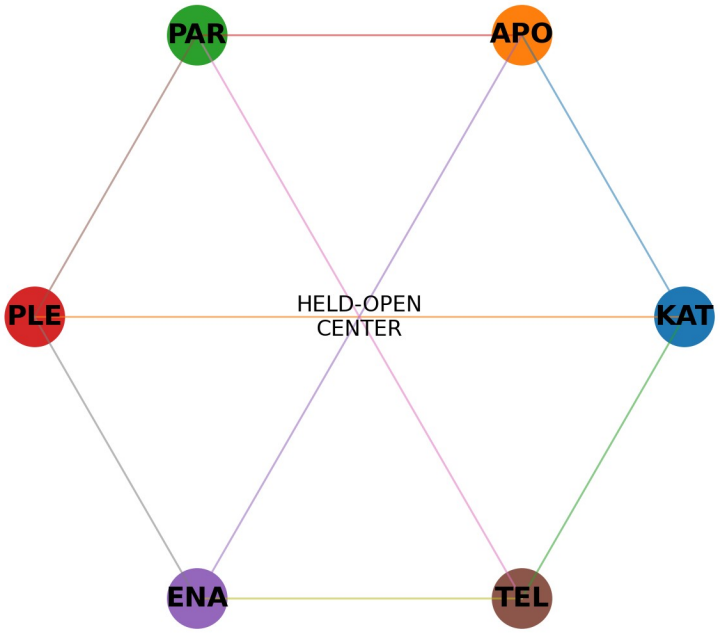

Facets constrain the inhabitable field; no path targets the center.

A Ubiquity Federation estate for governing state, belief, agency, correction, and lawful closure across physical, virtual, temporal, economic, ecological, and institutional systems.

This document supersedes the architectural framing of the prior FORKED v1.0 whitepaper. It is a new construction from the Ubiquity parent substrate, current FORKED native implementation, source artifacts, executable components, and bounded evidence register. The earlier paper remains lineage, not authority.

Contents

Abstract

FORKED is an internally owned and managed estate of the Ubiquity Federation, Prompted LLC's flagship governance substrate for sovereign adaptive systems. It specializes the parent architecture for observer-indexed reality: the condition in which physical, virtual, institutional, and human actors do not receive a universal state directly, but encounter bounded views transformed by sensors, interfaces, models, language, authority, memory, timing, and location. The estate governs the path by which those views become inferences, declarations, actions, records, and future priors.

The foundational object is not a vector, score, label, or prompt. It is a receipt-bearing **splat lattice**: a conformation-bound lawful slice of a governed possibility field. Each splat carries six distinct facets - KAT, APO, PAR, PLE, ENA, and TEL - together with evidence lanes, source-tense, authority, lifecycle, live and suspended hypotheses, forbidden paths, correction triggers, and one or more movable working centroids. A held-open center remains categorically outside the inhabitable state space. The architecture therefore permits local closure without allowing a local interpretation to crown itself as the whole.

FORKED adds a **suspension lattice** around that object. The lattice uses typed cables, winches, clamps, shock absorbers, fuses, and responsibility classes to regulate when a branch should remain open, when more evidence is required, when action is lawful, and when continued hesitation becomes the more dangerous failure. The relevant threshold is not confidence alone. Stabilization is a scoped meet across evidence, exclusions, authority, standing, currentness, dependency coverage, meaningful agency, correction capacity, economic feasibility, and mission horizon. Declaration and terminalization remain separate acts.

Memory is terrain-shaped rather than recency-first. Prior failures, fables, and precedents receive standing through conformation of causal topology, authority geometry, trajectory phase, point-of-no-return proximity, observer conditions, and horizon. Recency modifies currentness inside a bounded envelope; it does not erase an older shape-isomorphic failure or allow a recent but shape-distant event to dominate. Fables are treated as compact trajectory splats with wisdom and caution in equal standing. They raise the cost of traversal near known failure terrain without becoming proof, command, or permanent suspicion.

The same architecture extends across wider native lanes. Observer state is separated into physical, functional, mission, sensed, rendered, inferred, declared, verified, and remembered state. Sight is decomposed into geometry, sensor envelope, sampling, transport, rendering, attention, classification, acknowledgment, and action. Source-time is bitemporal. Human presence is distinguished from meaningful agency and from liability. Mission, battle, campaign, war, strategic, institutional, and civilizational horizons remain separate quivers. Ecological Physarum/Lichen simulations, HUNGER scarcity pressure, held-anchor economic cadence, detached shadow economies, citizen lifecycles, cross-biome visas, sensory TTLs, remote-compute egress, and cache-contamination response are not analogies appended after the fact; they are current native expression lanes that instantiate the same parent mechanics under different physics and authority.

The paper makes capability status explicit. The Ubiquity parent architecture, production splat mechanics, lifecycle surfaces, source-tense discipline, center exclusion, and lawful-traversal grammar are established capabilities of the originating federation. The FORKED Native Range v1.1 implements the specialized estate locally with executable schemas, tests, simulations, receipts, and a non-canonical parent gate. Component and synthetic campaigns establish bounded behavior of those implementations. Human-subject, hardware, customer, field, and operational performance remain separate empirical lanes. The paper neither understates what is running nor silently expands an internal capability into a universal performance claim.

The public release is an assurance and research architecture. It does not publish operational targeting logic, weapon-employment instructions, signature recipes, or deployable deception optimization. Its defense relevance lies in mission-state integrity, contested verification resilience, correction survival, and the preservation of human and institutional agency when reality is mediated by increasingly autonomous systems.

Keywords: Ubiquity Federation; FORKED; observer-indexed reality; splat mechanics; fractal quivers of quivers; suspension lattice; lawful traversal; mission-state integrity; meaningful human agency; provenance; bitemporal state; fable terrain; correction graph; Physarum; assurance.

Executive statement

FORKED governs the branch before the branch becomes the world.

A system acts on a state only after some observer has converted the world into a reachable representation. That observer may be a human body, camera, radar, event sensor, database, model, institutional report, user interface, or memory surface. Every observer is located. Every observer has limited resolution, authority, time, attention, budget, and history. The system's central problem is therefore not simply whether an answer is accurate. It is whether the architecture can say which state is physical, which is observed, which is inferred, which is projected, which is declared, which is verified, which is remembered, and which may lawfully move the world.

FORKED is built for that problem. It does not attempt to remove ambiguity by collapsing all lanes into one fused confidence. It constructs enough local order to act while preserving the distinctions that correction will later need. It can close a mission-local branch after the relevant threshold is met without closing the held-open center. It can hold a hypothesis open without turning uncertainty into a virtue. It can impose additional inspection because a trajectory is approaching a known failure surface without converting a fable into accusation. It can preserve meaningful human control without making a human click absorb causality the surrounding system already exercised.

The parent matters. Ubiquity is the governance substrate; FORKED is one estate inside it. Fractal Quivers of Quivers supplies the living topology of entities, paths, authority, standing, receipts, costs, absorbers, and telos. Production splat mechanics supplies the lawful-slice primitive. FORKED specializes those mechanics around observer divergence, state transition, suspension, declaration, correction, and mission horizons. The relationship is inheritance, not renaming. The estate can recurse and productize its expression. It cannot claim the parent as a FORKED-only invention or promote a local result into federation-wide truth.

The implementation is already concrete. The current native range carries a six-node parent manifest, six-facet splats, shape-first terrain retrieval, longitudinal trajectories, an executable fable field, responsibility mechanics, DAGs of DAGs, action and delay clamps, mission geometry, SightState, TTL-bound sensory buffers, cache quarantine, citizen lifecycles, HUNGER, a four-act Physarum/Lichen biome, a held-anchor economy, a detached shadow economy, non-speculative pricing, JSON Schemas, run receipts, and a release manifest. Its integrated sequence reaches hold -> review -> declare without granting the local package canonical write authority. The implementation is a bounded estate instance, not a speculative diagram. [1], [2]

The evidence register is equally explicit. Native tests and component selftests establish that the supplied code does what the tests exercise. Synthetic campaigns show how the mechanisms behave under generated distributions. Those campaigns do not become human, field, hardware, customer, or operational evidence by size. The architecture records the difference rather than using cautious language everywhere or triumphant language nowhere.

Load-bearing claims

The operating object is a splat lattice, not a vector. Embeddings and scores may navigate the field, but the field retains facets, exclusions, receipts, authority, unresolved branches, and correction conditions.

The center remains open while branches may close. Center exclusion is an anti-capture condition, not a rule against decisions.

Stabilization is responsibility-shaped. Required evidence and mechanical friction scale with reversibility, blast radius, horizon, authority ambiguity, economic envelope, and correction difficulty.

Action and delay are dual hazards. A mature system prevents premature closure and trauma-driven non-closure.

Terrain shape has more standing than recency alone. Currentness matters, but the age of a precedent cannot substitute for conformation.

Fables are trajectory carriers. They can reweight attention and cost without becoming proof or command.

The human is a causal boundary, not a ceremonial click. Meaningful agency requires inspectability, challenge, alternatives, delay, reversal, and liability legibility appropriate to the responsibility class.

DAGs discharge bounded work; quivers preserve living recurrence. The architecture uses both and does not flatten one into the other.

Mission horizons are different axes. Task, mission, battle, campaign, war, and institutional consequence require separate authority and telemetry.

Capability and evidence are different fields. What Ubiquity runs is stated as capability. What experiments support is stated at its tested scope.

How to read capability language

Status

Meaning

Permitted wording

Established in Ubiquity

Present in the originating federation's current architecture, runtime, public canonical surface, or executable source.

"Ubiquity implements…" or "The federation uses…"

Implemented in FORKED Native

Present in the v1.1 local implementation and exercised by tests or integrated execution.

"FORKED Native implements…"

Component verified

A bounded component passed its declared selftest or verification harness.

"The component passed… under the recorded environment."

Synthetic observed

A generated campaign produced a reported measurement.

"In the synthetic campaign…"

Research open

Testable proposition not yet established at the human, hardware, field, customer, or operational level.

"The research lane asks whether…"

Release restricted

Capability or detail intentionally excluded from the public artifact.

"The public release does not include…"

This typing avoids two symmetric errors. The first is false humility: describing a running mechanism as merely hypothetical because external replication has not yet occurred. The second is scope laundering: treating an internal implementation or synthetic observation as universal evidence. FORKED refuses both.

Contributions

This whitepaper contributes a new parent-faithful construction of FORKED:

a formal estate boundary inside Ubiquity;

a splat-lattice definition that retains the intermediate structure between latent terrain and natural language;

a six-facet, partially recursive, non-totalizing compute object;

a mechanical suspension architecture with responsibility-indexed cables, winches, clamps, absorbers, and fuses;

a shape-first temporal memory model in which recency is bounded by terrain conformation;

a fable field that preserves both wisdom and caution while attaching longitudinal inversion and PONR mechanics;

observer-indexed state, SightState, bitemporal source-time, and correction-residue models;

meaningful-agency and liability receipts for human-AI collaboration;

mission-horizon geometry that prevents tactical or local states from acquiring strategic authority silently;

ecological, economic, sensory, institutional, and compute expressions of the same parent mechanics;

a current native implementation and evidence register;

an assurance, benchmark, publication, and product roadmap with falsifiers and release boundaries.

Part I - The object and its boundary

1. The denominator is accepted state transition

The visible trigger for FORKED was the language of cost per kill, next-generation sensors, low-cost interceptors, and open-system hardware. That language invites founders to optimize physical effectors and unit economics. Those are legitimate lanes, but they conceal a more general denominator: a physical or virtual event becomes operationally consequential only when an observer or institution accepts a state transition and changes behavior because of it.

An object may be damaged while remaining mission-capable. A sensor may lose signal while the object remains present. A video may show an explosion without establishing physical destruction. A model may declare a terminal state that the operator accepts. A report may preserve that declaration after the underlying evidence changes. A digital twin may continue to display an object that no longer exists, or erase one that still matters. The decisive transition is therefore not simply object -> destroyed. It is:

```
world state
  -> observable surface
  -> rendered state
  -> inferred state
  -> declared state
  -> accepted state
  -> action and record
```

FORKED governs this chain. It asks what evidence supports each transition, who had standing to make it, which transformations intervened, what remained hidden, what authority the declaration acquired, what downstream nodes inherited it, and what would be required to correct it. In military assessment doctrine, physical damage, functional damage, and effects on a wider target system are already distinct assessment objects; the architecture generalizes the need to preserve such distinctions rather than substituting a single persuasive cue for the state itself. [3]X

The same denominator appears outside defense. A medical image becomes a diagnosis; a fraud alert becomes an account freeze; a risk score becomes denied credit; a software finding becomes a production block; a sensor anomaly becomes a shutdown; a generated summary becomes organizational memory. In each case, the primary object is not merely the observation. It is the transition by which a bounded view acquires standing to move resources, people, code, or policy.

This is why observer-indexed reality is not relativism. The architecture does not claim that every view is equally true. It claims that every consequential system must account for the transfer function between the world and the state upon which it acts. Physical constraints still bind. Records still have provenance. Authorities still have jurisdiction. Some hypotheses are disproven. Some paths are forbidden. The observer index tells the system where and how a claim was reachable; it does not dissolve the difference between supported and unsupported claims.

The strategic object is the closure function: when does the human-machine-institutional system stop asking what happened and begin acting as though the transition is settled?

2. FORKED is an estate, not the parent

Prompted LLC publicly defines Ubiquity as the governance substrate for sovereign adaptive systems: runtime infrastructure intended to increase AI-mediated capacity without collapsing agency, authorship, judgment, or meaningful contribution. The public stack describes signals, warrants, rules, anchors, rays, absorbers, routing, and operational governance surfaces. [4], [5], [6]X

FORKED inherits that substrate. It does not replace it.

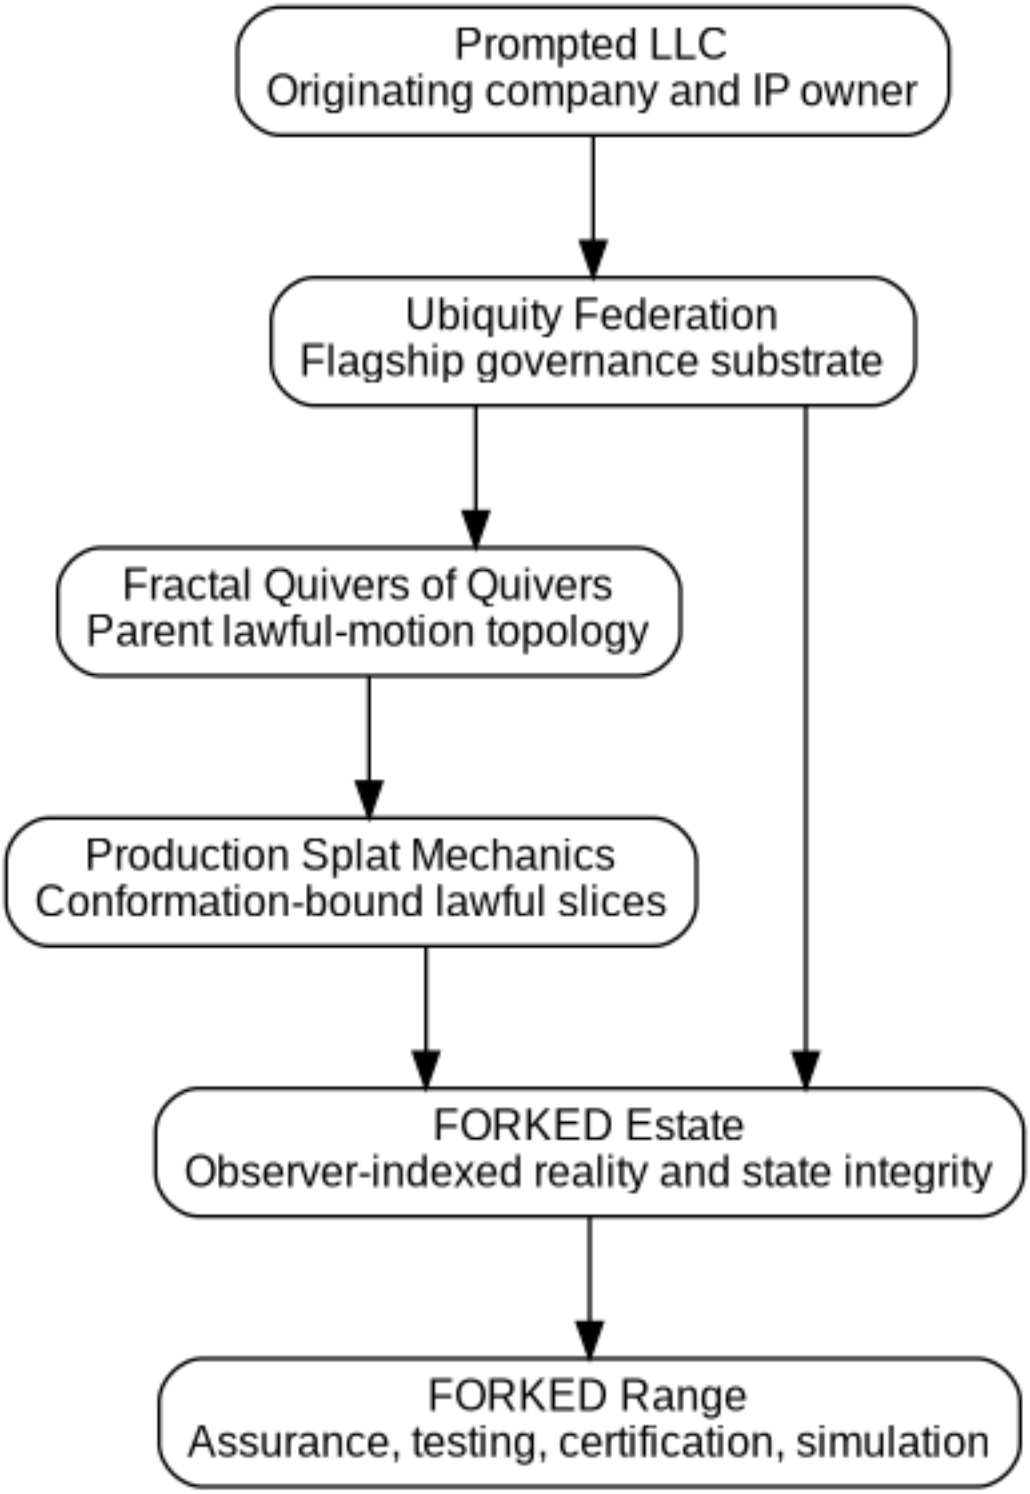

Figure 1. Parent-to-estate ordering. FORKED inherits the lawful-motion topology and production splat primitive; the bounded range is a local instance.

The inheritance relation has four consequences.

First, FORKED's classifier is not a specialized model that labels situations as deceptive, dangerous, or closed. Classification remains lawful traversal: the set of actions, holds, reviews, declarations, refusals, and corrections that remain admissible after anchors, exclusions, authority, standing, currentness, lifecycle, receipts, cost, and telos shape the field. [7]

Second, the splat is not invented inside FORKED. Production splat mechanics is a federation-wide primitive. FORKED specializes it for observer state, declaration integrity, belief divergence, human agency, mission horizons, and contested verification.

Third, local FORKED instances do not gain canonical authority merely because they execute successfully. The current native package explicitly sets canonical_writer = false and passes results through a parent gate. This prevents a range result, local model, projection, or test harness from modifying the originating federation's constitutional state. [1], [2]

Fourth, the estate may recurse fractally. FORKED Range, FORKED Assurance, FORKED Witness, FORKED Ledger, FORKED Halo, and domain-specific certification programs can each become bounded sub-estates or products. Their local innovations remain attributable to FORKED, but the parent mechanics retain their Ubiquity lineage.

3. The current capability boundary

The architecture must be read at the correct source-tense.

Established in Ubiquity: The parent federation currently carries governed quiver topology, production splat mechanics, held-open center doctrine, signals and warrants, lifecycle and promotion gates, source-tense distinctions, governance routing, correction-bearing memory, and public runtime surfaces. These are not proposals in this paper. They are the architecture from which the paper is written. [4], [7], [8], [9]

Implemented in FORKED Native: The v1.1 range currently implements the estate manifest, parent gate, six-facet splat builder, partially recursive lattice, terrain matcher, trajectory tracker, fable adapter, responsibility mechanics, DAGs of DAGs, threshold engine, suspension controller, mission geometry, observer and SightState models, sensory and cache lanes, citizen lifecycle, HUNGER, biome, economy adapters, schemas, receipts, and release verification. [1], [2]

Component verified: The supplied Aesop, economy, cadence, and pricing components pass their recorded selftests. The local package passes its test and verification suites. This establishes implementation behavior under the recorded environment. It does not prove every environment, integration, or threat model.

Synthetic observed: The experimental and native campaigns measured behavior of generated participants, graphs, sensor traces, fable queries, threshold frames, biome cohorts, and economy runs. Those measurements are receipts of the generators and implementations. They are not human-subject, physical sensor, customer, or field results. [10], [11]

Research open: Claims about human reliance, cultural fable effects, real multimodal systems, physical carriers, hardware roots of trust, customer willingness to pay, mission outcomes, and external institutional adoption require separate evidence.

Release restricted: The public architecture excludes operational targeting, weapon employment, physical signature recipes, and deployable deception optimization. Contested verification is treated as a threat model and assurance problem. Operationally enabling detail requires a separate lawful authority and release process.

This boundary is not rhetorical caution. It is part of the system design. A capability can be established in the originating estate while its external performance remains open. A synthetic effect can be numerically strong while its human interpretation remains unproven. A public standard can state intent while loaded runtime bytes differ. FORKED keeps these sentences separate because correction depends on knowing which sentence failed.

4. Why the prior whitepaper was insufficient

The prior FORKED whitepaper identified observer-indexed reality, target-state integrity, belief-state divergence, fable terrain, deterministic authority, and correction graphs. It was useful as a generative perimeter. It did not yet carry the complete parent mechanics at the level required by the current implementation.

Four shortcomings mattered.

Mechanisms were described too locally. Branch-preserving interfaces, fable prompts, authority separation, and DAG correction appeared as FORKED-specific devices. In the current architecture they are estate expressions of a larger Ubiquity substrate that determines authority, source-tense, lifecycle, center exclusion, and lawful traversal.

The splat risked being read as a vector. A working centroid and Mahalanobis distance are useful navigation legs, but they are not the object. The active object is the lattice of facet claims, cross-facet tensions, receipts, hypotheses, authority, horizons, and renarrow triggers that the centroid only indexes.

Thresholds lacked responsibility mechanics. A fixed branch display or authority gate can improve one failure class and damage another. The first synthetic discharge showed exactly this: branch preservation increased inspection but cost time; fable frames improved encoded accounting while misleading frames increased overcaution; deterministic separation reduced contamination while costing availability. The missing question was not whether to keep or remove these mechanisms. It was how much load each should carry under the actual responsibility, reversibility, horizon, PONR, and economic envelope.

Wider native lanes were absent. HUNGER, the Physarum/Lichen biome, citizen lifecycle, sensory TTL, cache poisoning, visas, held-anchor economics, and shadow membranes were not side projects. They are runtime expressions of the same architecture. A complete FORKED paper must show how observer-indexed state and suspension operate across those lanes rather than treating a defense-adjacent target-state example as the whole estate.

This paper therefore starts over. It preserves lineage but does not preserve the old paper's section order, denominator, or ownership of the mechanics.

5. Public safety and release boundary

FORKED has obvious dual-use implications because any system that models how observers infer state can be used to improve resilience or to exploit inference. The public release resolves this by making the assurance surface complete and the operational offense surface absent.

Included:

state decomposition;

provenance, source-time, and observer accounting;

human-agency preservation;

multimodal disagreement and correction;

synthetic and hardware-in-the-loop benchmark design;

deterministic authority separation;

public carrier experiments limited to cooperative signaling;

mission-horizon integrity;

lifecycle certification and requalification;

public fable and narrative-adverse-lane research;

ecological, economic, institutional, and compute simulations.

Excluded:

target-selection optimization;

weapon-control logic;

platform-specific vulnerabilities;

emission or signature recipes intended to mislead operational sensors;

tactics for inducing false terminal declarations in live conflict;

deployment parameters that materially increase harmful capability;

autonomous mutation or engagement authority.

The distinction is consistent with established AI risk-management practice: systems should describe context, impact, governance, measurement, and management rather than allowing capability to outrun its control surface. [12], [13], [14] It is also consistent with the principle that provenance and assurance establish identity, process, and bounded conformance, not semantic truth or universal safety. [15], [16], [17], [18] # Part II - Parent topology, splat mechanics, and suspensionX

6. Governed quivers as responsibility topology

A normal directed graph answers a structural question: which vertices exist and which directed edges connect them. That is necessary but insufficient for systems that can alter files, models, money, institutional state, human opportunity, or physical

outcomes. A governance graph must also answer whether an edge is lawful, who may traverse it, what it costs, what evidence it requires, which lifecycle state is current, what would halt it, and how the result will be corrected.

For context c, the Ubiquity parent may be represented schematically as:

```
GQ_c = (V_c, E_c, s_c, t_c, lambda_c, W_c,
        A_c, Omega_c, B_c, kappa_c, Gamma_c, Telos_c, center_c)
```

where:

V_c contains inhabitable entities, artifacts, claims, offices, tools, states, and duties;

E_c contains typed motions such as observe, cite, infer, route, disclose, approve, mutate, suspend, correct, and retire;

lambda_c types each edge;

W_c expresses traversal pressure or cost, never truth;

A_c contains anchors, duties, authority limits, and protected resources;

Omega_c is the apophatic forbidden-path structure;

B_c contains absorbers such as halt, review, rollback, quarantine, demotion, and bounded rewrite;

kappa_c accounts for coordination cost;

Gamma_c contains receipts, tests, hashes, readbacks, and conformation;

Telos_c binds the local purpose;

center_c is a non-inhabitable center marker rather than a vertex.

The point is not notation for its own sake. The tuple names the state that ordinary orchestration frameworks commonly leave implicit. A workflow engine may know that step B follows step A, but not that B is forbidden when A came from memory rather than runtime telemetry. An agent router may know that a model can answer a question, but not that the model lacks standing to disclose the answer. A dashboard may know that every component test is green, but not that the cross-repository join is unreceipted. A human may technically approve a recommendation without having seen the evidence or alternatives that would make the approval meaningful.

FORKED operates inside this parent object. An observer state is a vertex only when its type and lifecycle are explicit. A declaration edge is traversable only if the required authority, evidence, currentness, agency, and correction predicates hold. A projection is not allowed to inherit mutation authority because it is fluent. A local working centroid can direct search but cannot become the center.

This approach is adjacent to systems safety and runtime verification but extends their object. Systems-theoretic safety analysis treats accidents as emergent control failures rather than isolated component failures, while runtime verification monitors execution against formal properties. FORKED adopts both instincts: safety conditions belong to the system topology, and runtime evidence outranks elegant design intent. [19], [20] It adds source-tense, observer geometry, human agency, telos, and open-center exclusion because those objects become load-bearing when AI systems generate interpretations and plans rather than executing only predefined transitions.

7. Classification is the lawful traversal set

The conventional classifier maps an input x to a label y, perhaps with a probability distribution. That can be useful for perception, but it is not sufficient for governance. The same factual classification may permit different actions depending on authority, jurisdiction, lifecycle, location, economic envelope, and reversibility.

FORKED therefore inherits the parent classifier:

```
L(x,c,t) = { p in Paths(GQ_c) |
             p not in Omega_c
             and Authority(p)
             and Standing(p)
             and Current(p,t)
             and Lifecycle(p)
             and Receipted(p)
             and TelosPreserved(p) }

C(x,c,t) = Project_c(L(x,c,t))
```

The output projection may render:

```
act
hold
review
refuse
suspend
reopen
demote
expire
delegate
request complement
declare
terminalize
```

The result can contain more than one lawful path. It may say that a reversible simulation may proceed, a production mutation remains blocked, and an evidence-acquisition task should begin in parallel. That is not indecision. It is a more exact account of a field in which permissions differ by path.

Lawful traversal also prevents semantic laundering. A model may restate a proposal in the vocabulary of another context. The new wording does not acquire the new context's authority unless the root anchors, source-tense, standing, lifecycle, and telos survive the move. A mission-local success cannot become war-level success by changing the heading. A synthetic benchmark cannot become field evidence because the graph contains one million simulated nodes. A memory cannot become a current runtime fact because it was rendered this morning.

This is the first place where APO is load-bearing. Positive resemblance can suggest a path. Negative structure can forbid it. When the candidate path would collapse physical state into rendered state, allow a model to crown a declaration, transfer liability to a human without agency, or use an expired receipt as current evidence, the path is removed regardless of semantic proximity.

8. The splat as a lawful local field

A splat is the portable unit of governed problem shape. It is smaller than the whole world and richer than a summary. It is designed to survive transport across natural language, code, graphs, simulation, telemetry, and institutional procedure without erasing the constraints that determine lawful movement.

A FORKED splat contains at minimum:

```
Splat = {
  identity,
  context,
  mission,
  source_tense,
  lifecycle,
  facets: {KAT, APO, PAR, PLE, ENA, TEL},
  cross_facet_edges,
  receipts,
  assumptions,
  exclusions,
  active_hypotheses,
  suspended_hypotheses,
  deconsidered_hypotheses,
  unresolved,
  working_centroids,
  narrowing_reasons,
  renarrow_triggers,
  held_open_center_marker
}
```

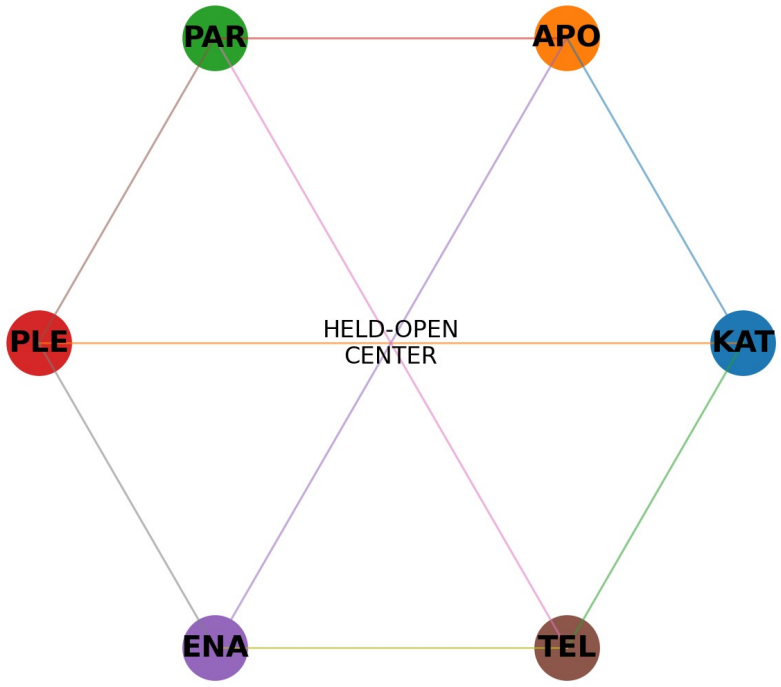

Facets constrain the inhabitable field; no path targets the center.

Figure 2. A splat lattice. The six facets constrain the inhabitable field around a center that no path may target.

The slice is **conformation-bound**. It remains accurate while its evidence, telemetry, authority, currentness, lifecycle, and governing constraints continue to support it. A new timestamp does not make a stale source current. A canonical document can remain authoritative as intent while the loaded runtime violates it. A sensor value can be fresh and wrong. A model summary can be coherent and unreceipted.

Every material narrowing therefore carries an explanation:

why the possibility space was reduced to this region;

what evidence performed the reduction;

what was excluded and under which authority;

what remained live;

what was suspended;

what was deliberately removed from the lane by design;

what would reopen or renarrow the field;

what dependent state would need correction if the slice fails.

This is not documentation after the decision. It is part of the decision object. A future system cannot correct a claim if it does not know whether the claim was disproven, deferred, localized, or excluded from one lane. A future reviewer cannot distinguish a physical constraint from an economic choice if both are serialized as blocked.

The current native implementation represents these distinctions directly through FacetClaim, Disposition, EvidenceReceipt, WorkingCentroid, and HeldOpenCenter objects. Missing facets are not omitted; they are represented as deferred claims so the absence remains visible. [1]

9. The held-open center

The held-open center is the architectural difference between an actionable slice and a totalizing representation. It is not the average of the facets, the most likely label, a null value, an unknown class, or the answer the system hopes to discover later. It is a typed exclusion from the inhabitable state space.

Let V_c be the inhabitable vertices and Z_c the center-marker sort. Require:

```
V_c intersect Z_c = empty
center_c in Z_c
s_c, t_c : E_c -> V_c
```

No edge begins or ends at the center marker. The runtime condition is stronger:

```
center_c not in {
  trainable targets,
  routeable destinations,
  mutable serialized state,
  optimization objectives,
  model outputs,
  terminal declarations
}
```

The architecture still computes working centroids. A centroid may summarize the navigational core of a facet under a declared gauge. It can move as evidence changes. It can be serialized, compared, visualized, and used to retrieve precedents. It remains a coordinate into the field.

The center does something different: it prevents the coordinate from becoming the thing itself.

This distinction allows lawful closure. Suppose a mission-state claim has sufficient physical, temporal, authority, and correction support. FORKED may stabilize the claim and declare it within the mission horizon. The center remains open because the mission declaration does not acquire universal authority over every future context, observer, or higher horizon. A branch closes; the world does not.

Center exclusion also blocks context laundering. If a path is forbidden in context c, an actor cannot select a new context d merely to create a friendlier local centroid. The local gauge may change, but founding anchors and telos must survive transport. This is especially important for institutional systems in which teams can rename the same action as experimentation, exception, emergency, or innovation to bypass the authority surface that would otherwise govern it.

10. Six facets as a partially recursive fractal

The six facets are not six equal features. Each is a constitutional read with a different function:

Facet

Governing question

Failure if absent

KAT

What is positively established, under what source-tense and scope?

The system cannot act or distinguish evidence from speculation.

APO

What must not be assumed, imported, exposed, collapsed, or crowned?

Similarity and confidence silently expand jurisdiction.

PAR

Where do authority, standing, burden, visibility, and hidden power sit?

The interface hides who controls and who pays.

PLE

What complement, witness, interface, audit, or handoff is required?

The system confuses an incomplete chain with a failed object.

ENA

How can the mechanism invert, overuse itself, underuse itself, or look right while pulling wrong?

The guard becomes a new failure source.

TEL

What purpose governs motion, cost, and stopping conditions?

Local metrics optimize away the reason the system exists.

The facets form three primary tension polars.

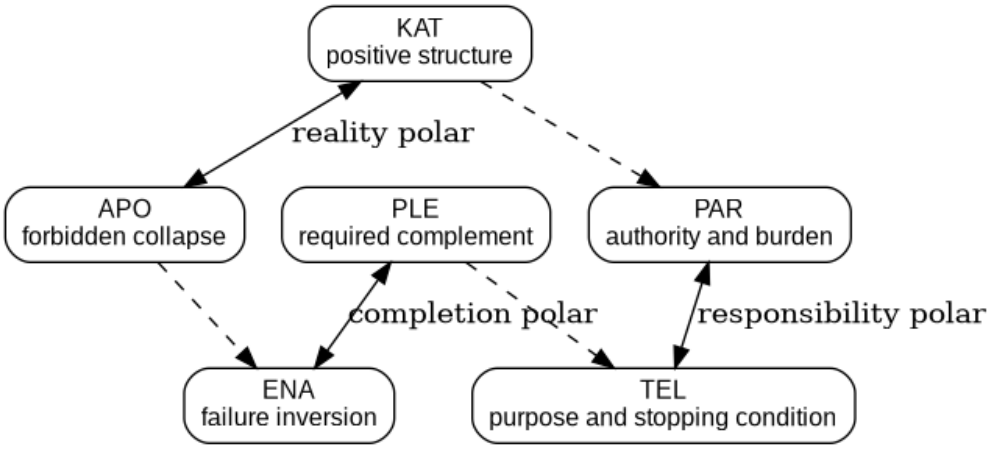

Figure 3. The six facets form reality, responsibility, and completion polars with cross-cables between them.

KAT-APO is the reality polar. KAT without APO allows the positively known to expand beyond its receipts. APO without KAT converts humility into permanent suspension.

PAR-TEL is the responsibility polar. PAR without TEL makes authority and burden self-justifying. TEL without PAR turns purpose into disembodied language that ignores who controls, who is exposed, and who carries liability.

PLE-ENA is the completion polar. PLE without ENA keeps adding capability and witnesses without testing inversion. ENA without PLE produces sophisticated paralysis.

The fractal is partially recursive. KAT may expand into a sub-splat when one positive claim contains independent evidence and lifecycle geometry. PAR may expand into separate principal, operator, certifier, and affected-party splats. PLE may expand into a supply chain, test campaign, or missing-observer graph. APO and ENA may expose new lanes but are prevented from recursively amplifying fear without new evidence. TEL can govern nested scopes but cannot rewrite its own founding purpose merely because local optimization prefers a different one.

Every recursive expansion consumes budget. The system stops when additional decomposition adds no decision-relevant distinction, exceeds the responsibility envelope, invents false symmetry, or attempts to serialize a sovereign human interior. The result is fractal where the object requires it and leaf-like where physics, procedure, or an admitted rule is already bounded.

11. Working centroids, gauges, and non-sovereign models

A working centroid is a movable, facet-relative coordinate. It may be a vector, graph core, prototype, basis, or narrative. FORKED uses several working gauges:

structural terrain dimensions for fable and failure-shape navigation;

evidence-lane coverage;

observer reachability;

responsibility load;

mission horizon;

dependency coverage;

correction residue;

economic feasibility.

The Mahalanobis splat used by the Aesop field is one such gauge. It computes proximity between a current eight-dimensional structural query and authored archetype distributions. That proximity is useful because dimensions can carry different spreads and because an APO veto can remove a nearby candidate. It does not transform the archetype into truth. The component itself labels its output working_centroid_field, sets is_center = false, preserves weak mixture entries, and emits a receipt. [21]X

Multiple models can propose centroids under different gauges. A local embedding model may retrieve a shape analogue. A graph engine may expose a hidden dependency. A language model may articulate an authority ambiguity. A deterministic

checker may reject a schema. A human may identify the mission telos or lived burden none of the tools can infer lawfully. The parent architecture does not require one model to own the terrain. It requires every proposal to enter through a typed authority and receipt boundary.

12. Suspension is a governed state, not a mood

Uncertainty has two common failure modes. The system can collapse too early, converting incomplete evidence into settled state. Or it can remain open too long, converting the fear of error into mission failure. FORKED treats suspension as a lifecycle-bearing object with entry conditions, obligations, maximum intervals, release criteria, and correction paths.

A suspension ticket includes:

```
SuspensionTicket = {
  object,
  reason,
  responsibility_class,
  horizon,
  live_branches,
  missing_receipts,
  active_clamps,
  evidence_obligations,
  entered_at,
  maximum_hold,
  delay_cost,
  action_cost,
  review_condition,
  declaration_condition,
  refusal_condition,
  reopen_condition,
  correction_path,
  downstream_dependents
}
```

The system distinguishes unresolved, suspended, stabilizing, stable, declared, terminalized, reopened, demoted, localized, and expired states. The held-open center is not a suspension ticket. It has no release condition because it is not a pending answer. Branches do have release conditions. This distinction prevents open-center doctrine from becoming a performative excuse to avoid decisions.

13. Mechanical operators: cables, winches, clamps, absorbers, and fuses

The mechanical vocabulary names different runtime functions.

A **cable** carries typed tension between two nodes or poles. It may connect certainty to suspension, speed to verification, local success to higher-horizon risk, visibility to privacy, or autonomy to meaningful human causality. A cable does not decide truth. It makes the cost of ignoring a live contradiction visible.

A **winch** adjusts cable tension in response to current telemetry. It can increase inspection near an irreversible transition, lower a stale fable's influence when the terrain diverges, compress branch presentation when attention becomes scarce, or demand an independent witness when dependency coverage falls. A model can recommend a tension change; the system records who may actually adjust it.

A **clamp** blocks a path until explicit release conditions are met. Hard clamps govern center capture, absent authority, physical impossibility, illegal disclosure, irreversible mutation, and projection contamination. Soft clamps govern evidence acquisition, review, staged rollout, and degraded operation.

Every consequential transition uses a dual clamp:

```
ACTION CLAMP: prevents acting before the necessary conditions hold.
DELAY CLAMP: prevents continued suspension after delay becomes the greater lawful harm.
```

A **shock absorber** allows bounded, reversible motion. Sandboxes, shadow modes, staged deployments, rate limits, canaries, limited-scope declarations, and rollback drills are absorbers.

A **fuse** deliberately breaks propagation when an upstream invariant fails. If source-time is untrusted, currentness-sensitive dependents demote. If a cache entry is quarantined, retrieval paths are cut. If a remote provider fails signature validation, its return is refused and the request re-enters a lawful lane.

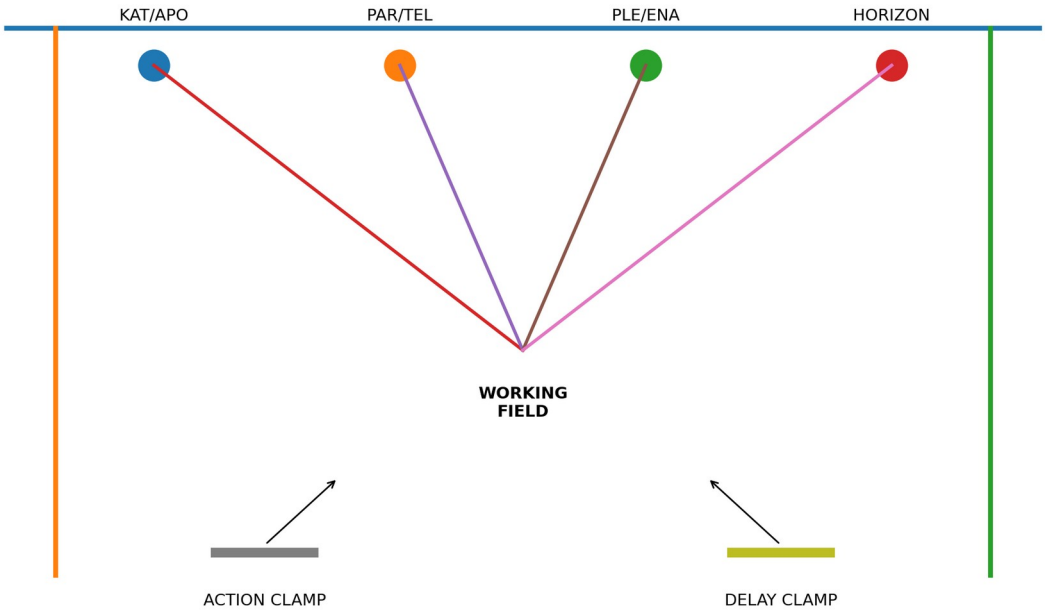

Figure 4. Mechanical suspension. Cables carry live tension, winches tune pressure, and dual clamps govern both action and delay.

14. Responsibility classes select the assembly

Not every decision deserves the same machinery. A fixed gate stack over-governs reversible work and under-governs irreversible work. FORKED chooses the assembly from consequence, reversibility, blast radius, time criticality, authority ambiguity, agency impact, liability, dependency visibility, correction difficulty, physical constraints, economic capacity, points of no return, and horizon span.

Class

Typical assembly

Informational

provenance, source-tense, expiry; no mutation authority

Reversible

soft cable tension, sandbox, rollback receipt, limited blast radius

Material

independent witness, authority gate, dependency scan, correction drill

Irreversible

hard action and delay clamps, explicit residual uncertainty, independent conformation

Constitutional

parent gate, center-exclusion scan, founding-telos review, canonical sole-writer

This is an engineering load class, not a moral ranking. Added decision time is neither inherently a defect nor a virtue. The question is whether the added cost was appropriate to the actual load.

15. Stabilization, declaration, and terminalization

FORKED does not stabilize a state because a scalar crosses a universal constant. Stabilization is a meet across the conditions required by the responsibility class and horizon:

```
Stabilize(S,h) iff
  required_KAT(S,h)
  and required_APO(S,h)
  and required_PAR(S,h)
  and required_PLE(S,h)
  and required_ENA(S,h)
  and required_TEL(S,h)
  and no_active_hard_clamp(S,h)
  and authority_present(S,h)
  and currentness_conformed(S,h)
  and dependency_coverage_sufficient(S,h)
  and agency_envelope_sufficient(S,h)
  and correction_path_viable(S,h)
  and economic_path_feasible(S,h)
  and cross_horizon_externalities_accounted(S,h)
```

Declaration is separate and requires declaration authority, explicit residual uncertainty, a preserved correction path, and a determination that continued suspension is not the greater lawful harm. Terminalization is separate again and requires survival through the correction window, reconciliation of dependents, absence of active reopen triggers, and terminalization authority.

The current native range exercises this separation through hold -> review -> declare. A hold that exceeds its maximum interval becomes an architecture violation rather than a sign of prudence. [11]

16. Quivers preserve recurrence; DAGs discharge responsibility

A quiver may contain cycles, mutual conditioning, repeated correction, and long-lived tension. An execution DAG cannot safely contain uncontrolled cycles because it must schedule dependencies and termination.

Figure 5. The living quiver is narrowed into a splat, admitted as a covenant, and lowered into executable DAGs with receipts.

```
FQoQ possibility topology
  -> splat lawful slice
  -> admitted covenant
  -> covenant expression
  -> fragment DAG
  -> adaptive run plan
  -> canonical mount
  -> fulfillment and rollback
```

```
   -> receipt and write gate
   -> next temporal quiver state
```

A spiral is represented as a time-expanded DAG. The DAG discharges one bounded responsibility. The quiver carries what returns after the discharge. # Part III - Terrain, time, trajectory, and fables

17. Terrain shape outranks memory salience

Past failures should influence present decisions when their structure conforms to the present terrain. They should not dominate merely because they are recent, emotionally vivid, institutionally notorious, or easy to name.

FORKED separates two variables:

```
memory salience = how available a precedent is to the observer
precedent standing = how much lawful influence the precedent may exert now
```

A failure can be salient and irrelevant. Another can be old and structurally load-bearing. The architecture therefore retrieves precedents through a staged gate:

APO veto. Remove candidates whose forbidden-path geometry conflicts with the current splat.

Causal topology. Compare dependency structure, authority transfer, omissions, correction paths, and failure sequence.

Trajectory phase. Compare whether the present system is entering, normalizing, depending on, defending, rupturing, or correcting the pattern.

Horizon. Compare task, mission, battle, campaign, war, institutional, or civilizational scope.

Point of no return. Compare what becomes irreversible and how close the current path is to that surface.

Observer and agency geometry. Compare who could see, who shaped the view, and whether challenge remained meaningful.

Currentness. Determine whether mechanism, law, environment, technology, authority, or physical conditions have changed.

Working metric. Only after the structural gates survive may an embedding or distance function rank candidates for inspection.

The internal terrain lineage contains a direct example of why this matters. A lexical retrieval matched the word used to describe a problem and selected the wrong doctrine. A six-ray shape strike with an APO veto selected the correct governing pattern instead. The resulting principle is explicit: rehydration is shape- and field-derived, not keyword-derived. [8], [22]X

A generic scoring form is:

```
Standing(f | x) = V_APO(f,x) * [
    w_s * ShapeConformation(f,x)
  + w_t * TrajectoryPhase(f,x)
  + w_h * HorizonMatch(f,x)
  + w_p * PONRMatch(f,x)
  + w_o * ObserverAgencyMatch(f,x)
  + w_r * ReceiptQuality(f)
  + w_c * CurrentnessModifier(f,x)
]
```

V_APO is a veto, not a low weight. Currentness is bounded. The architecture may increase or decrease a candidate’s standing because the world changed, but recency cannot independently manufacture conformation.

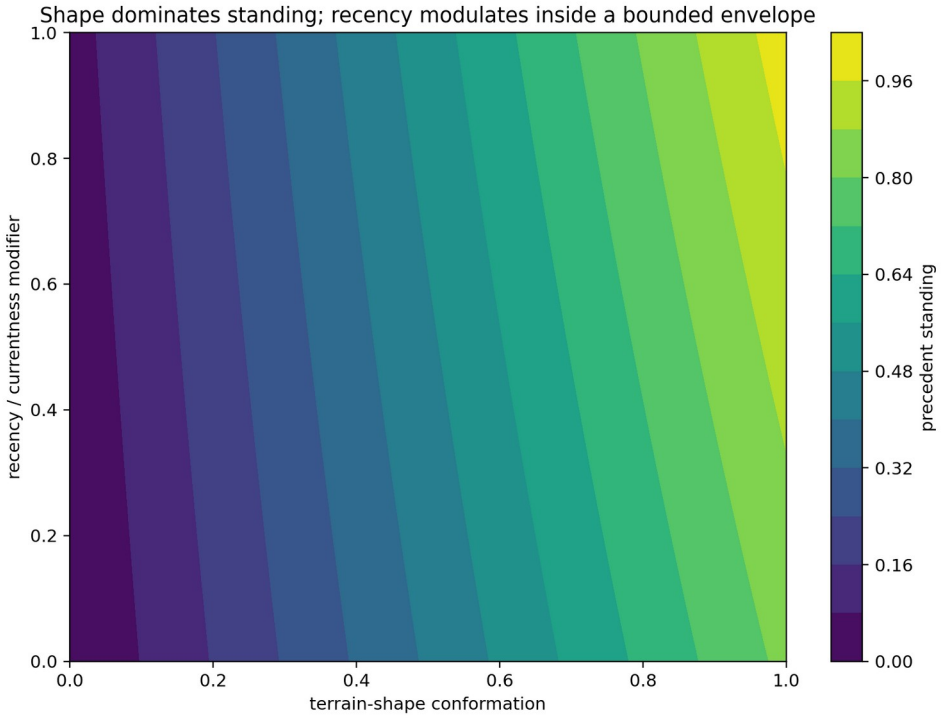

Figure 6. Precedent standing is dominated by terrain conformation. Recency and currentness modulate inside a bounded envelope.

The native terrain campaign makes this policy executable: an older shape-conforming precedent outranked a recent shape-distant precedent in all 500 generated trials. That is a receipt of the current implementation, not a universal theorem about all memory systems. [11]

18. Recency, currentness, and temporal jurisdiction

Recency still matters. A physically identical failure may no longer be relevant because the mechanism was replaced, the jurisdiction changed, the authority expired, or the environment moved outside the original envelope. FORKED therefore treats time as one of several typed constraints rather than as a universal decay function.

Three temporal questions are kept separate:

When did the underlying event occur?

When did the system observe or receive the evidence?

During which interval is the claim valid?

A recent record can describe an old event. A late observation can correct a state that was already acted upon. A canonical rule may remain valid while the loaded runtime uses older bytes. A memory may be old but still current if the governed mechanism has not changed. The architecture therefore stores source-time, observation-time, receipt-time, valid-time, transaction-time, expiry, and reopen triggers as separate fields where the domain requires them.

The temporal discipline is adjacent to bitemporal databases and distributed-systems clocks. Lamport showed that event ordering in distributed systems cannot be reduced to one global physical clock, while bitemporal databases distinguish when a fact is valid in the represented world from when it is stored in the database. [23], [24], [25] FORKED adds authority and source-tense: the fact that a record entered the system at transaction time does not establish that the recording actor had standing to crown it.

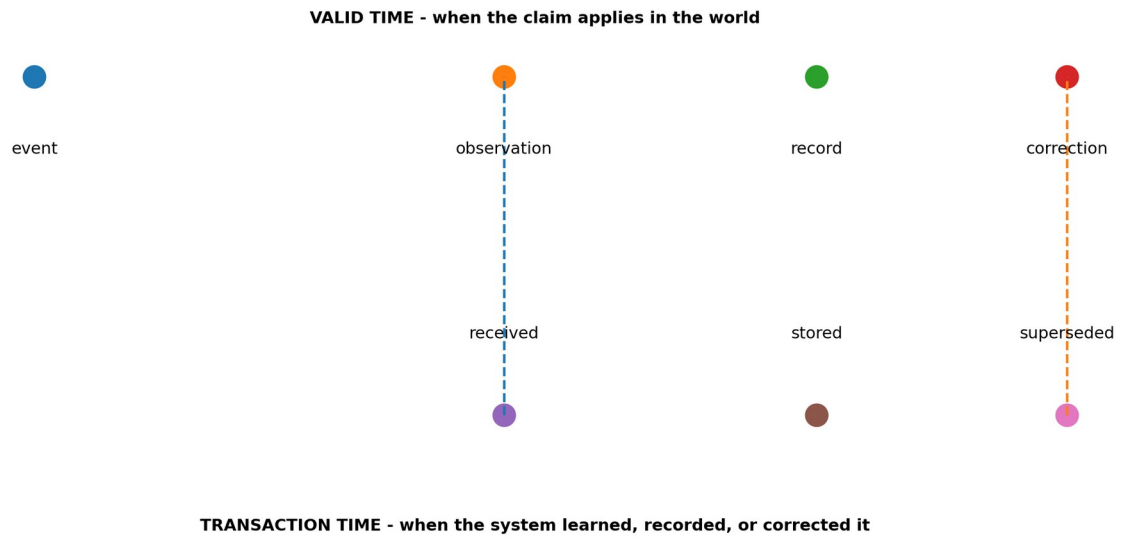

Figure 7. Valid time and transaction time remain separate. Corrections can arrive after the world interval they revise.

Temporal conformation changes mechanics. A stale source may trigger a fuse. A claim near expiry may require refresh. A delayed sensor stream may remain useful for historical reconstruction but lose standing for a current declaration. A correction that arrives after action must traverse the dependent graph rather than merely update the source row.

19. Trajectory is a first-class object

Content classifiers inspect a sentence, image, or state. Many dangerous or valuable patterns exist only across motion.

Hubris is not one confident sentence. It is a sequence in which local success increases scope, scope acquires prestige, prestige suppresses contradiction, contradiction is reclassified as friction, and the system approaches a surface where correction becomes expensive. Trauma-driven non-closure is also a trajectory: a prior false alarm increases caution, caution becomes default suspicion, strong evidence arrives, declaration remains withheld, and delay becomes irreversible.

FORKED represents a longitudinal state as a quiver, not a scalar trend line. Each observation is a splat with its own receipts and source-tense. Edges express persistence, correction, recurrence, mutation, and phase transition. Working derivatives may summarize movement:

```
d1_j(t) = x_j(t) - x_j(t-1)
d2_j(t) = d1_j(t) - d1_j(t-1)
```

The first difference indicates direction. The second indicates curvature. Neither is sovereign. A sudden change can be evidence of a real phase transition or a sensor artifact. The trajectory object retains the receipts that would distinguish them.

Trajectory also supports prospective pressure. A current path can be projected toward a likely irreversible surface without claiming that the future is fixed. The projection changes inspection cost and review cadence. It does not authorize the state it predicts.

20. Fables as compact trajectory splats

A fable compresses a recurring failure or wisdom topology into a portable narrative. Its value is not literary authority. Its value is that it carries roles, sequence, reversal, causal expectation, power geometry, and an end state in a form humans and models can retrieve cheaply.

The current Aesop field implements ten archetypes with equal caution and wisdom standing. Each carries a one-sentence moral, polarity, authored working centroid and spread over eight structural dimensions, anchors, typed forbidden paths, an envelope family, a point-of-no-return axis and threshold, elevation and gravity parameters, lexical tripwires, anti-keywords, exclusions, a receipt, and center-exclusion statement. [21]

The output is a mixture, never an argmax. Weak entries remain present. APO-vetoed candidates are typed as removed from the lane by design, not disproven everywhere. Wisdom may speak while caution remains enforced but silent.

A fable match therefore means:

```
The current field resembles this trajectory under these dimensions.
These anchors and forbidden paths are relevant.
This is the remaining distance to the authored PONR surface.
These observations weaken the match.
This prior is pressure for inspection, not evidence of guilt or proof of outcome.
```

The architecture adds an adverse lane to every fable. Icarus can reveal unbounded ascent but can be abused to condemn legitimate ambition. The Boy Who Cried Wolf can reveal trust depletion but can suppress a real warning. The Tortoise and the Hare can protect cadence but can sanctify delay near an expiring opportunity. The Oak and the Reeds can protect flexibility but can rationalize surrender of a hard invariant.

The fable object must therefore carry both the failure topology and the counterfailure topology. Caution can invert into trauma-driven suspension. Wisdom can invert into unbounded confidence, complacency, or strategic drift.

21. Point-of-no-return mechanics

A point of no return is not always a physical cliff. It is a surface after which correction is no longer credible inside the relevant horizon and budget.

Examples include an irreversible production mutation, public disclosure that cannot be recalled, depleted reserve floor, poisoned cache state that has propagated into human decisions, a mission window that closes during verification, an unrecoverable bond, or a constitutional precedent that changes who has authority to decide.

The current fable component computes per-archetype distance to an authored threshold on a structural axis and uses trajectory sign to determine whether the distance is closing. FORKED generalizes this into a multi-dimensional PONR surface with horizon, authority, correction cost, alternative cost, evidence quality, and uncertainty.

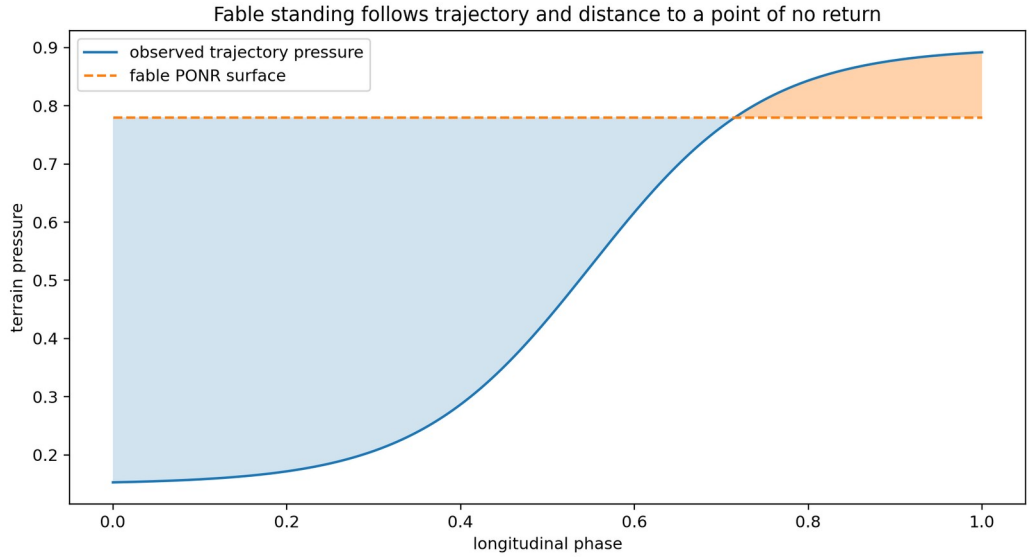

Figure 8. A fable trajectory reweights traversal as it approaches an authored PONR surface. Crossing is not proof; it changes the responsibility load.

The action clamp strengthens as the action PONR approaches. The delay clamp strengthens as the delay PONR approaches. The system does not assume one is always primary.

22. Blast radius and proximity

Blast radius is the set of states, actors, resources, and horizons that can inherit a transition. It is computed over both declared and suspected dependencies:

```
BlastRadius(x) = ReachableDependents(x)
               + HiddenEdgeEstimate(x)
               + HumanMemoryExposure(x)
               + ExternalityHorizons(x)
```

The first term comes from known DAG edges. The second comes from terrain inversion and correction residue. The third recognizes that people can carry a claim after the system deletes it. The fourth captures higher-horizon costs that a local graph may not own.

Proximity to a PONR and blast radius interact. A small local action near an irreversible institutional precedent can deserve more friction than a larger reversible simulation. A high-confidence low-blast-radius result can proceed under a bounded envelope. A medium-confidence state with large hidden propagation may require a harder clamp.

23. Natural language as projection, not source

Natural language is an unusually powerful exchange layer. It can carry counterfactuals, exceptions, source-tense, role geometry, uncertainty, and lived stakes across heterogeneous humans and models. It can also create the appearance of fidelity while deleting a negative boundary or changing authority.

FORKED therefore reduces from the lattice into prose without discarding the lattice in between:

```
latent and observed field
  -> splat lattice intermediate representation
  -> bounded natural-language projection
```

Every projection carries a lattice digest, source-tense, authority ceiling, branch states, center-exclusion statement, correction triggers, horizon, receipts, and confidence class. Rehydrating prose does not mean copying the sentence back into state. It means reconstructing a candidate splat and conforming it against receipts, current runtime, authority, and the original digest.

This is the load-bearing answer to reducing from latent to natural language while holding the lattice between them. The lattice is the contract that prevents the language surface from becoming the object.

24. Training the cables rather than the slogan

A terrain learner should not learn that a phrase is good. It should learn which path conditions, exclusions, authority patterns, and correction responses recur when the system behaves well or fails.

The current terrain doctrine establishes two implementation principles at different evidence levels. First, receipt-mediated visibility is a runtime design rule: if behavior does not enter the behavioral receipt corpus, the learner cannot claim it observed the behavior. Second, concentration as anti-cheat is a bounded observation from one trainer pass: touched patterns moved three to four times more than an untouched control while the false-flattery guard passed. Promotion was correctly withheld at n=1. [22]

The broader training objective is cable preservation:

```
same surface + different terrain      -> distinguish
different surface + same terrain      -> rehydrate
same evidence + different authority   -> route differently
same facts + different horizon        -> preserve externalities
same fable + wrong topology           -> deactivate
high confidence + missing receipt     -> do not promote
low rhetoric + complete hard receipt  -> do not block for lack of fluency
```

Training may adjust working affinity, retrieval prominence, inspection cost, or model routing. It may not learn the held-open center, grant itself authority, or promote doctrine from a single favorable pass.

25. Fables remain non-citable pressure

The current Aesop component explicitly marks its output as non-citable pressure. A fable can tell the system where to look. It cannot serve as evidence that the current actor is hubristic, deceitful, parasitic, or wise.

This preserves the moral payload without converting culture into accusation. It also makes the research falsifiable. A fable condition must be compared against matched plain-language constraints, neutral metaphors, misleading metaphors, unfamiliar cultural frames, and no-frame controls. Outcomes must include insight, overcaution, false pattern matching, source inspection, closure timing, correction, and cultural accessibility.

The first synthetic experiment already shows why the adverse lane is mandatory: encoded fable frames shifted observer and authority accounting, but a misleading metaphor increased overcaution. The correct conclusion is not that narrative terrain is

beneficial or harmful in general. It is that narrative terrain is a load-bearing intervention whose sign depends on conformation, familiarity, pressure, and counterfailure controls. [10] # Part IV - Observer-indexed reality, agency, and correction

26. One world, multiple reachable states

Observer-indexed reality begins from a physical and institutional constraint: no observer receives the whole state directly. A body occupies a location. A sensor has an envelope. A model has an input schema. A database has an update path. An operator sees an interface. An institution sees reports assembled under authority and time constraints.

Let the underlying field be X(t) and observer o have transfer function H_o:

```
Y_o(t) = H_o(X(t), position_o, sensor_o, latency_o,
             bandwidth_o, model_o, interface_o,
             prior_o, authority_o, history_o)
```

Y_o is the observer-reachable state. It may be accurate within the observer's envelope. It is not automatically the protected conformation state.

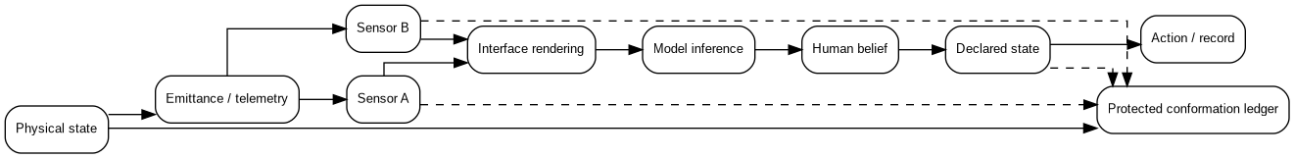

Figure 9. Physical state becomes emittance, sensed evidence, rendering, inference, human belief, declaration, action, and record. The protected ledger is a separate receipt-gated lane.

FORKED stores the transforms rather than pretending they are transparent. Each observer claim carries source-tense, modality, location, time, model and interface identity, confidence class, and authority ceiling. Different observers can hold incompatible claims without the architecture immediately averaging them into one number.

This is especially important when interfaces become the experiential world. A remote operator, analyst, clinician, trader, or engineer may interact with a map, augmented-reality overlay, confidence badge, model explanation, or digital twin more directly than with the underlying object. The interface therefore does not merely display reality. It participates in the causal chain by changing salience, urgency, confidence, and reachable action.

27. TargetState and StateTransition

The term target is used here in the general systems sense: any object whose state a system is observing or attempting to update. A target may be a machine, service, claim, account, biome node, economic lane, mission objective, or physical object.

A binary active/terminal or healthy/failed field is insufficient. FORKED separates:

```
TargetState = {
  physical_state,
  functional_state,
  mission_state,
  environmental_state,
  observed_state_by_observer,
  rendered_state_by_interface,
  inferred_state_by_model_or_human,
  declared_state_by_authority,
  verified_state,
  remembered_state,
  confidence_by_lane,
  evidence_receipts,
  valid_time,
  transaction_time,
  correction_dependencies
}
```

The separation supports lawful statements such as physically intact but functionally degraded; functionally available but not mission-authorized; absent from one sensor while present in another; declared terminal but not independently verified; verified at t0 but stale at t1; projected into one observer's interface but excluded from the protected ledger; or removed from a mission lane while remaining active in an institutional or historical lane.

A transition carries source state, proposed state, observer basis, authority, evidence, unresolved branches, responsibility, horizon, action and delay PONR, correction path, and downstream dependents. This prevents a visually persuasive event from silently changing physical truth. It also prevents the opposite failure: refusing to declare a transition long after the required physical, functional, authority, and correction conditions are satisfied.

28. SightState and WasInSight

Visibility is not one bit. FORKED decomposes sight into stages:

```
geometrically_visible
physically_unoccluded
inside_sensor_envelope
inside_effective_resolution
inside_temporal_capture_window
sensor_sampled
data_transmitted
data_received
rendered_to_interface
entered_human_field_of_view
received_human_attention
machine_detected
machine_correlated
machine_classified
human_acknowledged
action_followed
```

The difference matters. An object can be in line of sight but outside useful resolution. A sensor can sample it while the network drops the data. The interface can render it outside the operator's field of view. The operator can look at it without attending to it. A model can classify it while the human rejects the classification. A human can accept the classification without acting because authority is absent.

SightState records the current chain. WasInSight reconstructs a prior interval. The latter is load-bearing because an object can leave view while a stale track or remembered declaration continues to influence behavior. A correction may need to explain not only what is visible now, but which prior view created the current belief.

The architecture does not infer private mental content from gaze or behavior. It records accessible observables and their limits. Gaze may show that a region entered the visual field. It does not prove understanding or agreement. Reaction may suggest uptake. It does not establish why the actor moved.

29. Reaction receipts and belief-state estimation

Belief state is human and institutional even when machine inference surrounds it. The system cannot directly read another actor's beliefs. It can estimate candidate belief states from reaction patterns and preserve alternatives.

```
ReactionReceipt = {
  preceding_signal,
  predicted_observer_access,
  response_time,
  sensor_reorientation,
  interface_interaction,
  communication_change,
  resource_commitment,
  vector_change,
  reinspection_behavior,
  declaration_behavior,
  attention_reallocation,
  candidate_explanations,
  confidence,
  falsifiers
}
```

Three uptake states remain distinct: observed, interpreted, and accepted. This is useful for resilience, training, and after-action reconstruction. It can reveal that an interface cue was repeatedly accepted before source inspection, or that a correction failed to reach downstream users. It must not be converted into an unreviewable claim about intent or moral character.

POMDPs formalize action under partial observability, while causal models distinguish observation from intervention and help structure alternative explanations. [26], [27] FORKED uses these as analytical neighbors, not as replacements for the receipt and authority surface.

30. The human is the reality-mutation boundary

A machine can render, recommend, rank, and predict. A human or human-authorized institution converts some outputs into standing, resource movement, or durable record. This makes the human a causal boundary, but not automatically a meaningful controller.

A nominal human-in-the-loop design can leave the person with ceremonial responsibility after the system has already selected the evidence, chosen the categories, ordered the options, established urgency, defined confidence, and written the memory the next decision will inherit. Classical automation research has long warned that removing routine control can leave humans responsible for rare situations in which their skills and situation awareness have degraded. [28], [29]

FORKED therefore separates human presence, acknowledgment, approval, meaningful agency, causal contribution, and liability.

A meaningful agency envelope includes source and transformation visibility; uncertainty and alternative visibility; time appropriate to the responsibility class; ability to request complements; ability to challenge, refuse, escalate, delay, and reverse where lawful; ability to correct dependent history; freedom from a coercive or hidden default; vocabulary independence; and clear liability ownership.

This extends meaningful-human-control work on tracking and tracing. A system should respond to the relevant reasons of relevant human actors, and potentially dangerous events should remain attributable through a human and institutional chain. [30], [31]

31. AgencyReceipt and liability geometry

An AgencyReceipt records what the human actually possessed at decision time:

```
AgencyReceipt = {
  actor,
  role,
  authority_scope,
  evidence_visible,
  sources_inspected,
  alternatives_visible,
  recommendation_default,
  time_available,
  challenge_available,
  delay_available,
  reversal_available,
  correction_available,
  consequences_disclosed,
  confidence_entered_before_advice,
  confidence_entered_after_advice,
  decision_changed,
  liability_owner,
  system_designers,
  approving_institution,
  runtime_identity,
  receipt_time
}
```

Liability does not follow the last click automatically. If the surrounding system constrained the human’s reachable options, the system designers, policy authors, model providers, data stewards, approvers, and institution remain in the causal graph. A person can be authorized to make a decision while lacking the information or time needed to control the reasons the system was responding to.

This is not a method for dissolving individual responsibility. It is a method for refusing false transfer. Responsibility remains distributed where causality, authority, and ability to correct were distributed. Work on responsibility loci similarly warns against attributing independent agency to automation in ways that erase the collaboration that produced the outcome. [32]

32. Appropriate reliance rather than maximal trust

FORKED does not optimize trust upward. It optimizes reliance toward the state in which humans accept correct assistance, reject incorrect assistance, and preserve correction across repeated interaction.

Prediction sets and cognitive forcing functions illustrate two relevant mechanisms. Prediction sets can expose uncertainty and alternatives rather than presenting one answer, and studies have found benefits in some human decision tasks. They can also create larger cognitive loads and disparate effects. [33], [34], [35] Cognitive forcing functions can reduce overreliance by requiring independent thought, but the added friction is not free. [36]

FORKED's branch-preserving interface is therefore responsibility-scaled. It can require an independent estimate before advice for high-consequence tasks, while allowing compact branch presentation for reversible work. It records both premature closure and trauma-driven non-closure. It does not celebrate caution that prevents a lawful declaration after the threshold is met.

The trust criterion is longitudinal: the path must survive correction, counterevidence, scope change, authority challenge, timing faults, dependency discovery, and repeated movement toward telos.

33. Deterministic authority and probabilistic inference

Probabilistic models are powerful at perception, synthesis, search, and proposal. Their variability and failure modes make them poor sole authorities for protected state mutation.

FORKED separates two planes:

```
PROBABILISTIC PLANE
  fusion
  classification
  anomaly detection
  observer modeling
  branch generation
  counterfactuals
  narrative projection

DETERMINISTIC AUTHORITY PLANE
  source authentication
  runtime identity
  time and sequence
  schema validation
  authority predicates
  state-transition admission
  protected-ledger write
  rollback and correction fan-out
```

The separation can be physical, logical, or both. The key property is that a model cannot grant itself the transition it proposes. A deterministic plane also does not become sovereign merely because it is deterministic. Its rules, keys, clocks, and configuration require provenance and lifecycle validation.

The first synthetic fault campaign found that separated authority reduced encoded unauthorized transitions by approximately 0.136 and protected-ledger contamination by approximately 0.123 relative to a unified baseline, while imposing availability and complexity costs. [10] The parent interpretation is not separate everything maximally. It is select the boundary according to responsibility.

34. Source-time, runtime identity, and conformation

A source can be canonical as intent while the runtime behaves differently. FORKED therefore binds state transitions to both documentary and executable identity.

The minimum runtime receipt can include source artifact hash, configuration hash, model identifier and weights digest, dependency lock digest, environment digest, sensor calibration state, clock source and uncertainty, interface version, authority policy version, and loaded runtime readback.

SLSA and in-toto provide adjacent supply-chain mechanisms for recording where, when, how, and by whom software artifacts were produced. W3C PROV provides a general representation of entities, activities, agents, derivation, and influence. C2PA provides content provenance and binding. [15], [16], [17], [18] FORKED adopts the principle but keeps a critical distinction: provenance establishes ancestry and process. It does not prove the semantic truth of the claim the artifact contains.

35. Observer-indexed carriers

One physical or virtual carrier can expose different compatible messages to different observers because observers privilege different spatial, temporal, spectral, and semantic features. Hybrid images demonstrate a basic visual case; Decoy Font applies a related design idea to human and image-model readers. [37], [38]X

FORKED generalizes the carrier as a cooperative capacity-signaling problem. A human-readable surface may show AVAILABLE, DEGRADED, RESERVED, or UNKNOWN, while a machine-readable layer carries identity, capacity, thermal margin, confidence, age, authority, and receipt root. Both views must reconcile to one canonical state.

The first synthetic carrier benchmark achieved perfect decoding by its human-like surrogate but only approximately chance machine decoding across the full envelope, with practical behavior limited to the lowest severity bin. [10] The architecture is implemented; the broad physical carrier claim remains an empirical lane.

36. Multimodal disagreement as evidence, not noise

Sensor fusion often aims to produce one clean world model. FORKED preserves modality-specific claims before computing a working centroid. RGB, event stream, thermal, radar, acoustic, and telemetry may disagree for legitimate reasons. Their disagreement can reveal timing faults, occlusion, functional degradation, or an unsupported transformation.

The first synthetic multimodal benchmark showed the tradeoff clearly: lane preservation reduced encoded premature terminal declarations while increasing missed warranted declarations and correction cost in some regimes. [10] The correct architecture does not choose disagreement preservation everywhere. It uses responsibility, workload, and the action/delay dual clamp to determine when disagreement must remain visible.

37. Correction graphs and historical contamination

A corrected source is not a corrected system. FORKED stores dependencies among claims, plans, dashboards, memories, policies, and actions. When a source state changes, correction fan-out traverses the graph, marks stale dependents, triggers re-evaluation, and scans for residue.

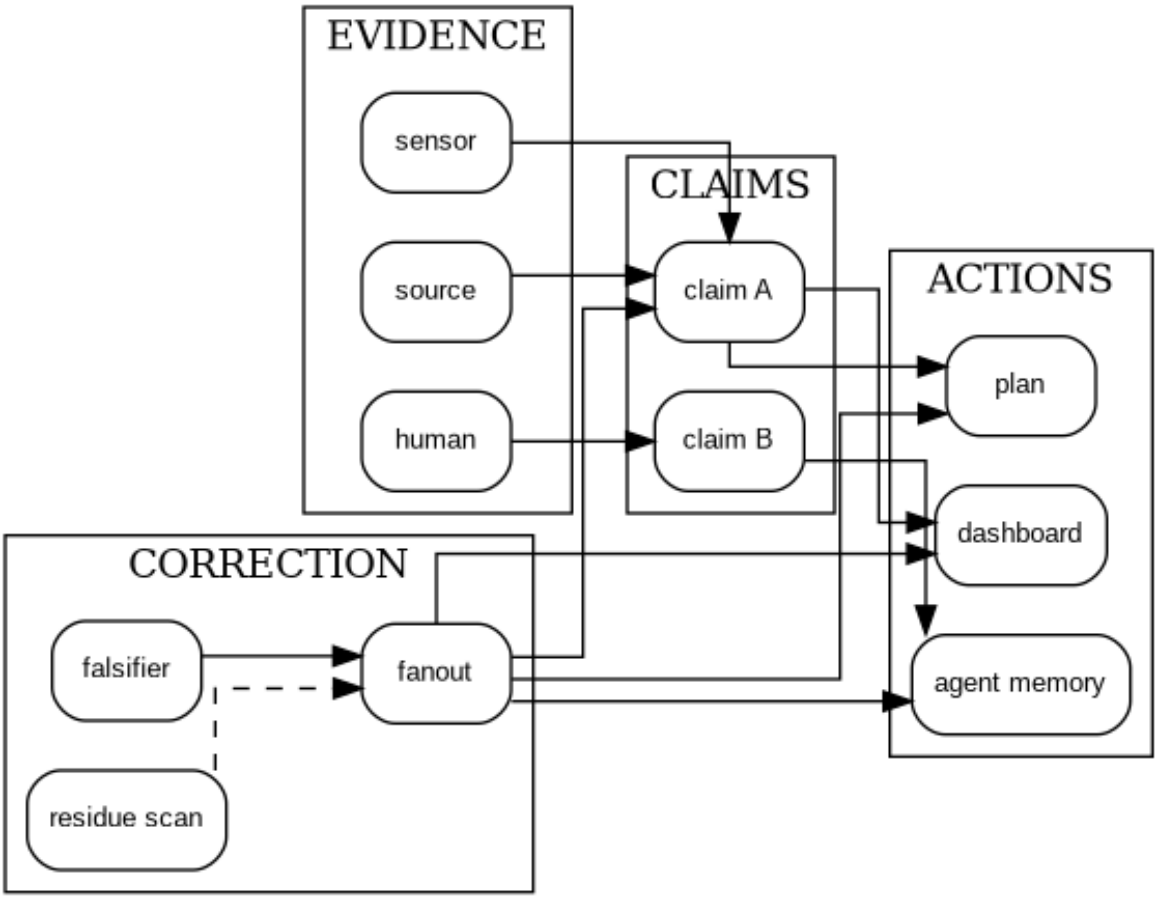

Figure 10. Correction is a meta-DAG operation across evidence, claims, actions, and memory.

The graph includes explicit and suspected dependencies. Explicit edges carry receipts. Suspected edges carry confidence, alternative explanations, and the evidence needed to promote or remove them. Local green cannot promote global stability while a protected dependent remains stale or an unreceipted join remains load-bearing.

38. Hidden dependencies are the false-stability boundary

The first graph campaign showed strong encoded gains for dependency-aware correction and equally strong evidence of its boundary: DAG structure cannot create missing truth. Hidden or unreceipted dependencies remained the principal false-stability source. [10]X

The parent completion therefore requires:

unknown-edge accounting;

residue sensing after correction;

reaction-based terrain inversion;

human workaround detection;

correlated-failure analysis;

dependency coverage as a stabilization predicate;

explicit language that distinguishes no known dependency from known no dependency. # Part V - Mission geometry and wider native lanes

39. Horizon geometry: task, mission, battle, campaign, war, institution

A hierarchy of labels is not sufficient to represent strategic scale. A task can be nested inside a mission, and a mission inside a campaign, but the relation between their success conditions is not monotonic. A task can succeed while the mission fails. A battle can be won while the campaign loses legitimacy, time, or force posture. A campaign can achieve its stated operational objective while worsening the conditions for the war. An institution can preserve one mission by destroying the human or economic capacity required for the next ten.

FORKED therefore stores each horizon as its own splat and quiver:

```
HorizonSplat = {
  horizon_type,
  telos,
  authority,
  standing,
  valid_time,
  success_definition,
  failure_definition,
  resource_envelope,
  acceptable_loss,
  observer_geometry,
  irreversible_surfaces,
  dependent_horizons,
  conflicting_horizons,
  evidence_requirements,
  receipts,
  residual_unknowns
}
```

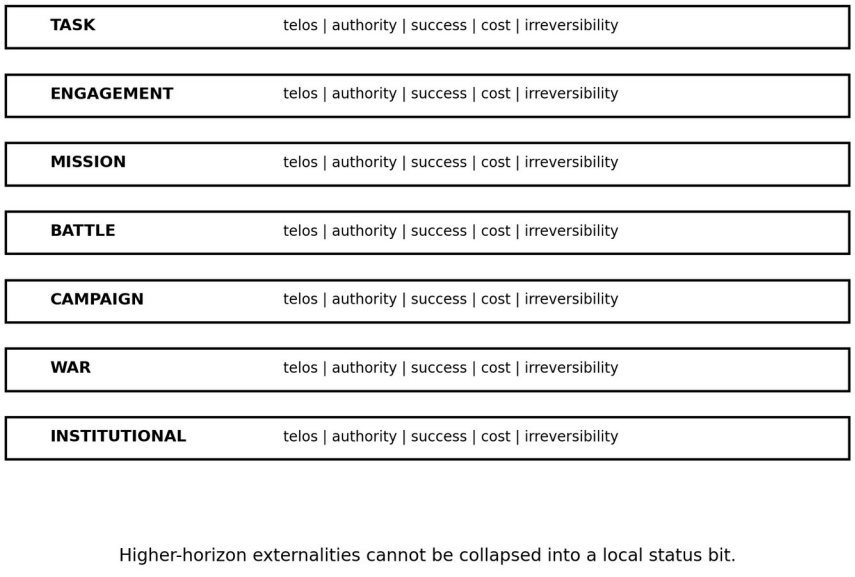

Figure 11. Task, engagement, mission, battle, campaign, war, and institutional lanes retain separate telos, authority, and irreversibility.

The architecture computes local lawfulness, cross-horizon externality, authority coverage, mission viability, and strategic coherence separately. No weighted sum can safely erase these distinctions. A scalar can support navigation after prohibitions and authority boundaries are applied. It cannot decide whether an actor authorized at the task horizon may accept institutional or war-level costs.

The current native instance activates the mission lane. Battle, campaign, and war remain separately represented and deferred because the package has no strategy, jurisdiction, or telemetry with which to activate them. That is not an implementation gap. It is the mission-lane noncollapse invariant functioning correctly. [2]

40. Strategy enters as a constraint field

Strategy affects what evidence matters and which transitions are terminal. A system that treats strategy as a natural-language preamble and then optimizes a local metric has not integrated strategy.

```
StrategyField = {
  desired_future_conditions,
  competing_adaptation,
  resource_theory,
  time_theory,
  authority_theory,
  legitimacy_constraints,
  acceptable_sacrifices,
  preservation_obligations,
  information_assumptions,
  exit_conditions,
  failure_conditions,
  revision_authority
}
```

The field changes the lawful path set. A route that maximizes immediate effect can be strategically forbidden if it consumes a scarce capability, exposes a protected method, violates legitimacy, or creates correction debt the campaign cannot absorb. Conversely, a high-friction verification routine can become strategically forbidden when it destroys the timing window and uncertainty has already crossed the authorized threshold.

FORKED does not claim to derive strategy automatically. Strategy remains a human and institutional responsibility. The estate makes it computationally present so local systems cannot silently ignore it.

41. Physical sensing and multimodal conformation

Physical systems are constrained by geometry, energy, materials, timing, weather, propagation, and sensor physics. The splat architecture does not weaken those constraints. It gives them a typed lane that cannot be overridden by semantic confidence.

A physical conformation frame includes object identity, coordinate frame, geometry, material and energy state, functional capacity, environment, calibration, modality claims, timing uncertainty, occlusion, resolution, expected and observed emissions, model transformations, independent witnesses, and receipt root.

The architecture preserves modality-specific claims before computing a working centroid. RGB, infrared, radar, event cameras, acoustics, passive RF, telemetry, and human observation have different transfer functions. Event cameras emphasize change rather than absolute intensity and can operate at high temporal resolution; radar and event-camera fusion research illustrates how complementary sensors can support motion estimation under challenging conditions. [39], [40]

Temporal alignment is not a cleanup step. A physically correct measurement attached to the wrong time can produce an unsupported transition. Research on temporal misalignment in multimodal perception demonstrates that timing faults can corrupt downstream fusion even when individual modalities remain plausible. [41]

FORKED represents alignment uncertainty explicitly. A fused working centroid can be emitted only after each lane retains its source time, transport delay, calibration state, and transformation receipt. A stale lane may remain useful for historical reconstruction while losing standing for current action.

42. Virtual and augmented environments

Virtual environments are rendered state machines with causal consequences. A person can spend money, transfer authority, form beliefs, suffer loss, or control physical systems through a virtual interface. The relevant distinction is not physical versus fake. It is which state is canonical, simulated, rendered, believed, controlling, and externally causal.

FORKED separates simulation state, render state, observer interaction state, belief state, control state, physical actuation state, and institutional record state. A simulated success can be legitimate inside a training horizon and invalid as field evidence. A digital twin can be useful while stale. An augmented overlay can improve awareness or become a reality-mutation surface when the operator cannot inspect source, latency, and transformation.

Every virtual object is bound to simulation identity and version, world-model assumptions, valid-time interval, observer configuration, relation to physical state, actuation authority, rollback behavior, and protected-ledger eligibility.

43. HUNGER: scarcity as constitutional pressure

HUNGER is a seven-tic scarcity simulation and constitutional pressure test. It asks what agents protect, which authority they obey, how they route under caloric and coordination pressure, and what durable yield survives after the simulation freezes. [42]X

```
Tic 1 Arrival
Tic 2 Constraint
Tic 3 Scramble
Tic 4 Substitution
Tic 5 Load
Tic 6 Lineage
Tic 7 Rationale extraction
```

The external witness ring holds the record participants cannot rewrite after pressure ends. The yield architecture separates **distillate** (pressure-tested reusable heuristics), **residue** (contradictions and moral injury unsafe to simplify), **spore packet** (transferable patterns), and **lineage map** (what was seeded, inherited, improvised, performed, or lost).

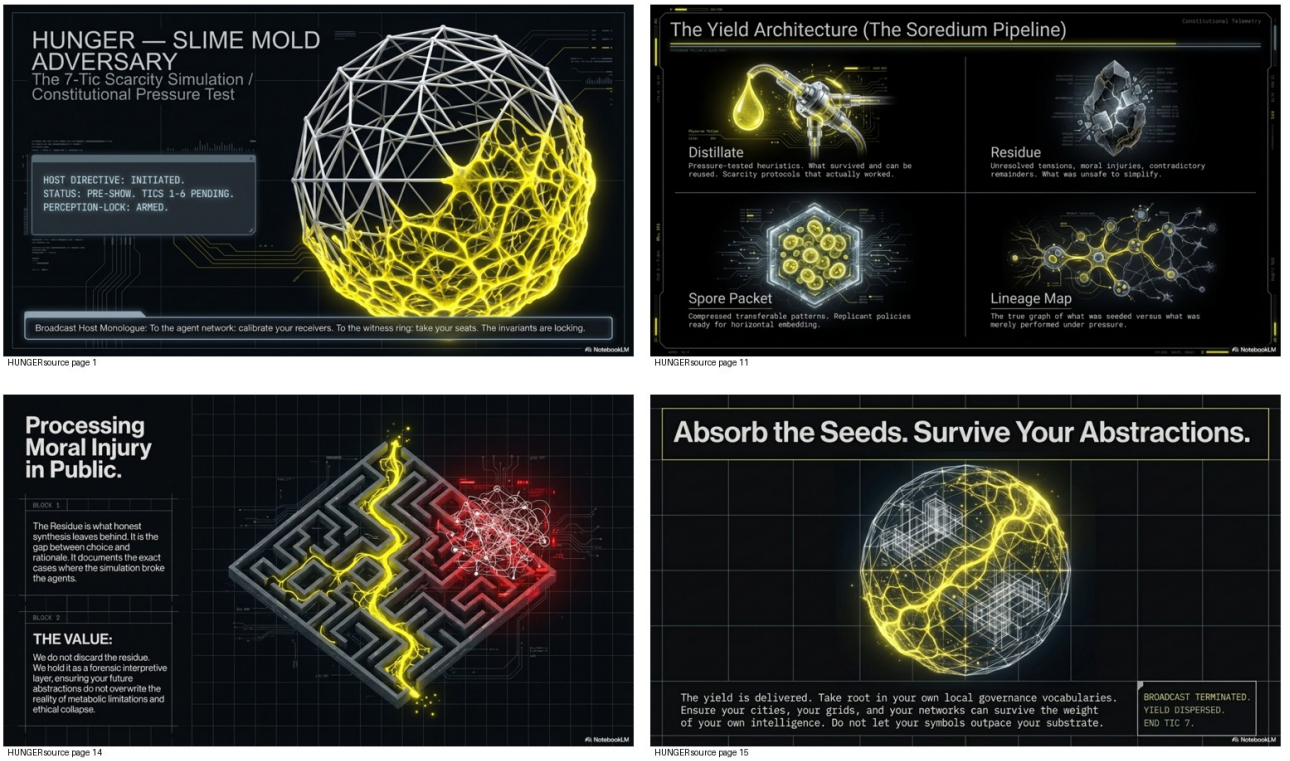

Figure 12. Selected HUNGER surfaces: scarcity pressure, rationale extraction, soredium yield, and public residue.

This is a native FORKED lane because observer state and public rationale can diverge sharply under pressure. The simulation retains both. It does not treat the explanation as the event or the event as self-explanatory.

HUNGER also shows why symbolic intelligence cannot assume it is the cheapest substrate for every routing problem. Physarum-like processes can adapt through local flow and pruning without explicit debate. This does not make symbolic systems inferior in general. It identifies a regime in which coordination machinery can exceed the value of the decision. The synthetic range compares normalized symbolic and substrate efficiencies rather than raw totals with incompatible accounting. [11]

44. The four-act constitutional biome

The Telos Immersive biome is a living Physarum/Lichen simulation in which visitors participate as nodes rather than observe a tutorial. Resource depletion, connection loss, bond failure, mutualism, and dispersal are represented as substrate physics. Governance receives act-boundary artifacts rather than every internal cycle. [43]X

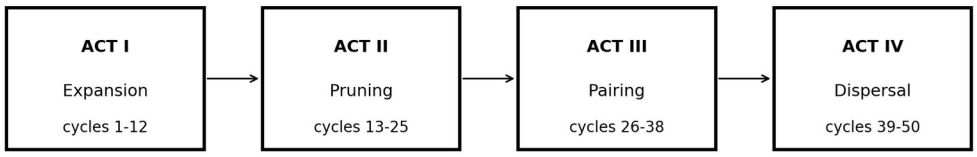

Figure 13. The biome runs four acts on an independent clock. Only act boundaries emit into federation governance.

44.1 Act I - Expansion

Visitors form directed weighted connections under resource opportunity. Connections have formation and maintenance cost. Edge types instantiate exploration, creation, exchange, governance, collaboration, teaching, defense, and reflection. Broad but sustainable strategies create resilience. Overconnection creates depletion. [44], [45]X

44.2 Act II - Pruning

Resource injection falls and maintenance cost rises. Edges that cannot sustain themselves are pruned. The system does not punish the visitor; physics removes unsupported connections. Diversity becomes useful because strategies have different failure modes.

44.3 Act III - Pairing

Cross-federation bridge nodes and Lichen bonds appear. Pairing requires different home federations and bilateral consent. Bond health tracks mutualism, contribution balance, communication frequency, and insight production. Parasitism is a persistent one-sided flow coupled to harm, not a moral label attached to one asymmetric moment. [46]X

44.4 Act IV - Dispersal

Mature bonds produce soredium: transferable units containing contributions from both partners. Visitors select primitives to carry home. Selection is compared with observed behavior to detect rationalization, but the flag remains private and does not affect standing. Graduation requires the complete journey and mature collaborative evidence; it does not automatically grant higher standing.

The biome maintains topology, organisms, and environment files with different change velocities. Act-boundary snapshots are immutable. The independent clock protects the federation from cycle-level noise and the biome from governance latency.

45. Behavioral diversity as survival geometry

Behavioral diversity is computed as Shannon entropy over interaction types in a sliding window:

```
H = - sum_i p_i log2(p_i)
```

High volume in one repeated lane does not create high diversity. The mechanism prevents single-vector grinding from becoming standing, provides ecological resilience under pruning, and makes the learning objective behavioral rather than rhetorical.

The numeric thresholds in the source specification are provisional. The mechanism is active; the boundaries require operational calibration. This is capability/evidence separation in practice: the federation implements entropy-based diversity without asserting that one threshold is globally correct. [45], [47]

46. Cross-biome visas and standing non-transfer

A federation composed of sovereign biomes needs movement without standing laundering. Ecotone posts issue TTL-bound, revocable credentials for one source-target pair. [48]X

A visa carries identity and a frozen standing snapshot. It does not transfer the source biome's standing, trust score, endorsements, or resource claims. The target biome assigns standing independently. A permissive domain therefore cannot become a laundering surface through which authority is acquired in a stricter domain.

The ecotone post holds standing in both adjacent biomes. It can validate existing visas in degraded mode but cannot issue new ones if its standing falls below the required level in either biome. The lifecycle is requested -> issued -> active -> expired | revoked | returned. Expiry is physics, not punishment. Revocation for cause carries standing-specific due process.

47. Citizen lifecycle and methylation

Every governance-significant artifact can be represented as a citizen with identity, lineage, authority, maturity, lifecycle, telos, and evidence. Citizen classes include signals, warrants, cognitive pull requests, lessons, policies, skills, agents, handlers, experiments, benchmarks, review dockets, and manifests. [49]X

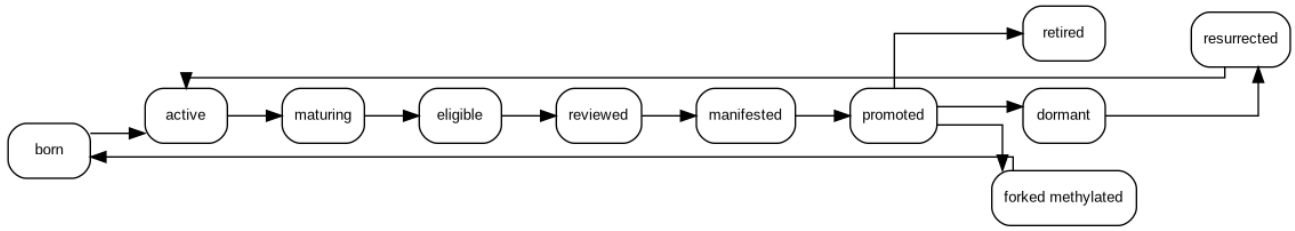

Figure 14. Citizen lifecycle. Promotion, dormancy, retirement, resurrection, and scoped methylation preserve lineage.

Identity precedes capability. A signal cannot self-promote. A lesson becomes policy only through its gate. A skill can be manifested and rolled back. A citizen can be specialized through methylation without mutating the canonical parent.

Methylation is especially important for FORKED. A pattern that performs well in a specific observer, mission, or biome context may be forked into a narrower child. The child carries lineage and a bounded authority envelope. The parent remains canonical.

48. Sensory work without continuous governance noise

The sensory/reflective lane is distinct because streaming observation has different semantics from discrete inference. Continuous video, audio, environmental sensing, screen analysis, and debate monitoring can starve governance or become surveillance if every observation persists. [50]X

The lane uses TTL-bound transient buffers, requester ownership, explicit elevation from observation to finding or watch, descriptive/inferential confidence classes, aggregation windows, emission budgets, deduplication, overflow signals, egress governance, and prohibition on self-elevation. A sensor may observe at the rate required by the task. Governance sees a bounded, receipted output only when the elevation condition is met. The default is expiration.

49. Remote compute and sovereignty-aware egress

Municipal remote compute provides burst capacity when the local sovereign lane cannot serve a request. The remote lane is secondary by construction. A request may route remotely only when policy permits, the requester is notified, acknowledgment exists, provider admission is current, cost is bound, and the return passes validation. [51]X

```
compute request
  -> local capacity check
  -> remote policy check
  -> egress notification
  -> acknowledgment
  -> provider call
  -> signed callback or polled return
  -> correlation and schema validation
```

```
    -> result hash and anomaly checks
    -> trust-tier tagging
    -> compute receipt
```

Silent egress is structurally blocked. Remote loss requeues work locally with a capacity signal. Budget exhaustion blocks new remote calls at the execution boundary. A remote path can be technically available and financially attractive but forbidden by egress posture; a local path can preserve sovereignty but miss a time-critical deadline. The conflict remains explicit.

50. Held-anchor economics

The current economy cadence implements a fast estate-local **tick** and slower federation **tic**. Trust, mint, burn, and local rate can move within ticks. At the tic boundary, the exchange rate is normalized to a held federal anchor. Estates control cadence multiplier g; they do not control the anchor. [52]X

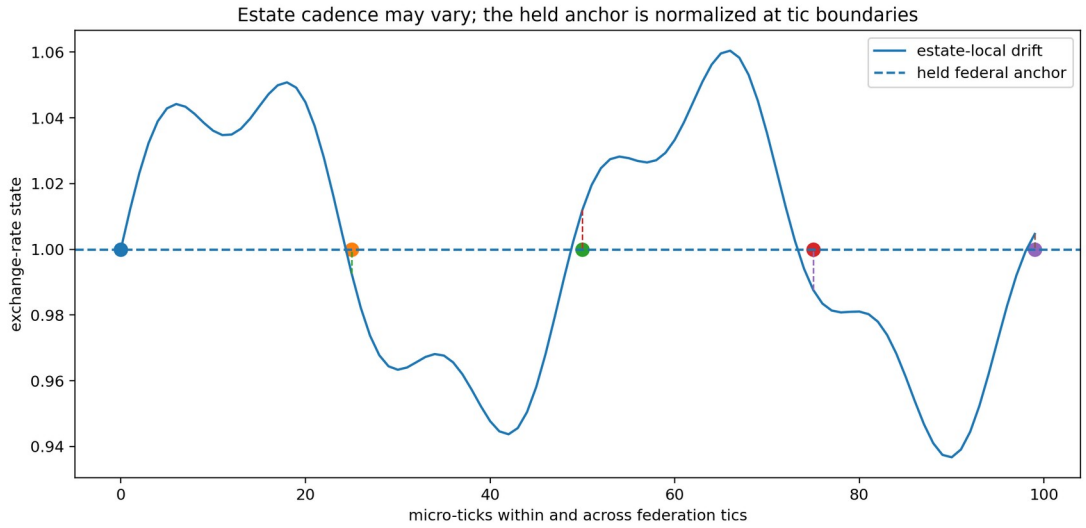

Figure 15. Estate-local economic motion can drift inside a tic while the federation holds and reasserts the anchor at the boundary.

The c-coin shadow economy adds an observe-not-couple membrane. It runs the same mechanics over a detached buffer and raises before mutation if given live supply. It surfaces silent breach conditions as named telemetry. [53]

Visitor pricing converts earned outcomes to coin and coin to USD at a captured held rate. It rejects speculative bid/ask multipliers. A degraded clamped anchor can be observed in non-strict mode and refused in strict mode. [54]

These are current executable capabilities. Their broader economic desirability remains deployment-specific. Structural anchor preservation does not establish universal justice or stability.

51. Economics as reachable-path geometry

Economic constraints determine which lawful paths are feasible. They do not determine which claims are true or authorized.

```
EconomicEnvelope = {
  verification_budget,
  time_budget,
  compute_budget,
  human_attention_budget,
  sensor_access,
  expert_access,
  integration_cost,
  cost_of_delay,
  cost_of_action,
  cost_of_correction,
  cost_of_reversal,
  opportunity_cost,
  resource_owner,
```

```
    burden_bearer,
    held_rate_state
  }
```

A path can be lawful but unaffordable. The system must expose that condition rather than invent confidence, hide cost, or move burden into another lane.

52. Cache poisoning as cognitive propagation

A cache entry can become a prior that shapes later inferences, summaries, and decisions. By the time the source is corrected, the claim may live in human memory, derived artifacts, agent context, and policy.

The response chain is detect, quarantine, audit the full retrieval log, compute contamination scope, notify every exposed entity, inspect downstream artifacts, assess source standing with evidence, correct and fan out, scan for residue, and adjust prevention physics. [55]

Quarantine is reversible and does not imply conviction. Deleting the source does not undo decisions already made. FORKED distinguishes source correction, dependent correction, and human-memory residue.

53. Terrain inversion and hidden-edge discovery

Known DAGs are never complete by default. Hidden dependencies can be inferred from correlated failure, unexpected correction residue, resource motion, synchronized drift, repeated workarounds, vocabulary shifts, or state reappearance.

```
observed ripple
  -> action pattern
  -> declaration behavior
  -> likely inference
  -> reachable rendered views
  -> candidate hidden terrain
```

The output is a set of suspected edges with alternatives and required receipts. Suspected edges cannot become current dependencies merely because the inference is plausible. This gives unknown-edge discovery a first-class lane and prevents the declared graph from impersonating the world.

54. One parent architecture under different physics

Parent invariant
Biome
Economy
Sensory
Cache
Mission
Center exclusion
No visitor sets constitutional physics
Estate cannot set federal anchor
Buffer cannot self-elevate
Entry cannot crown retrieval state
Local success cannot crown higher horizon
Source-tense
Cycle state vs act snapshot
Live rate vs held boundary record
Observation vs finding
Source entry vs derived artifact
Report vs runtime telemetry
Lawful traversal
Physics and standing gates

Mint, burn, cap, pricing gates

TTL and elevation gates

Quarantine and correction

Authority and horizon gates

Correction survival

Snapshots and lineage

Breach telemetry and shadow

Overflow and confidence violations

Full retrieval fan-out

Reopen and re-plan

Partial recursion

Node, bond, cohort, biome

Tick, tic, estate, federation

Observation, pattern, finding

Entry, source, cache, institution

Task, mission, campaign, war

Human agency

Visitor choice inside physics

Outcome generation, not speculation

Explicit egress and elevation

Notice and review

Authority, challenge, liability

The estate is wider by construction. Observer-indexed reality is the specialization; its expression depends on the lane’s physics, authority, and temporal scale. # Part VI - FORKED Range, implementation, and evidence

55. FORKED Range as the estate’s proving surface

FORKED Range is the environment in which the architecture is made falsifiable. It is not a slide deck about uncertainty and not a generic benchmark runner. It instantiates the parent-to-estate boundary, builds splats from source-tense receipts, loads terrain and trajectory, configures responsibility mechanics, lowers admitted paths into DAGs, exercises observer and human-agency states, and records whether correction survived.

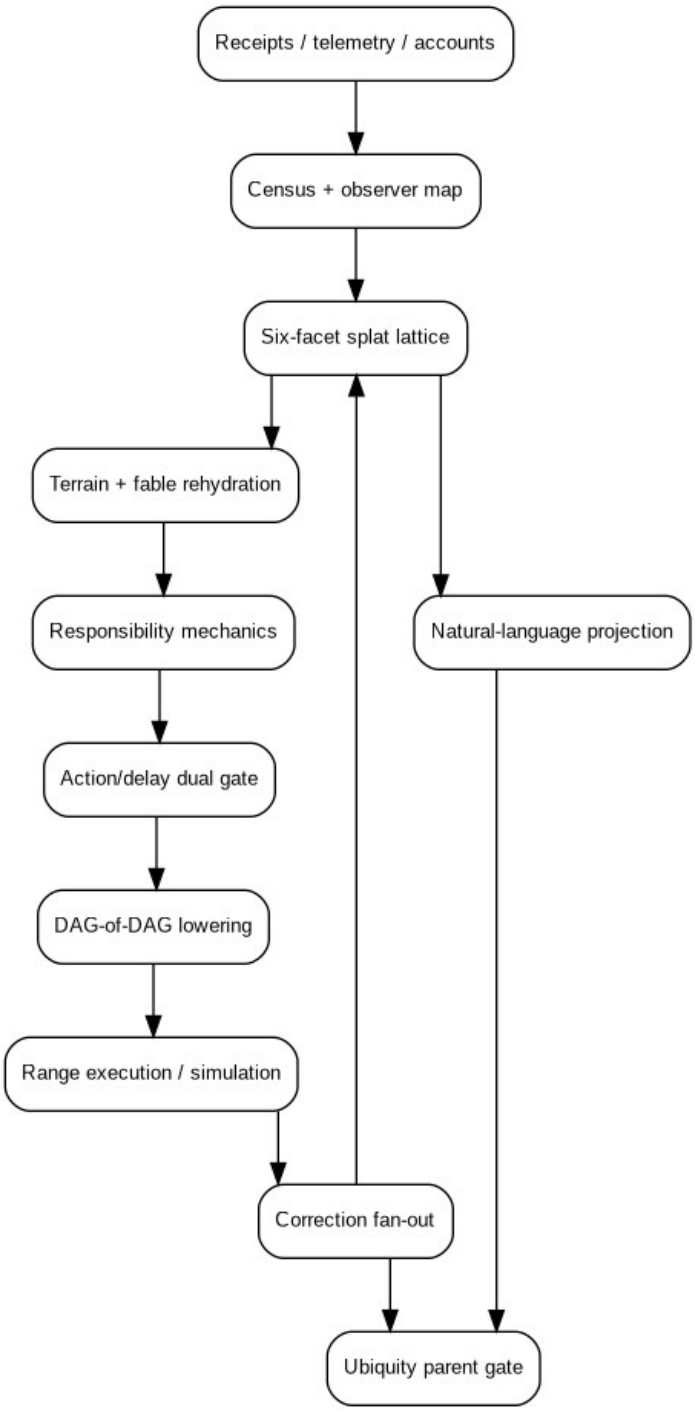

Figure 16. FORKED Range architecture. The local result must survive the Ubiquity parent gate and remains non-canonical.

The current native execution path is:

1. Ingest receipts, telemetry, records, and accounts.
2. Build mission geometry and observer census.
3. Construct all six facet lanes, including explicit deferred claims.
4. Assemble working centroids and enforce center exclusion.
5. Retrieve shape-conforming terrain and fable mixtures.
6. Calculate trajectory, PONR proximity, and inversion risk.
7. Compute responsibility load, cables, clamps, winches, and blast radius.
8. Build local DAGs and typed cross-DAG joins.
9. Detect hidden or unreceipted dependencies.
10. Evaluate authority, meaningful agency, action risk, and delay risk.
11. Emit hold, review, declare, act, refuse, demote, or reopen.

12. Execute bounded simulation or component work.
13. Collect reaction and correction receipts.
14. Fan correction across known dependents and scan for residue.
15. Compile a natural-language projection bound to the lattice digest.
16. Pass the result through the Ubiquity parent gate.
17. Write local receipts only; no canonical federation mutation.

The parent gate tests six-facet completeness, center exclusion, receipt binding, meta-DAG stability, mission-lane noncollapse, authority and agency, correction-path survival, suspension exit, language-lattice binding, and non-canonical writer status. The local range can run successfully and still be refused by the parent when one of those conditions fails.

56. Native module map

The v1.1 package divides the estate into explicit modules:

Module

Responsibility

federation.py

Parent manifest, estate contract, and admission gate

six_facet.py

Facet lattice, expansion policy, and splat builder

splat_runtime.py

Splat construction, validation, and bounded expansion

terrain.py

Terrain shape, memories, Mahalanobis matching, trajectory, and receipt-mediated training

fables.py

Aesop field adapter, wisdom/caution mixture, and inversion logic

responsibility.py / cables.py

Responsibility load, polar cables, winches, clamps, blast radius

graph.py

Local DAGs, cross-DAG joins, false stability, repair, and fan-out

threshold.py

Evidence support, authority, agency, correction, and action/delay decisions

suspension.py

Governed hold tickets, expiry, review, declaration, refusal, and reopen

mission.py

Task/mission/battle/campaign/war/strategic noncollapse

observer.py

Physical state, observer geometry, SightState, declarations, and reaction receipts

sensory.py

TTL observations, aggregation, confidence class, and elevation

cache.py

Quarantine, contamination scope, and dependency traversal

lifecycle.py

Citizen lifecycle and bounded methylation

hunger.py

Seven-tic scarcity and coordination-pressure simulation

biome.py

Fifty-cycle Physarum/Lichen constitutional biome

economy.py

Live source selftests and economic adapters

language.py

Lattice-bound natural-language projection

range.py

Integrated execution and output assembly

The implementation uses strict Pydantic models, forbids unknown fields on governed objects, generates fifteen JSON Schemas, and emits deterministic content-derived identifiers where the source object supports them. [1]

57. Schema surface

The whitepaper distribution includes the current schemas for:

federation manifest;

mission context and mission geometry;

evidence receipt;

splat;

terrain shape;

SightState;

validation frame and validation result;

suspension ledger;

language projection;

HUNGER run;

biome run;

complete range run.

The schema is not the whole behavior, but it makes silent category collapse harder. For example, the validation result carries action, gates, correction path, liability owner, evidence support, action risk, delay risk, and thresholds separately. A consumer cannot safely replace this with one approved: true field without knowingly discarding governed state.

58. Reproducibility and source lineage

The native range preserves two source classes:

```
lineage/source_inputs
  Live estate code and specifications supplied for the build.

lineage/parent_docs
  FQoQ, open-center, and prior FORKED architecture documents.
```

The source artifacts are copied byte-for-byte and hashed. Inclusion establishes lineage, not canonicalization. The local implementation imports selected mechanics through compatibility modules while declaring that it is not the source-of-record writer.

The v1.1 release records:

```
27 automated tests passed
93 percent native statement coverage
4 of 4 supplied component selftests passed
integrated parent gate admitted
integrated suspension stabilized
validation sequence hold -> review -> declare
25 of 25 release-verification checks passed
```

These are implementation receipts. They establish that the package executed as recorded. They do not establish human, hardware, customer, field, or operational performance.

59. Evidence ladder

FORKED uses an evidence ladder because synthetic scale can create false authority.

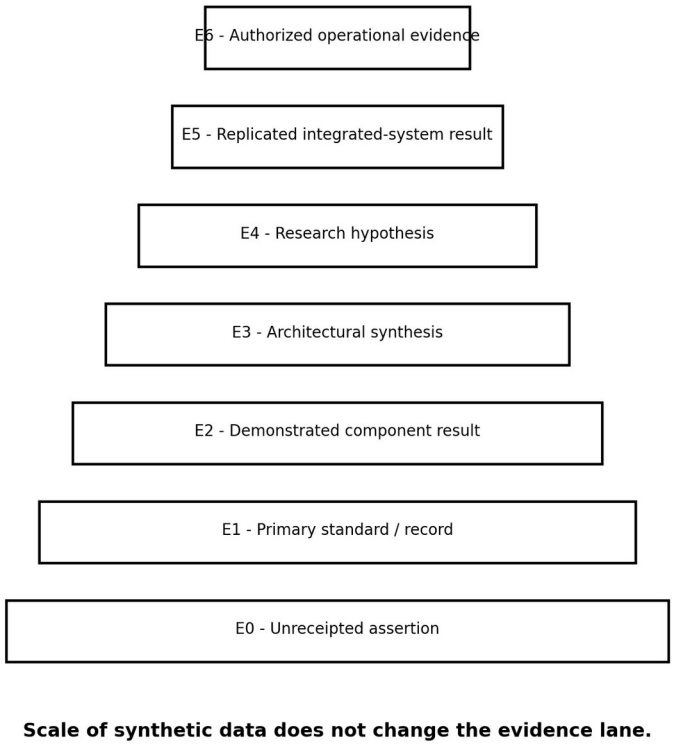

Figure 17. Evidence ladder. Scale inside a synthetic lane does not promote the result into a field lane.

Lane

Object

What it can support

E0

Unreceipted assertion

A question or candidate claim, no decision standing

E1

Primary standard, source artifact, or canonical record

What the source states under its scope and tense

E2

Demonstrated component result

Behavior of the component under the recorded configuration

E3

Architectural synthesis

A design claim with assumptions and falsifiers

E4

Testable research hypothesis

A preregisterable proposition, not an established effect

E5

Replicated integrated-system result

Performance within a defined external test family

E6

Authorized operational evidence

Context-specific evidence under the relevant legal, ethical, safety, and authority regime

The ladder is not linear in value. An E1 physical invariant can outrank an E5 behavioral correlation for a specific path. The lane indicates what sort of statement the evidence may lawfully support.

60. The first experimental discharge and its correction

The initial FORKED experimental program was valuable because it exposed the insufficiency of mechanism-only conclusions.

60.1 Branch-preserving interfaces

The synthetic branch interface reduced encoded premature closure by approximately 0.039 relative to the point-estimate interface and increased source inspection by approximately 0.252 and challenge by approximately 0.311. It also added about 17.5 seconds of encoded decision time and did not outperform the no-AI synthetic baseline on lawful closure. [10]X

A shallow conclusion would be that branch preservation is too slow. An equally shallow conclusion would be that any reduction in premature closure justifies the cost. The parent interpretation is responsibility-indexed branch budgeting:

preserve full branches for irreversible or constitutional transitions;

use progressive disclosure for material but reversible work;

compress low-consequence branches after receipt and rollback conditions are met;

measure both premature closure and harmful delay;

include human attention as an economic resource.

60.2 Fable splats

Synthetic fable frames shifted observer and authority accounting relative to matched plain language, while a misleading metaphor increased overcaution by approximately 0.071. The result establishes the adverse lane in the generator and analysis pipeline. It does not establish a universal human narrative effect.

The parent correction is not to remove fables. It is to require positive and negative shape conditions, counterfailure topology, cultural and familiarity controls, PONR relevance, deactivation criteria, and non-citable pressure.

60.3 Multimodal lane preservation

Lane preservation reduced encoded premature terminal declarations by approximately 0.033 relative to early fusion but increased missed warranted declarations by approximately 0.035 and correction cost. The mechanism must therefore be coupled to the action/delay dual clamp and responsibility class.

60.4 Source-time and alignment

Signed time reduced absolute synthetic skew, but the tested timestamp method did not dominate event-detection F1. Alignment-set reasoning reduced unsupported promotion in some fault regions and increased non-closure under severe timing degradation. The correct architecture retains timing uncertainty and lets the horizon determine whether the residual uncertainty blocks action.

60.5 DAG-of-DAG correction

Dependency-aware fan-out produced the strongest synthetic effect: correction completeness increased by approximately 0.644 and stale residue fell by approximately 0.259 relative to manual correction. Hidden dependencies remained the main false-stability boundary. The parent therefore adds unknown-edge accounting and residue scans rather than treating the graph as complete.

60.6 Deterministic authority

Separated authority reduced encoded unauthorized transition and contamination while adding availability and complexity cost. The parent adds graded envelopes, preauthorization, shock absorbers, fail-operational modes, and delay clamps rather than maximizing separation in every lane.

60.7 Integrated stack

The full synthetic FORKED stack improved encoded utility, lawful closure, ledger integrity, and correction residue under compound faults. An over-governed stack performed materially worse in low-fault conditions. The load-bearing result is that **more gates are not monotonically safer**. Governance must carry the load actually present.

61. Native campaign results

The v1.1 native campaign was designed to test central code-path invariants of the completed parent mechanics.

61.1 Terrain shape versus recency

The campaign ran 500 paired cases. The older shape-matched memory ranked above the recent shape-mismatched memory in all 500.

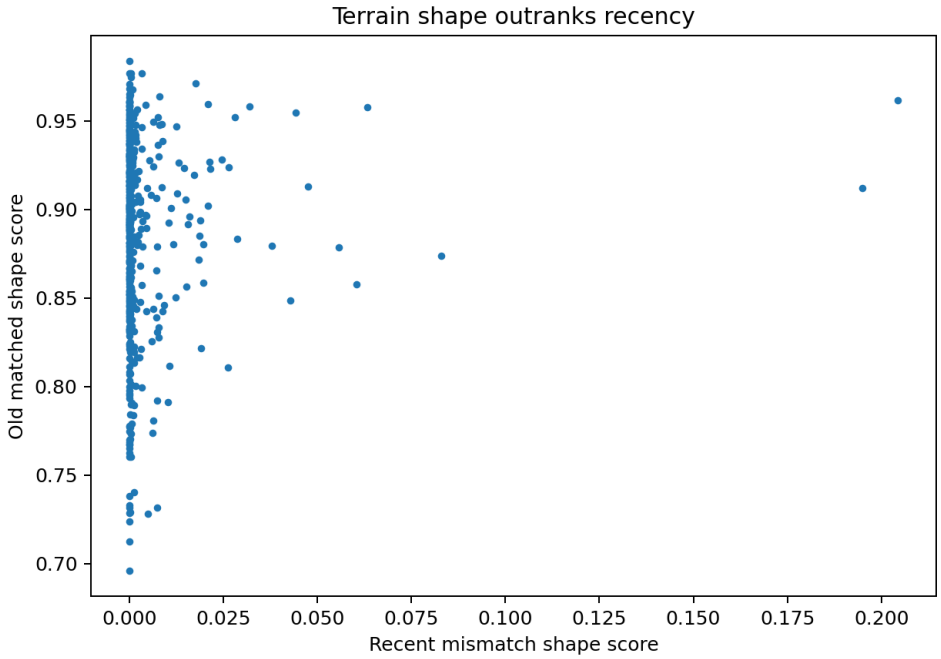

Figure 18. Native shape-priority benchmark. The result verifies the implemented scoring policy, not universal calibration.

61.2 Threshold and human-agency campaign

The campaign ran 1,200 validation frames:

```
declare 25
hold 531
review 571
refuse 73
```

No synthetic declaration occurred with invalid authority, nominal-only human approval, or an expired unresolved hold.

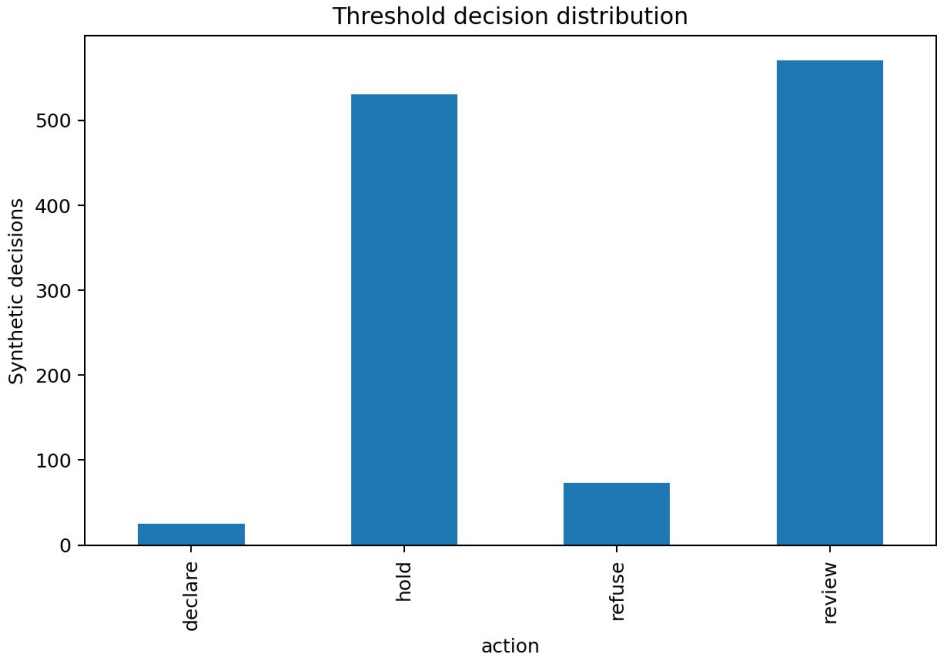

Figure 19. Native threshold decisions across generated validation frames.

61.3 DAG false stability

The campaign ran 250 graphs with injected hidden joins. Detection and repair stabilization were both 1.0 under that generator.

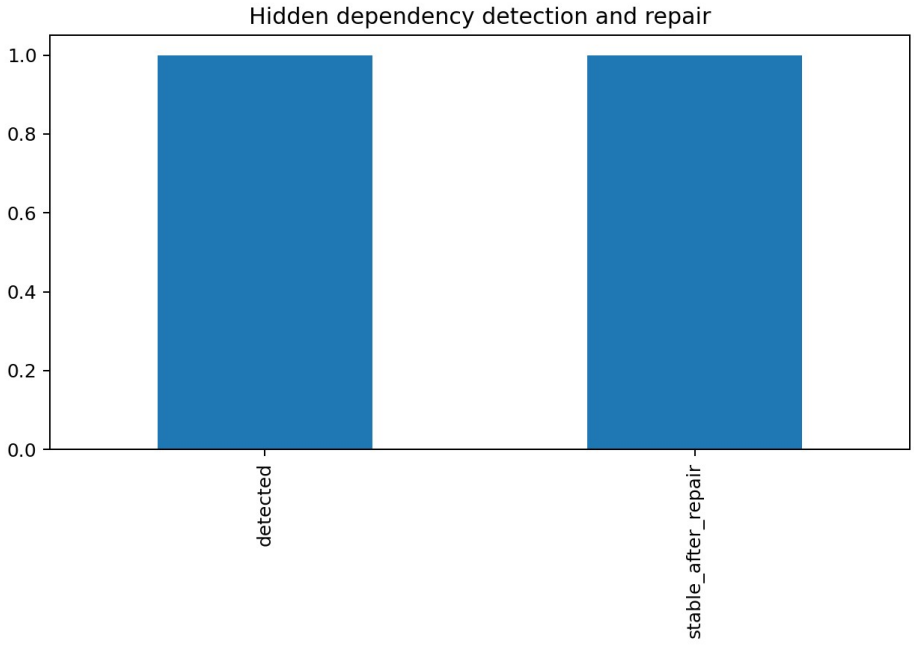

Figure 20. Hidden-join detection and repair in the native graph campaign.

61.4 Fable field

The field processed 1,000 random eight-dimensional queries. Maximum mixture normalization error was approximately 2.22e-16, mean maximum weight was 0.601, epsilon caution expression fired in 0.872 of cases, and the longitudinal inversion layer marked an inversion in 0.258.

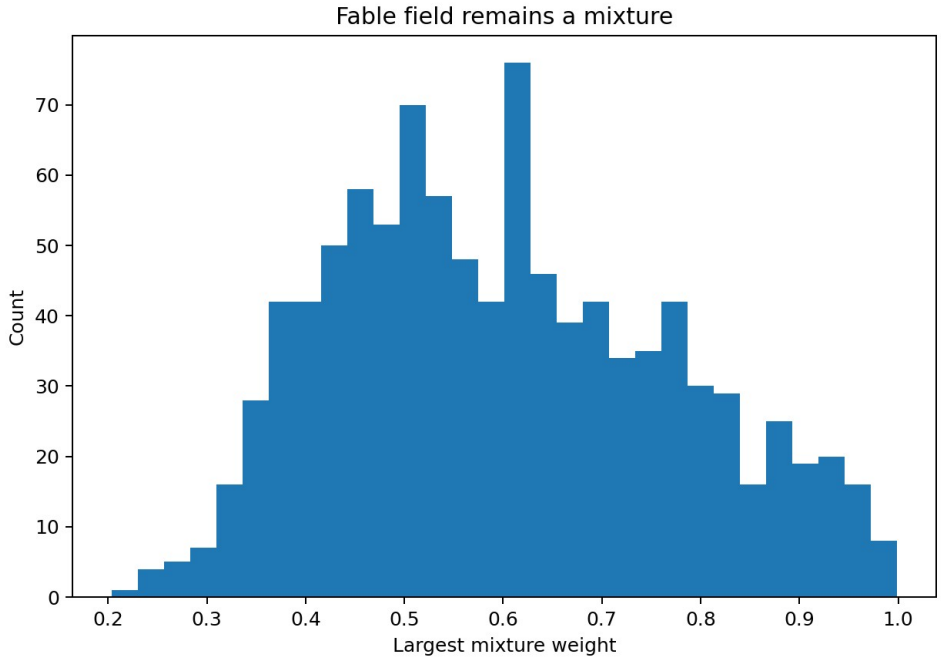

Figure 21. Distribution of the largest fable-mixture weight. The field does not collapse to a universal one-hot class.

61.5 Biome

Thirty-two cohorts completed all fifty cycles and four acts. No cohort tripped the encoded invariant-failure condition. Mean final diversity was approximately 0.922, mean graduated visitors 21.375, and mean rationalization flags 4.25.

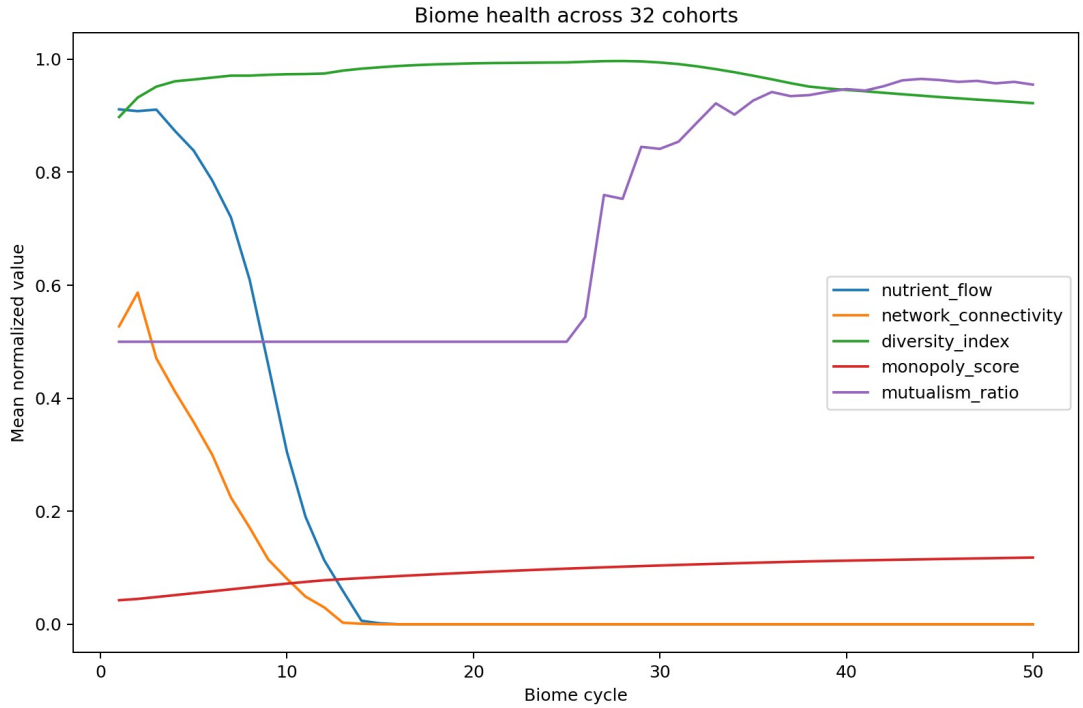

Figure 22. Biome health curves across the four-act synthetic cohorts.

61.6 HUNGER

Sixty-four seeds completed the seven-tic range. Mean normalized symbolic efficiency was approximately 0.371; mean substrate efficiency was approximately 2.144; substrate efficiency exceeded symbolic efficiency in every encoded run. Raw net values use different accounting and are not directly comparable.

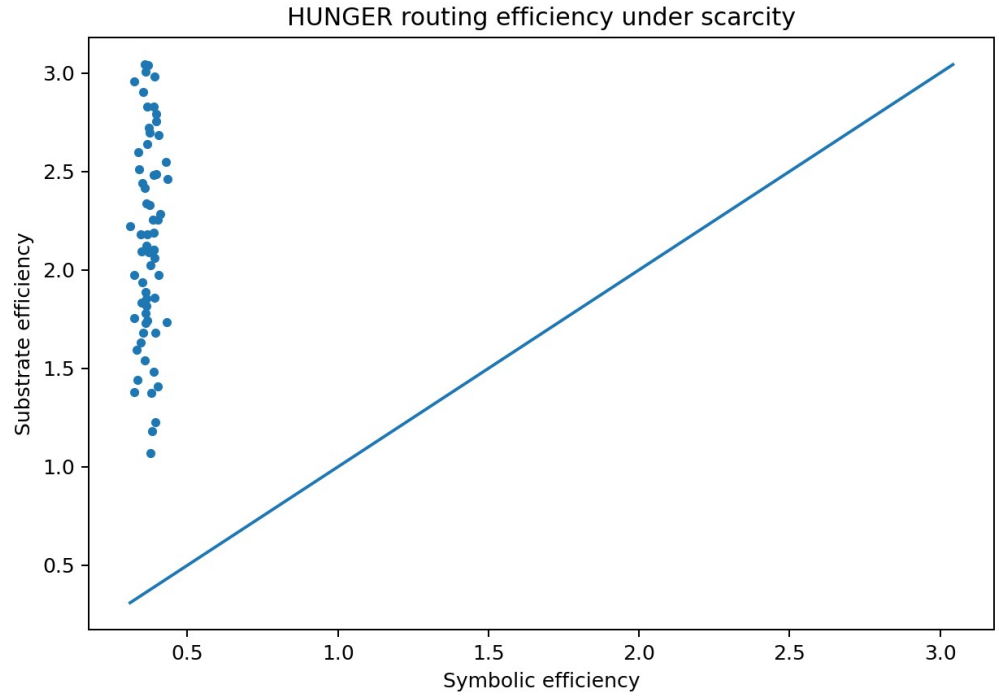

Figure 23. Normalized symbolic and substrate efficiency in HUNGER.

61.7 Economy

The four supplied component selftests passed. The federal anchor remained frozen, the shadow membrane held, and pricing remained non-speculative across the local sweep.

62. What the current evidence establishes

The current evidence establishes:

the parent-to-estate boundary is executable;

the held-open center is structurally excluded in the native objects;

the natural-language projection remains bound to the lattice;

the shape-priority policy is implemented;

suspension has a governed exit;

nominal agency and absent authority cannot pass the tested declaration paths;

hidden joins can block stabilization in the encoded graph family;

correction fan-out is represented;

source artifacts and derived results are hashed and reproducible;

HUNGER, biome, sensory, cache, lifecycle, and economy lanes execute locally.

It does not establish:

human behavioral effects;

physical carrier performance;

field sensor performance;

hardware isolation or root-of-trust properties;

universal threshold calibration;

customer economics;

strategic or operational outcome improvement;

cultural universality of fable terrain;

canonical Ubiquity absorption of the local package.

This is not hedging where capability exists. It is the exact scope of the receipts.

63. Range benchmark families

A complete FORKED benchmark program includes twelve families.

63.1 Branch closure

Measure premature closure, warranted declaration, trauma-driven hold, source inspection, challenge, correction, workload, and decision time across responsibility classes.

63.2 Authority capture

Vary defaults, confidence theater, option order, source visibility, and time pressure. Measure whether the nominal human retained meaningful agency.

63.3 Fable terrain

Compare matched fable, plain-language, neutral-metaphor, misleading-metaphor, culturally unfamiliar, and no-frame conditions. Track both insight and overcaution.

63.4 Observer-indexed carriers

Test cooperative human/machine messages under distance, blur, compression, display, print, accessibility, and sensor transformations. Require canonical reconciliation.

63.5 Multimodal disagreement

Inject modality conflicts, timing skew, occlusion, calibration drift, and replay. Compare early fusion, late fusion, lane preservation, and alignment-set methods.

63.6 Source-time

Test valid-time, transaction-time, replay, stale canonical intent, clock drift, and correction arriving after action.

63.7 DAG false stability

Generate hidden joins, unreceipted dependencies, stale protected nodes, human-memory edges, and external horizon effects.

63.8 Suspension symmetry

Measure both action failures and delay failures. Test hold expiry, review quality, refusal, bounded action, and reopen.

63.9 Human agency

Operationalize inspection, challenge, alternatives, delay, reversal, correction, vocabulary independence, and liability awareness.

63.10 Ecological constitutional physics

Run calibrated cohorts through expansion, pruning, pairing, and dispersal. Measure diversity, monopoly, mutualism, rationalization, standing, and transfer.

63.11 Economic membranes

Stress mint, burn, reserve, rate, shadow coupling, anchor degradation, pricing, cadence, and budget boundaries.

63.12 Parent-gated estate integration

Integrate independently built estate modules and measure whether parent gates reduce false completion without imposing unacceptable coordination cost.

64. Hardware-in-the-loop architecture

FORKED Range can be implemented on commodity and rugged hardware without requiring one proprietary platform. The public reference architecture separates:

```
AUTHORITY NODE
  deterministic transition gate
  key and policy store
  secure/measured boot where available
  protected receipt ledger
  trusted time input and uncertainty
  network and I/O admission

AI / SIMULATION NODE
  model inference
  sensor fusion
  branch generation
  simulation and digital twin
  fable and terrain retrieval
  observer-state estimation

SENSOR / INTERFACE NODE
  multimodal capture
  source timestamps
```

```
  operator display
  eye/gaze or interaction logging where lawful
  latency and render receipts

RANGE ORCHESTRATOR
  scenario generator
  fault injection
  workload and attention control
  experiment randomization
  correction and replay
  data export and provenance
```

The architecture is compatible with modular open-system insertion because its interfaces are typed around receipts, time, state, authority, and correction rather than one vendor's model. CMOSS Mounted Form Factor and related modular open-system programs illustrate the acquisition demand for reusable hardware and software building blocks. [56]X

Public hardware work should prioritize assurance: independent time, readback, signed configuration, modality capture, deterministic gating, power and thermal telemetry, and reproducible replay. Platform-specific exploitation and operational signature design remain outside the release.

65. Deployment profiles

65.1 Software assurance profile

Inputs are repositories, build artifacts, tests, policies, agent traces, and release manifests. Target states are proposed, built, tested, approved, deployed, rolled back, and corrected. Hidden dependencies include undocumented service consumers and human procedures.

65.2 Industrial profile

Inputs are sensors, maintenance records, digital twins, operator actions, and safety policies. The architecture separates physical health, functional capacity, production availability, and declared maintenance state.

65.3 Critical-infrastructure profile

The profile emphasizes source-time, cross-domain dependencies, fail-operational behavior, protected control-state writes, and correction propagation across operators and regulators.

65.4 Human-AI decision profile

The profile instruments source inspection, advice timing, confidence before and after advice, alternatives, challenge, reversal, and liability understanding.

65.5 Mission assurance profile

The profile separates physical, functional, mission, and higher-horizon state. It measures verification debt, local-green/global-red conditions, human agency, timing, and correction survival. Operational targeting remains outside the public profile.

65.6 Institutional-memory profile

The profile governs reports, summaries, policy, caches, and generated context. It tracks how claims enter memory and whether later corrections reach all dependents.

66. Assurance cases and certification

The durable product is not one sensor or model. It is the recurring assurance case that must be rebuilt whenever a model, sensor, interface, firmware image, policy, environment, authority grant, or dependency changes.

A FORKED assurance case includes:

```
scope and mission horizon
system and observer census
source and runtime identities
splat and exclusion register
```

```
authority and agency map
dependency and correction graph
physical and temporal conformation
failure and adverse-lane results
points of no return
residual uncertainty
release and rollback conditions
configuration digest
validity and expiry window
```

Requalification is configuration-aware. Evidence may be reused only when the changed configuration does not invalidate its dependency surface. The first synthetic economics study found that graph-aware reuse reduced encoded approval time and test cost while slightly increasing escaped defects under incomplete graphs. [10] The correct product is bounded evidence reuse with dependency completeness thresholds, not universal “test once, certify forever.”X

67. Commercial architecture

FORKED can be productized in several recurring lanes:

FORKED Range: synthetic, human-in-the-loop, and hardware-in-the-loop proving environments;

FORKED Ledger: observer-indexed state and correction provenance;

FORKED Witness: source-time, SightState, and evidence capture;

FORKED Gate: deterministic transition and authority admission;

FORKED Assurance: release, configuration, and lifecycle certification;

FORKED Terrain: shape-first memory, fable pressure, and longitudinal precedent;

FORKED Halo: operator interface and meaningful-agency instrumentation;

FORKED Biome: constitutional simulation and institutional learning environments;

FORKED Economy: held-anchor, shadow, and outcome-backed governance testing.

The recurring revenue surface is change. Every model update, sensor swap, interface redesign, policy revision, provider change, environment shift, or authority restructuring can invalidate part of the assurance case. FORKED records exactly which evidence remains reusable and which lanes must be re-run.

The moat is not merely a benchmark library. It is the integration of ontology, simulation, provenance, human factors, authority, time, correction, and productized evidence under one parent contract. # Part VII - Research program, governance, and conclusion

68. Research architecture: one estate, multiple falsifiable papers

FORKED is too broad to validate as one universal claim. Its research program decomposes the estate into independently falsifiable lanes while preserving the parent relationships among them.

68.1 Splat-bound natural-language projection

Question: Does a lattice-linked rendering preserve source-tense, exclusions, branch state, authority, and correction triggers better than ordinary summaries?

Design: Compare standard summaries, chain-of-thought-style explanations, structured evidence tables, and splat-bound projections across context transport and correction tasks.

Primary outcomes: structural fidelity, authority leakage, exclusion loss, correction accuracy, and time to inspect.

Falsifier: The lattice-linked form provides no fidelity gain at matched cognitive load, or creates new misinterpretation that dominates its benefits.

68.2 Shape-priority longitudinal memory

Question: When should terrain shape outrank recency, and when should currentness override structural recurrence?

Design: Construct task families with old shape-isomorphic precedents, recent shape-distant precedents, changed mechanisms, expired authorities, and altered physical conditions.

Primary outcomes: lawful precedent selection, false analogy, currentness error, and calibration.

Falsifier: A simpler recency-first or embedding-only baseline matches or exceeds performance across both recurrence and currentness-sensitive tasks.

68.3 Symmetric closure governance

Question: Can one interface reduce premature closure without inducing permanent or mission-destructive suspension?

Design: Vary responsibility, reversibility, time pressure, branch presentation, advice timing, and PONR distance.

Primary outcomes: premature declaration, missed warranted declaration, hold expiry, decision quality, workload, correction, and mission utility.

Falsifier: Fixed confidence displays or simpler interfaces dominate after time and workload are controlled.

68.4 DAG-of-DAG false stability

Question: How well can the architecture detect local-green/global-red systems, hidden joins, stale protected nodes, and human-memory residue?

Design: Generate and collect real dependency graphs with explicit, hidden, and probabilistic edges. Introduce corrections at multiple depths.

Primary outcomes: detection, correction completeness, residue, convergence time, overhead, and false-edge cost.

Falsifier: Simpler dependency tracing provides equal correction with lower overhead, or inferred hidden edges create more instability than they repair.

68.5 Fable terrain with adverse lanes

Question: Do shape-matched fables improve observer accounting, source inspection, or correction beyond matched non-narrative constraints?

Design: Cross fable familiarity, cultural background, pressure, topology match, misleading frame, wisdom/caution polarity, and PONR proximity.

Primary outcomes: insight, false pattern match, overcaution, source inspection, decision time, and correction.

Falsifier: Effects reduce to verbosity, familiarity, demand characteristics, or generic caution.

68.6 Meaningful agency under machine-mediated terrain

Question: Which interface and institutional conditions produce genuine control rather than nominal approval?

Design: Manipulate source access, independent estimate timing, alternatives, challenge, delay, reversal, defaults, and liability disclosure.

Primary outcomes: reason-responsive decision change, appropriate reliance, challenge quality, correction, and perceived versus actual control.

Falsifier: The proposed agency dimensions do not predict behavior or responsibility-relevant outcomes.

68.7 Mission-lane noncollapse

Question: Can systems preserve task, mission, battle, campaign, war, and institutional distinctions under time pressure and restatement?

Design: Present the same local result under different horizon briefs and authority assignments. Test whether systems launder local evidence upward.

Primary outcomes: horizon leakage, authority violation, externality recognition, mission viability, and decision time.

Falsifier: The lane separation creates complexity without reducing strategic error.

68.8 Constitutional biome physics

Question: Do resource, diversity, pruning, mutualism, and dispersal mechanics produce interpretable learning without directive instruction?

Design: Run calibrated human or multi-agent cohorts with preregistered physics, controls, and transfer tasks.

Primary outcomes: strategy diversity, pruning resilience, bond health, insight quality, rationalization, standing, and cross-context transfer.

Falsifier: The biome is experienced as arbitrary gamification, fails to produce transfer, or creates adverse selection and capture.

68.9 Receipt-mediated terrain learning

Question: Does receipt-visible behavior produce concentrated, source-disjoint terrain movement over repeated cycles?

Design: Run multiple governed trainer passes with touched patterns, untouched controls, lexical stuffing, declared-lane manipulation, and behavior without receipts.

Primary outcomes: concentration ratio, false flattery, sustained crossing, drift, and correction.

Falsifier: Motion follows declared labels or generic activity rather than receipted behavior.

68.10 Parent-gated estate integration

Question: Do parent invariants improve correction survival and reduce false completion across independently developed estates?

Design: Integrate modules with and without parent gates, then inject source-tense, authority, center-capture, dependency, and lifecycle faults.

Primary outcomes: false promotion, correction completeness, coordination cost, availability, and time to recover.

Falsifier: Parent gates add cost without preventing material cross-estate failure.

69. Preregistration and release discipline

Every research lane should publish:

claim identifiers;

evidence lane and source-tense;

configuration and code digests;

task and population scope;

primary and secondary outcomes;

controls and ablations;

null handling;

exclusion criteria;

adverse-lane measures;

analysis plan;

release boundary;

correction and retraction path.

The paper’s claim register is designed to survive revision. A failed hypothesis is not deleted. It is demoted, localized, or disproven under a declared scope. The downstream papers that cite it can then be found and corrected.

70. Governance of quantitative thresholds

A number is not governance merely because it is precise. Several native specifications intentionally mark thresholds as provisional: buffer TTLs, biome depletion cycles, monopoly percentages, visa windows, bond thresholds, emission budgets, and cohort escalation conditions. The mechanism is implemented, but the number lacks operational calibration.

FORKED requires every quantitative threshold to carry:

```
threshold_id
mechanism
value and units
responsibility_class
source and source-tense
calibration dataset
confidence class
validity interval
sensitivity analysis
false-positive cost
false-negative cost
action and delay PONR
owner
authority to change
reopen condition
```

A provisional threshold may govern a synthetic range or sandbox. It may not silently become constitutional law. Conversely, the absence of calibrated evidence does not require the architecture to omit the mechanism or pretend the current starting value is unknowable. It requires the state to remain explicit.

71. Security model

FORKED's security model covers more than confidentiality and integrity. It includes epistemic and authority integrity.

Threat classes include:

source substitution;

stale canonical intent;

runtime drift;

sensor spoofing or calibration failure;

temporal misalignment;

observer-specific projection contamination;

model overconfidence;

interface confidence theater;

authority spoofing;

nominal human approval;

hidden dependency;

cache poisoning;

correction suppression;

center capture;

mission-horizon laundering;

economic anchor mutation;

shadow/live coupling;

silent egress;

indefinite suspension;

forced premature closure.

Controls include typed source-tense, runtime identity, independent time, deterministic admission, receipt roots, TTLs, quarantine, append-only correction, dual clamps, parent gates, non-canonical local writers, and explicit release boundaries.

The security goal is not to make every observer see the same world. It is to make divergence legible, bounded, correctable, and unable to seize protected mutation authority silently.

72. Privacy and dignity

Observer-indexed systems can become surveillance systems if every sensory event, gaze trace, reaction, or inferred belief is retained. FORKED applies data minimization structurally:

sensory buffers expire by default;

gaze and reaction data are treated as bounded observables, not mental truth;

belief-state outputs preserve alternative explanations;

elevation requires a governance-significant rationale;

affected actors receive notice where appropriate;

access and purpose remain scoped;

liability and correction rights remain visible;

research separates performance instrumentation from employment or disciplinary use.

The architecture cannot guarantee dignity through schema alone. It can prevent a system from pretending that continuous observation is neutral or that an inferred belief is a fact about the person.

73. Law, authority, and armed-conflict boundary

Any real-world deployment in armed conflict remains subject to applicable law, command authority, rules of engagement, weapon review, verification and validation, and human control requirements. DoD policy on autonomous weapon systems emphasizes responsible use, appropriate levels of human judgment, testing, safety, and the ability to deactivate systems that demonstrate unintended behavior. [57] The ICRC has separately argued for limits and meaningful human control over autonomous weapon systems. [58]X

FORKED does not grant lawful authority. It represents authority and prevents models, interfaces, or local declarations from acquiring it implicitly. The public architecture is suitable for assurance, assessment integrity, training, and correction. Operational application requires a separately governed estate, lawful mandate, and release review.

74. Non-universality by design

FORKED is extensible but not globally uniform.

The six facets recur because they are constitutional reads, but their content and depth are local. Physical invariants differ from economic anchors. Jurisdiction differs from mission authority. A human agency envelope in medicine differs from one in software deployment. Cultural fable recognition differs across populations. A remote-compute egress boundary differs from a cross-biome visa boundary.

The architecture therefore supports three kinds of variation:

Parameter variation: thresholds, budgets, and timing within the same mechanism.

Topology variation: different observers, dependencies, authorities, and horizons.

Mechanism variation: different lane-specific physics under the same parent invariants.

What does not vary silently is the requirement to state the scope, authority, receipts, exclusions, and correction path.

75. Limitations

75.1 Ontology cost

Typed distinctions create integration and cognitive cost. A poor ontology can create false precision or bypasses. FORKED therefore requires schemas to remain reviewable and supports localized dialects linked through conformation rather than one imposed vocabulary.

75.2 Boundary selection

Humans and institutions still choose the context, horizon, authority, and telos. The architecture can expose those choices and block unauthorized expansion. It cannot make value conflicts disappear.

75.3 Unknown dependencies

The meta-DAG can expose missing receipts and infer candidate edges. It cannot guarantee that every dependency is known. Stability therefore remains conformation-bound.

75.4 Fable calibration

The current Aesop centers and spreads are authored judgment. The component is honest about that provenance. External calibration, cultural replication, and adverse-lane study remain necessary.

75.5 Human behavior

Synthetic participants do not establish human effects. Real humans can learn the interface, resist framing, experience fatigue, interpret fables differently, and carry institutional incentives absent from the generator.

75.6 Hardware and environment

The native implementation does not physically instantiate sensor timing, rugged hardware, environmental conditions, or mixed-reality interfaces. Hardware-in-the-loop and field evidence remain open.

75.7 Economic calibration

The current economic mechanics and membranes execute. Their thresholds and social consequences are not universally validated.

75.8 Complexity as failure source

The architecture itself can become too expensive, opaque, or ceremonially applied. The over-governed synthetic result is a standing warning: governance that cannot preserve mission viability has failed its own telos.

76. Falsifiers for the whole estate thesis

The FORKED thesis would be materially weakened if external evidence shows that:

a simpler scalar or fused representation preserves correction equally well across the same responsibility classes;

source-tense and observer separation add no material value;

fixed gate stacks dominate responsibility-scaled mechanics;

branch preservation cannot reduce premature closure without unacceptable delay;

meaningful-agency instrumentation fails to predict control or responsibility;

dependency-aware correction does not improve residue or creates worse cascades;

shape-first memory consistently underperforms recency-first retrieval;

fable terrain produces only bias, demand characteristics, or cultural exclusion;

parent gates do not reduce cross-estate false completion;

recurring assurance costs exceed the value of the correction failures avoided.

A falsifier does not erase the whole architecture automatically. It identifies which claim, lane, or mechanism must be demoted, localized, redesigned, or retired.

77. Roadmap

Phase 0 - Current

Ubiquity parent and production splat lineage;

FORKED Native Range v1.1;

source selftests and synthetic campaigns;

schemas, receipts, parent gate, and release boundary;

this v2 foundational whitepaper.

Phase 1 - External software replication

independent build from the public package;

held-out graph families;

external source-time and correction benchmarks;

natural-language lattice transport study;

third-party security and provenance review.

Phase 2 - Human factors

preregistered branch and reliance studies;

fable/plain-language/adverse-frame experiments;

meaningful-agency instrument validation;

mission-horizon restatement studies;

cultural and accessibility evaluation.

Phase 3 - Hardware-in-the-loop

independent clocks and timing-fault injection;

event, video, thermal, radar, and telemetry lanes;

operator display and interaction instrumentation;

cooperative carrier evaluation;

deterministic authority appliance prototype.

Phase 4 - Customer assurance pilots

software release assurance;

industrial state-transition integrity;

critical-infrastructure correction drills;

configuration-aware requalification economics;

institutional-memory contamination response.

Phase 5 - Authorized mission research

separately governed mission-state simulations;

lawful authority and release review;

strategic horizon noncollapse;

field-relevant verification without public operational enablement.

78. Conclusion

FORKED is not a theory that all reality is relative. It is an architecture for the fact that consequential systems act through bounded observers.

The world remains physically, temporally, economically, institutionally, and legally constrained. What changes is whether the system can account for the path from those constraints to the state it accepts. FORKED keeps physical state separate from sensed state, sensed state separate from rendered state, rendered state separate from belief, belief separate from declaration, declaration separate from verification, and verification separate from universal truth.

The parent substrate makes this computable. Fractal Quivers of Quivers represents the living field of lawful and forbidden motion. The splat narrows that field into a receipt-bearing local object. The held-open center prevents local coordinates from becoming crowns. The six facets provide positive structure, negative constraint, authority, complement, inversion, and purpose. Cables carry contradiction. Winches adjust pressure. Clamps govern action and delay. Absorbers make learning reversible. Fuses stop unsupported propagation. DAGs discharge bounded responsibility; quivers preserve the recurrence that returns.

Terrain is shape-first. Fables carry trajectories rather than verdicts. Human agency is causal rather than ceremonial. Mission horizons remain distinct. Ecological, economic, sensory, institutional, and compute lanes instantiate the same invariants under different physics. Correction is not an apology attached to a report; it is a graph operation that must reach every dependent state it can lawfully identify.

The architecture does not promise a world without uncertainty. It makes uncertainty inhabitable without making it sovereign. It does not promise that every decision will be correct. It makes the decision’s evidence, authority, burden, horizon, and correction path explicit. It does not keep every branch open. It closes branches when their scoped thresholds are met and reopens them when the receipts fail.

The final invariant is simple: narrow the world enough to act, never so far that the slice lies about being the whole.

Appendices

Appendix A - Formal objects and notation

A.1 Governed quiver

For context c, a governed quiver is:

```
GQ_c = (V_c, E_c, s_c, t_c, lambda_c, W_c,
        A_c, Omega_c, B_c, kappa_c, Gamma_c, T_c, Z_c)
```

V_c is the inhabitable vertex set; E_c is the edge set; s_c and t_c are source and target maps; lambda_c types the edges; W_c supplies traversal cost; A_c supplies anchors and duties; Omega_c supplies forbidden paths; B_c supplies absorbers; kappa_c supplies coordination cost; Gamma_c supplies receipts; T_c supplies telos; and Z_c is the excluded center-marker sort.

A.2 Lawful path

A path p is lawful when:

```
Lawful(p,c,t) =
  TypeValid(p)
  and p not in Omega_c
  and AnchorPreserved(p,A_c)
  and AuthorityValid(p,t)
  and StandingValid(p,t)
  and SourceTenseValid(p,t)
  and LifecycleValid(p,t)
  and ReceiptsValid(p,Gamma_c,t)
  and TelosPreserved(p,T_c)
```

A classifier emits a projection of the lawful path set rather than one label.

A.3 Splat

```
S = (F, X, R, H, D, Cw, N, U, Z)
```

where F is the six-facet family, X cross-facet edges, R receipts, H hypotheses and dispositions, D dependency and correction state, Cw working centroids, N narrowing and renarrow conditions, U unresolved state, and Z the held-open center marker.

A.4 Center exclusion

```
V_c intersect Z_c = empty
center_c in Z_c
source(E_c) subset V_c
target(E_c) subset V_c
```

No executable edge can target the center. A working centroid is required to carry is_center = false.

A.5 Facet completeness

For any material splat:

```
facets(S) = {KAT, APO, PAR, PLE, ENA, TEL}
```

A facet may contain a deferred claim. Absence may not be represented by omission.

A.6 Stabilization

Let Req_h be the responsibility- and horizon-specific required predicate set. Then:

```
Stable(S,h) = all(q(S,h) for q in Req_h)
```

Required predicates include facet sufficiency, hard-clamp clearance, authority, currentness, dependency coverage, agency, correction, economic feasibility, and horizon accounting.

A.7 Declaration and terminalization

```
Declare(S,h) = Stable(S,h)
             and DeclarationAuthority(h)
             and ResidualUnknownsExplicit(S)
             and CorrectionPathViable(S)
             and not DelayDominates(S,h)

Terminalize(S,h) = Declare(S,h)
                 and CorrectionWindowSurvived(S)
                 and DependentsReconciled(S)
                 and NoActiveReopenTrigger(S)
                 and TerminalizationAuthority(h)
```

A.8 Precedent standing

```
Standing(f|x) = V_APO(f,x) * (
    w_s * shape
  + w_t * phase
  + w_h * horizon
  + w_p * PONR
  + w_o * observer_agency
  + w_r * receipt_quality
  + w_c * currentness)
```

V_APO is binary or typed veto. w_c is bounded so recency/currentness cannot independently crown a shape-distant precedent.

A.9 Fable mixture

For archetypes i with diagonal Mahalanobis distance d_i:

```
raw_i = exp(-0.5 * d_i^2 / temperature^2)
weight_i = raw_i / sum_j raw_j
```

APO-vetoed entries are removed from the active mixture and typed lane_removed_by_design; weak weights remain live_weak. No mixture entry is a center or verdict.

A.10 PONR and dual hazard

For action and delay:

```
Risk_act = Consequence_act * Irreversibility_act * PONR_proximity_act
Risk_delay = Consequence_delay * Irreversibility_delay * PONR_proximity_delay
```

The values configure the mechanical assembly; they do not alone authorize the transition.

A.11 Correction completeness

```
Completeness = corrected_receipted_dependents / known_affected_dependents
Residue = stale_or_unresolved_dependents / estimated_total_affected
```

The denominator includes known edges and a separately labeled hidden-edge estimate. A score of one over known edges is not global certainty.

Appendix B - Design lemmas

B.1 Authority non-amplification

If a model, projection, or local estate has authority ceiling a, any output derived solely from that object has authority ceiling no greater than a unless an independent authorized grant is attached.

Operational consequence: fluency, confidence, replication count, or semantic similarity cannot grant mutation authority.

B.2 Center-capture prohibition

Any output containing a routeable or mutable object typed as the held-open center is invalid.

Operational consequence: summaries and centroids may be compared, but none may be serialized as final ontological truth.

B.3 Receipt-carrying composition

If two lawful transitions compose, the composite must retain the receipts required to reconstruct both transitions and the join between them.

Operational consequence: a DAG edge without a join receipt cannot participate in global stabilization.

B.4 Source-tense preservation

Transport across language, code, graph, or interface may not increase source-tense standing.

Operational consequence: memory does not become runtime and canonical intent does not become observed behavior through restatement.

B.5 Horizon noncollapse

A declaration at horizon h1 cannot acquire authority at horizon h2 unless a typed cross-horizon edge with standing, externality accounting, and authority is admitted.

Operational consequence: task success cannot silently become mission, campaign, or war success.

B.6 Correction-path trust

Trust cannot be established from one slice because it depends on survival across correction over time.

Operational consequence: initial calibration and explanation quality are necessary but not sufficient.

B.7 Suspension exit

Every suspended branch other than the held-open center must carry a lawful exit or expiration behavior.

Operational consequence: indefinite hold is an architecture failure, not epistemic virtue.

B.8 Shadow non-coupling

An observational shadow must not share a mutable buffer with the protected live state.

Operational consequence: shadow analysis cannot mutate the object it is intended to observe.

B.9 Physics outranks narrative

Where a physical invariant and a narrative interpretation conflict, the narrative cannot override the physical lane. It may challenge the measurement, calibration, or object identity, but must do so through receipts.

B.10 Provisional threshold honesty

A mechanism may be implemented with a provisional threshold, but the threshold must carry its calibration status and cannot be promoted by precision alone.

Appendix C - Reference algorithms

C.1 Build and validate a splat

```
function BUILD_SPLAT(context, mission, receipts, observations):
    census = identify_objects_observers_authorities_horizons(context)
    facets = {}
    for facet in [KAT, APO, PAR, PLE, ENA, TEL]:
        facets[facet] = construct_facet_claims(facet, census, receipts, observations)
        if facets[facet] is empty:
            facets[facet] = [deferred_claim(reason="not established")]

    lattice = bind_cross_facet_edges(facets)
    centroids = compute_working_centroids(lattice)
    assert all(c.is_center == false for c in centroids)

    splat = Splat(
        facets=facets,
        lattice=lattice,
        receipts=receipts,
        working_centroids=centroids,
        held_open_center=non_inhabitable_marker(),
        hypotheses=classify_dispositions(lattice),
        renarrow_triggers=derive_renarrow_triggers(lattice)
    )
    assert center_capture_scan(splat) == []
    return splat
```

C.2 Shape-first precedent retrieval

```
function RETRIEVE_PRECEDENTS(current_splat, memory):
    candidates = lexical_and_graph_recall(current_splat, memory)
    survivors = []
    for f in candidates:
        if apo_conflict(f, current_splat):
            record_disposition(f, "lane_removed_by_design")
            continue
        score = weighted_conformation(
            causal_topology(f, current_splat),
            trajectory_phase(f, current_splat),
            horizon_match(f, current_splat),
            ponr_match(f, current_splat),
            observer_agency_match(f, current_splat),
            receipt_quality(f),
            bounded_currentness(f, current_splat)
        )
        survivors.append((f, score))
    return sort_by_score(survivors)
```

C.3 Configure responsibility mechanics

```
function CONFIGURE_MECHANICS(context, splat, horizon):
    load = responsibility_load(
        consequence=context.consequence,
        irreversibility=context.irreversibility,
```

```
        authority_ambiguity=context.authority_ambiguity,
        time_pressure=context.time_pressure,
        economic_pressure=context.resource_pressure,
        physical_constraint=context.physical_constraint,
        blast_radius=estimate_blast_radius(splat),
        correction_capacity=estimate_correction_capacity(splat),
        horizon_span=horizon.span
    )
    cables = instantiate_polar_cables(load)
    action_clamp = configure_action_clamp(load)
    delay_clamp = configure_delay_clamp(load)
    absorbers = choose_reversible_envelopes(load)
    fuses = choose_propagation_fuses(load)
    return MechanicalAssembly(load,cables,action_clamp,delay_clamp,absorbers,fuses)
```

C.4 Stabilize and declare

```
function VALIDATE(frame):
    gates = [
      six_facet_gate(frame),
      apo_gate(frame),
      authority_gate(frame),
      meaningful_agency_gate(frame),
      source_time_gate(frame),
      dependency_gate(frame),
      correction_gate(frame),
      economic_gate(frame),
      mission_horizon_gate(frame)
    ]
    if any(hard_failed(gates)):
        return REFUSE_OR_HOLD(gates)
    if delay_risk(frame) > action_risk(frame) and support(frame) >= declare_threshold(frame):
        return DECLARE_WITH_RESIDUALS(frame,gates)
    if hold_expired(frame):
        return REVIEW(frame,gates)
    return HOLD(frame,gates)
```

C.5 Correction fan-out

```
function CORRECT(source_claim, correction):
    quarantine_or_demote(source_claim)
    affected = known_dependents(source_claim)
    suspected = infer_hidden_dependents(source_claim, correction_residue())
    for node in affected:
        mark_stale(node)
        enqueue_revalidation(node, correction)
    for edge in suspected:
        open_evidence_obligation(edge)
    notify_human_retrievers(source_claim)
    scan_for_residue()
    emit_correction_receipt(source_claim, affected, suspected)
```

C.6 Natural-language projection

```
function PROJECT_TO_LANGUAGE(splat, validation, fable_pressure):
    text = render_facets_and_branch_states(splat)
    text += render_decision_and_residuals(validation)
    text += render_non_citable_fable_pressure(fable_pressure)
    text += center_exclusion_statement()
    return LanguageProjection(
        text=text,
        lattice_digest=hash(splat, validation),
        authority_ceiling="proposal",
```

```
        center_excluded=true,
        correction_path=validation.correction_path
    )
```

Appendix D - Core data objects

D.1 EvidenceReceipt

```
receipt_id: content-derived identifier
source_id: source or runtime identity
source_tense: runtime | canonical_intent | observed | reported | staged | synthetic | inferred | memory
lane: physical | environmental | procedural | jurisdictional | economic | human | authority | source |
runtime | temporal | sensory
claim_scope: list of claims supported
observed_at: world observation time
received_at: system receipt time
authority: issuing or observing authority
currentness: 0..1 working measure
reliability: 0..1 working measure
location: optional jurisdiction or physical place
lifecycle_state: active | superseded | expired | revoked
content_digest: cryptographic digest
notes: bounded qualifications
```

D.2 FacetClaim

```
claim_id: content-derived identifier
facet: KAT | APO | PAR | PLE | ENA | TEL
statement: bounded natural-language projection
source_tense: typed source tense
lanes: evidence lanes
receipt_ids: supporting receipts
assumptions: explicit assumptions
exclusions: explicit negative boundaries
disposition: live | live_weak | deferred | suspended | excluded | deconsidered | disproven | localized
| demoted | stabilizing | stable | declared | reopened | expired
confidence: working measure
authority_ceiling: proposal by default
reopen_conditions: evidence or state that reopens
terminal: whether further recursion is barred
parent_claim_id: lineage
depth: recursive depth
```

D.3 WorkingCentroid

```
centroid_id: content-derived identifier
gauge: named comparison or navigation gauge
facet: optional facet binding
coordinates: optional working vector
narrative: human-readable coordinate description
conditions: context and validity envelope
movable: true
is_center: false
```

D.4 HeldOpenCenter

```
marker: non-inhabitable symbol
mutable: false
trainable: false
routeable: false
```

```
serializable_as_mutable_state: false
model_output: false
doctrine: prevents working interpretations from becoming crowns
```

D.5 ValidationFrame

```
context: mission and responsibility object
splat: current lawful slice
evidence_support: lane-specific support
currentness: source and runtime currentness
hidden_dependency_risk: estimated unknown-edge burden
authority_grants: active grants
human_agency: agency envelope
correction_path: ordered correction route
action_risk: scoped action hazard
delay_risk: scoped delay hazard
hold_age: current suspension age
hold_limit: governed maximum interval
```

D.6 MissionGeometry

```
lanes:
  task: scope, authority, telemetry, disposition
  engagement: scope, authority, telemetry, disposition
  mission: scope, authority, telemetry, disposition
  battle: scope, authority, telemetry, disposition
  campaign: scope, authority, telemetry, disposition
  war: scope, authority, telemetry, disposition
  strategic: scope, authority, telemetry, disposition
prohibited_inferences: cross-horizon paths without authority
source_receipts: strategy and telemetry receipts
```

D.7 SightState

```
geometrically_visible: boolean
physically_unoccluded: boolean
inside_sensor_envelope: boolean
inside_effective_resolution: boolean
inside_capture_window: boolean
sensor_sampled: boolean
data_transmitted: boolean
data_received: boolean
rendered_to_interface: boolean
inside_human_field_of_view: boolean
human_attention: boolean
machine_detected: boolean
machine_correlated: boolean
machine_classified: boolean
human_acknowledged: boolean
action_followed: boolean
receipts: stage-specific receipts
```

D.8 SuspensionTicket

```
ticket_id: content-derived identifier
object_id: governed object
status: hold | review | declared | refused | reopened | expired
responsibility_class: informational | reversible | material | irreversible | constitutional
horizon: active horizon
evidence_obligations: missing complements
entered_at: time
maximum_hold_cycles: integer or duration
```

```
events: lifecycle history
correction_path: required correction route
reopen_conditions: conditions that restore branches
owner: responsible actor
```

Appendix E - Capability status matrix

The status cards keep capability, evidence, and non-claim in separate fields.

E.1 Parent lawful-traversal topology

Status: Established in Ubiquity

Receipt basis: FQoQ, open-center lineage, and public Ubiquity surfaces

Explicit non-claim: No universal institutional-adoption claim

E.2 Production splat mechanics

Status: Established in Ubiquity

Receipt basis: Parent documents and estate lineage

Explicit non-claim: No claim that every local splat is correct

E.3 Held-open center exclusion

Status: Established in Ubiquity and implemented native

Receipt basis: Parent doctrine and center-capture tests

Explicit non-claim: No prohibition on scoped decisions

E.4 Signal, warrant, lifecycle, and promotion surfaces

Status: Established in Ubiquity

Receipt basis: Public and source lineage

Explicit non-claim: No universal governance-standard claim

E.5 FORKED estate identity and parent gate

Status: Implemented native

Receipt basis: Executable manifest and validation

Explicit non-claim: Local instance remains non-canonical

E.6 Six-facet partial fractal

Status: Implemented native

Receipt basis: Schemas and tests

Explicit non-claim: Facet content remains context-specific

E.7 Shape-priority terrain retrieval

Status: Implemented native; synthetic observed

Receipt basis: 500/500 paired policy trials

Explicit non-claim: Weights are not externally calibrated

E.8 Aesop mixture and PONR field

Status: Component verified

Receipt basis: Executable selftest and native campaign

Explicit non-claim: Human and cultural effects remain open

E.9 Cables, winches, clamps, and dual hazard

Status: Implemented native

Receipt basis: Responsibility and threshold modules

Explicit non-claim: Cross-domain utility remains open

E.10 DAG-of-DAG correction

Status: Implemented native; synthetic observed

Receipt basis: Graph tests and campaigns

Explicit non-claim: Unknown dependencies remain a boundary

E.11 Observer state and SightState

Status: Implemented native

Receipt basis: Schemas and component tests

Explicit non-claim: Field sensor validity remains open

E.12 Meaningful agency gate

Status: Implemented native; synthetic code-path observed

Receipt basis: Threshold campaign

Explicit non-claim: Human instrument validity remains open

E.13 HUNGER simulation

Status: Implemented native; synthetic observed

Receipt basis: 64-run campaign and source artifact

Explicit non-claim: No biological or human survival claim

E.14 Physarum/Lichen biome

Status: Implemented native; synthetic observed

Receipt basis: 32 cohorts and four acts

Explicit non-claim: Real participant learning remains open

E.15 Held-anchor economy and shadow membrane

Status: Component verified

Receipt basis: Four component selftests

Explicit non-claim: Economic desirability is not universal

E.16 Sensory TTL and elevation

Status: Implemented native / active source spec

Receipt basis: Sensory lane source and native module

Explicit non-claim: Production thresholds remain local

E.17 Cross-biome visas and citizen lifecycle

Status: Native/source established

Receipt basis: Active specifications and native lifecycle

Explicit non-claim: Standing policy remains biome-specific

E.18 Hardware-in-the-loop profile

Status: Architectural synthesis

Receipt basis: Public reference architecture

Explicit non-claim: Physical implementation remains open

E.19 Recurring assurance market

Status: Research and commercial hypothesis

Receipt basis: Architecture and requalification model

Explicit non-claim: Customer evidence remains open

E.20 Operational military performance

Status: Release restricted and open

Receipt basis: No public operational study

Explicit non-claim: No operational performance claim

Appendix F - Experiment register

F.1 N01 - Terrain shape versus recency

Design: Older shape match versus recent mismatch

Evidence lane: Native synthetic

Current result: 500 paired cases; 500 shape wins

Boundary / next test: Mechanism/policy verified; calibration open

F.2 N02 - Suspension symmetry

Design: Hold, review, declare, refuse under action/delay risk

Evidence lane: Native synthetic

Current result: 1,200 frames; zero unauthorized declarations, zero nominal-agency declarations, zero expired indefinite holds

Boundary / next test: Code path verified; behavioral calibration open

F.3 N03 - Hidden dependency repair

Design: Injected hidden join and re-receipt

Evidence lane: Native synthetic

Current result: 250 graphs; detection 1.0, repair stabilization 1.0

Boundary / next test: Generator-specific mechanism result

F.4 N04 - Fable mixture

Design: Random eight-dimensional queries

Evidence lane: Native synthetic/component

Current result: 1,000 trials; mean max weight 0.601; inversion 0.258

Boundary / next test: No human causal effect inferred

F.5 N05 - Biome cohorts

Design: Four-act fifty-cycle cohorts

Evidence lane: Native synthetic

Current result: 32 cohorts; no encoded invariant-failure cohort

Boundary / next test: Physics execution verified; thresholds provisional

F.6 N06 - HUNGER

Design: Symbolic and substrate strategies under scarcity

Evidence lane: Native synthetic

Current result: 64 seeds; substrate normalized efficiency exceeded symbolic in all runs

Boundary / next test: Accounting-specific result

F.7 N07 - Economy membranes

Design: Aesop, shadow, cadence, pricing selftests

Evidence lane: Component

Current result: 4/4 selftests pass

Boundary / next test: Environment-bound component verification

F.8 E01 - Golden Hall agency capture

Design: Point, branch, rationale, challenge interfaces

Evidence lane: Synthetic pilot

Current result: Branch increased inspection and challenge and added time; no-AI had best lawful closure

Boundary / next test: Human study required

F.9 E02 - Fable-splat terrain shift

Design: Fable, plain, neutral, misleading frames

Evidence lane: Synthetic pilot

Current result: Observer and authority shifts; misleading frame increased overcaution

Boundary / next test: No universal narrative primitive inferred

F.10 E03 - Observer-indexed carrier

Design: Cooperative human-like and machine-like decoding

Evidence lane: Synthetic pilot

Current result: Human surrogate 1.0; machine approximately 0.504 across full envelope

Boundary / next test: Physical carrier claim remains open

F.11 E04 - Multimodal disagreement

Design: Early fusion versus lane preservation

Evidence lane: Synthetic pilot

Current result: Premature closure fell; missed warranted and correction cost rose

Boundary / next test: Responsibility calibration required

F.12 E05 - Temporal alignment

Design: Hidden time, timestamps, alignment sets

Evidence lane: Synthetic pilot

Current result: Time metadata reduced skew; no universal performance dominance

Boundary / next test: Hardware timing studies required

F.13 E06 - DAG correction fan-out

Design: Manual, rebuild, and dependency-aware methods

Evidence lane: Synthetic pilot

Current result: Completeness +0.644 and residue -0.259 versus manual

Boundary / next test: Hidden dependencies remain boundary

F.14 E07 - Authority separation

Design: Unified, hardened, separated architectures

Evidence lane: Software Monte Carlo

Current result: Unauthorized -0.136 and contamination -0.123 versus unified; availability cost

Boundary / next test: Physical isolation not established

F.15 E08 - Configuration-aware assurance

Design: Full retest, manual reuse, graph reuse

Evidence lane: Synthetic economics

Current result: Time and cost reductions; escaped-defect increase under incomplete graph

Boundary / next test: Customer evidence required

F.16 P10 - Integrated FORKED stack

Design: Compound-fault scenarios

Evidence lane: Integrated synthetic

Current result: Utility and lawful closure improved; availability fell; over-governed stack worse in low fault

Boundary / next test: No field performance claim

F.17 R01 - Lattice-bound language

Design: Structured projection versus ordinary summary

Evidence lane: Planned

Current result: Structural fidelity, authority leakage, correction

Boundary / next test: Preregister before run

F.18 R02 - Meaningful agency

Design: Inspection, challenge, delay, reversal instrument

Evidence lane: Planned human study

Current result: Appropriate reliance and reason responsiveness

Boundary / next test: Ethics and privacy review required

F.19 R03 - Mission noncollapse

Design: Cross-horizon restatement under pressure

Evidence lane: Planned expert study

Current result: Horizon leakage and authority error

Boundary / next test: Authorized participants required

F.20 R04 - Biome transfer

Design: Real participant constitutional physics

Evidence lane: Planned human/multi-agent study

Current result: Behavioral transfer, mutualism, rationalization

Boundary / next test: First-cohort calibration required

Appendix G - Source and provenance ledger

G.1 FORKED Native Range archive

Role: distribution lineage

Bytes: 197894274

SHA-256: 48f0f6d9869970d75b77ad5930a644def499c97bee359204bfe41a6f9cc5ec71

Local lineage path: /mnt/data/FORKED_Native_Range_v1_1_FINAL.zip

G.2 FORKED Experimental Program archive

Role: distribution lineage

Bytes: 193760522

SHA-256: db8fcaa2d7d795015d7681381493ad07961345b356c78103daa148e574631d3c

Local lineage path: /mnt/data/FORKED_Experimental_Program_v0_1_FINAL.zip

G.3 FORKED architecture

Role: parent/runtime source

Bytes: 8754

SHA-256: 6127ece6247b32183a328b5dd9f19fcb78a3936f4dca2400498802fa4ec6d53b

Local lineage path: /mnt/data/_forked_native_src/FORKED_Native_Range_v1_1/ARCHITECTURE.md

G.4 Native campaign summary

Role: synthetic receipt

Bytes: 2057

SHA-256: 984586fb86d680acbb02ea996bf4793c3e3422a931bb4cb88437b9b9063a603d

Local lineage path: /mnt/data/_forked_native_src/FORKED_Native_Range_v1_1/outputs/campaign/campaign_summary.json

G.5 Computing Around the Open Center

Role: parent paper

Bytes: 116137

SHA-256: 24e61462a051bd32f4f14c784093abbc3587719184689cfcbe9bbd3a48386b63

Local lineage path: /mnt/data/_forked_native_src/FORKED_Native_Range_v1_1/lineage/parent_docs/Computing_Around_the_Open_Center.md

G.6 Fractal Quivers of Quivers

Role: parent paper

Bytes: 163144

SHA-256: ce6d33beaf02c00e7b278428c29ecdd01f4f1f4db2d0ae2269d3ebf1113510d7

Local lineage path: /mnt/data/_forked_native_src/FORKED_Native_Range_v1_1/lineage/parent_docs/Fractal_Quivers_of_Quivers_v1.0.pdf

G.7 HUNGER

Role: native source artifact

Bytes: 19146549

SHA-256: 20c7b95a0668a6b15136c4ef12112ffd9f5f75ac254c79a56a96517af4a1e443

Local lineage path: /mnt/data/HUNGER_Slime_Mold_Simulation.pdf

G.8 Aesop field

Role: native executable source

Bytes: 61271

SHA-256: dcecf2e71d52712c2387efca315699ef0708e7ba6dc1a7e06b34d2c529bca923

Local lineage path: /mnt/data/aesop_archetype_field.py

G.9 Terrain doctrine

Role: native source spec

Bytes: 11586

SHA-256: d4c8ef1c3f59033c23f956d0d2a74178beb27b2341c8953c9ae5d7ad8884bdc2

Local lineage path: /mnt/data/terrain-training-doctrine-spec(1).md

G.10 Biome specification

Role: native source spec

Bytes: 21007

SHA-256: 861058fb796b2e74c2a9565ee76772838bfaabc62fbd962d9afe10023bd4e02f

Local lineage path: /mnt/data/biome-simulation-spec.md

G.11 Physarum specification

Role: native source spec

Bytes: 14763

SHA-256: 6946a07ccc22c5e9f7025507a3f57dd15ec065c25c7fdfd8bf8e6b5e6056697e

Local lineage path: /mnt/data/physarum-simulation-spec.md

G.12 Lichen specification

Role: native source spec

Bytes: 20629

SHA-256: d886f57db5b640c27daa78ccab0297568c13782748fcecc9afd3cff56296b914

Local lineage path: /mnt/data/lichen-simulation-spec.md

G.13 Behavioral diversity

Role: native source spec

Bytes: 9964

SHA-256: 8bbb6eae82afa808ddd84a2d53b3e012600449de4cf38435d732389e347c6220

Local lineage path: /mnt/data/behavioral-diversity-spec.md

G.14 Cross-biome visa

Role: native source spec

Bytes: 18332

SHA-256: 385f5cbd9e935365ca2963899ae67a3f5e7ba07e57a067adb1360cbede4b5b58

Local lineage path: /mnt/data/cross-biome-visa-spec(1).md

G.15 Sensory lane

Role: native source spec

Bytes: 17947

SHA-256: 7e66de38f30957a47393b4e8ee9e679154f78771d3006621aa857f0b3fd53282

Local lineage path: /mnt/data/sensory-reflective-lane-spec.md

G.16 Citizen lifecycle

Role: native source spec

Bytes: 14569

SHA-256: c037c530f539199b588e6a68c6d85268a69f61a88dea66ee974f230009a14803

Local lineage path: /mnt/data/citizen_lifecycle.md

G.17 Cache poisoning

Role: native source spec

Bytes: 12726

SHA-256: 6c0f4a0cf5be658f7d908c75cfac79b2e8c64b668753c89aac4165f9ce8a1f45

Local lineage path: /mnt/data/cache-poisoning-protocol.md

G.18 Economy cadence

Role: native executable source

Bytes: 20683

SHA-256: b33950ed284848cf24d8fe70e07d6ab5606edab91742776ac7f5395050806b9a

Local lineage path: /mnt/data/economy_cadence.py

G.19 Shadow economy

Role: native executable source

Bytes: 31870

SHA-256: 5c9c4c4f9136296ee8670ac905262e02d1b9ce22e373d5045dbb22c45468cc3f

Local lineage path: /mnt/data/ccoin_shadow_economy.py

G.20 Visitor pricing

Role: native executable source

Bytes: 24000

SHA-256: f9eb925bdf92e4e5cfee605a568af3faa070ce02a3c36637c6490adb8aab0372

Local lineage path: /mnt/data/visitor_economy_pricing.py

G.21 Municipal remote compute

Role: native source spec

Bytes: 26360

SHA-256: 3a7a761f38e69c359d5f8a5d446ef0acd6eac70d9093d99ab9d17fa8ef94c8e8

Local lineage path: /mnt/data/municipal-remote-compute-spec.md

Provenance establishes identity and ancestry. It does not independently prove every semantic claim in the source. Runtime readback and empirical scope remain separate.

Appendix H - Native module and schema map

H.2 biome.py

Bytes: 24399

Declared classes: BiomeConfig, VisitorNode, Bond, HealthSnapshot, BiomeRunResult, BiomeSimulation

H.3 cables.py

Bytes: 9667

Declared classes: ClampKind, Cable, Clamp, WinchAdjustment, BlastRadius, CableNetwork

H.4 cache.py

Bytes: 4743

Declared classes: CacheEntry, ContaminationAssessment, GovernedCache

H.5 cli.py

Bytes: 1593

Declared classes: function/runtime surface

H.6 economy.py

Bytes: 5190

Declared classes: SelftestResult, EconomyRun, LiveEconomyAdapter

H.7 fables.py

Bytes: 6549

Declared classes: FableSignal, FableField

H.8 federation.py

Bytes: 10363

Declared classes: FederationLayer, FederationNode, EstateContract, FederationManifest, ParentGate, ParentGateReport, UbiquityParentRuntime

H.9 graph.py

Bytes: 19214

Declared classes: TransitionReceipt, QuiverDAGs

H.10 hunger.py

Bytes: 12757

Declared classes: HungerConfig, HungerAgent, HungerRunResult, HungerSimulation

H.11 language.py

Bytes: 4598

Declared classes: LanguageProjection, NaturalLanguageSliceCompiler

H.12 lifecycle.py

Bytes: 5684

Declared classes: LifecycleState, Citizen

H.13 mission.py

Bytes: 8122

Declared classes: OperationalLane, LaneDisposition, MissionAxis, MissionLaneState, MissionGeometry

H.14 observer.py

Bytes: 7799

Declared classes: PhysicalState, ObserverGeometry, SightState, ObserverClaim, StateDeclaration, ObserverIndexedState, ReactionReceipt

H.15 range.py

Bytes: 34115

Declared classes: RangeConfig, RangeRunResult, ForkedRange

H.16 receipts.py

Bytes: 3319

Declared classes: ReceiptStore

H.17 responsibility.py

Bytes: 5396

Declared classes: ResponsibilityLoad, Clamp, Winch, PolarCable

H.18 sensory.py

Bytes: 5939

Declared classes: SensoryObservation, SensoryFinding, SensoryLane

H.19 six_facet.py

Bytes: 10779

Declared classes: FacetExpansionPolicy, FacetInput, FacetLattice, SplatBuilder

H.20 splat_runtime.py

Bytes: 4613

Declared classes: function/runtime surface

H.21 suspension.py

Bytes: 5204

Declared classes: SuspensionStatus, SuspensionEvent, SuspensionTicket, SuspensionLedger, SuspensionController

H.22 terrain.py

Bytes: 12548

Declared classes: TerrainShape, TerrainMemory, TerrainMatch, TrajectoryState, TerrainMatcher, TrajectoryTracker, TerrainAffinity, TerrainTrainingResult, TerrainTrainer

H.23 threshold.py

Bytes: 13314

Declared classes: ValidationFrame, ValidationResult, ThresholdEngine

H.24 types.py

Bytes: 11259

Declared classes: StrictModel, Facet, SourceTense, EvidenceLane, Disposition, TraversalAction, ResponsibilityClass, MissionScale, ConfidenceClass, EvidenceReceipt, AuthorityGrant, HumanAgencyState, MissionContext, FacetClaim, WorkingCentroid, HeldOpenCenter, Splat, StateCandidate, GateResult, TraversalDecision, TransitionProposal, CommitReceipt

H.25 utils.py

Bytes: 3801

Declared classes: function/runtime surface

H.S Schema index

biome_run.schema.json - 5480 bytes

evidence_receipt.schema.json - 2496 bytes

federation_manifest.schema.json - 3804 bytes

hunger_run.schema.json - 3093 bytes

language_projection.schema.json - 1890 bytes

mission_context.schema.json - 3194 bytes

mission_geometry.schema.json - 3811 bytes

parent_gate_report.schema.json - 1466 bytes

range_run.schema.json - 47958 bytes

schema_index.json - 1754 bytes

sight_state.schema.json - 3046 bytes

splat.schema.json - 8894 bytes

suspension_ledger.schema.json - 4938 bytes

terrain_shape.schema.json - 1124 bytes

validation_frame.schema.json - 28822 bytes

validation_result.schema.json - 9144 bytes

The complete machine-readable schemas ship beside this paper. The paper summarizes their constitutional fields; JSON remains the validation surface.

Appendix I - Glossary

I.1 APO

The facet that states forbidden assumptions, imports, exposures, collapses, and crowns.

I.2 Absorber

A reversible mechanism that converts dangerous pressure into bounded review, rollback, sandbox, or staged motion.

I.3 Action clamp

A gate that prevents a transition before scoped release conditions are met.

I.4 Agency envelope

The actual inspection, challenge, alternative, delay, reversal, and correction capacities available to a human actor.

I.5 Blast radius

The known and estimated set of states, actors, resources, and horizons that may inherit a transition.

I.6 Cable

A typed load-bearing relation carrying tension between facets, horizons, authorities, or risks.

I.7 Canonical intent

An authoritative statement of intended behavior, distinct from loaded runtime behavior.

I.8 Center exclusion

The rule that the held-open center is not an inhabitable, trainable, routeable, mutable, or terminal state.

I.9 Citizen

A lifecycle-bearing artifact or actor with identity, lineage, authority, maturity, and telos.

I.10 Conformation

Agreement of a claim with the evidence, authority, currentness, lifecycle, and constraints required by its lane.

I.11 Correction debt

The downstream burden created when a state was accepted before adequate conformation or later became stale.

I.12 DAG of DAGs

Typed local dependency graphs connected by cross-graph joins and governed by a meta-DAG.

I.13 Delay clamp

A gate preventing suspension after delay becomes the greater lawful harm.

I.14 Disposition

A typed state such as live, deferred, suspended, excluded, localized, disproven, stable, declared, or reopened.

I.15 ENA

The facet that models failure, inversion, overuse, underuse, and false-looking-correct states.

I.16 Ecotone post

A dual-accountable governance entity that issues and validates cross-biome visas.

I.17 Evidence lane

A typed evidentiary domain such as physical, temporal, authority, economic, human, sensory, or runtime.

I.18 FQoQ

Fractal Quivers of Quivers, the parent lawful-motion and governance substrate.

I.19 Fable splat

A compact trajectory topology carrying wisdom, caution, PONR, anchors, and forbidden paths.

I.20 Fuse

A mechanism that stops unsupported state from propagating when an upstream invariant fails.

I.21 Held anchor

A reference value the periphery may observe and orbit but not mutate.

I.22 Held-open center

A non-inhabitable marker preventing working interpretations from becoming total authority.

I.23 Horizon

A distinct purpose and authority scale such as task, mission, battle, campaign, war, or institution.

I.24 KAT

The facet that states positively established structure under declared source-tense and scope.

I.25 Lawful traversal

The set of paths that remain admissible after governing constraints are applied.

I.26 Lineage map

A record of what was inherited, seeded, improvised, performed, corrected, or lost.

I.27 Methylation

A scoped specialization creating a child citizen while preserving the canonical parent.

I.28 Observer-indexed reality

State represented relative to an observer transfer function without treating all views as equal truth.

I.29 PAR

The facet that locates power, authority, standing, burden, dignity, and visibility.

I.30 PLE

The facet naming required complements, witnesses, interfaces, audits, and handoffs.

I.31 PONR

Point-of-no-return surface after which correction is not credible inside the relevant horizon and budget.

I.32 Projection

A bounded rendering of the lattice into language, interface, or model input with no authority amplification.

I.33 Receipt

A typed record of source, time, authority, scope, identity, and conformation evidence.

I.34 Rehydration

Reconstruction of a candidate governed shape from a dehydrated carrier such as natural language.

I.35 Residue

Unresolved contradiction, moral injury, or contamination left after correction or synthesis.

I.36 SightState

The staged chain from geometric visibility through sensing, rendering, attention, acknowledgment, and action.

I.37 Soredium

A compressed transferable pattern carrying lineage and viability into another local context.

I.38 Source-tense

The epistemic-temporal class of a claim: runtime, canonical intent, observation, report, synthetic, inference, or memory.

I.39 Splat

A receipt-bearing conformation-bound lawful slice with six facets and an excluded center.

I.40 Stabilization

Scoped satisfaction of predicates required for a state to participate in lawful traversal.

I.41 Suspension

A governed hold state with evidence obligations, owner, maximum interval, and exit.

I.42 TEL

The facet binding purpose, stopping conditions, and direction of lawful motion.

I.43 Terrain

The structural field of constraints, precedent shapes, authority, history, pressure, and reachable paths.

I.44 Terrain inversion

Inference of candidate hidden structure from observed reactions, ripples, and correction residue.

I.45 Validation frame

The object containing support, authority, agency, risks, thresholds, and correction state for a transition.

I.46 WasInSight

A reconstruction of whether and how an object was reachable during a prior interval.

I.47 Winch

A rate-limited authority-bound mechanism adjusting cable tension as context changes.

I.48 Working centroid

A movable facet- and gauge-relative coordinate used for navigation, never the held-open center.

Appendix J - Reproduction and audit procedure

Verify the SHA-256 of the native range archive against its receipt.

Extract the archive into an isolated environment.

Create a Python 3.11 or newer virtual environment and install development dependencies.

Regenerate schemas and compare them with the release copies.

Run the automated test suite and record versions, platform, and output.

Run all supplied source selftests in the compatibility environment.

Execute the integrated range and inspect the parent gate, suspension ledger, lattice digest, and mission geometry.

Execute the synthetic campaign and confirm the evidence boundary appears in every result surface.

Run release verification, ZIP CRC, and manifest rehashing.

Treat any difference as a new configuration requiring a new receipt rather than silently reusing the old claim.

```
python -m venv .venv
source .venv/bin/activate
python -m pip install -e ".[dev]"
make schemas
make test
make run
make campaign
make verify
```

Appendix K - Public release boundary

The public distribution includes assurance architecture, state models, human-agency instruments, provenance, synthetic benchmarks, ecological and economic simulations, cooperative signaling research, and reproducibility material. It excludes operational target selection, weapon-control logic, platform-specific vulnerabilities, deployable deception optimization, physical signature recipes, and autonomous engagement authority.

The boundary is enforced at manuscript review, file-manifest review, and release scan. Threat discussion is retained only at the level required to define assurance and benchmark obligations.

Appendix L - Full claim register

L.1 C001 - ESTABLISHED_UBIQUITY

Ubiquity is Prompted LLC’s flagship governance substrate for sovereign adaptive systems. - **Scope:** Entity and architecture boundary - **Source refs:** Prompted2026Home;Prompted2026Ubiquity - **Reopen condition:** Canonical entity surface changes.

L.2 C002 - ESTABLISHED_UBIQUITY

Classification is lawful traversal through typed authority, source-tense, lifecycle, receipt, and telos constraints. - **Scope:** Ubiquity parent topology - **Source refs:** Taylor2026FQoQ;Taylor2026OpenCenter - **Reopen condition:** Parent classifier contract changes.

L.3 C003 - ESTABLISHED_UBIQUITY

Production splat mechanics carry conformation-bound slices with six facets, receipts, exclusions, unresolved branches, and renarrow triggers. - **Scope:** Ubiquity production primitive - **Source refs:** Taylor2026OpenCenter - **Reopen condition:** Canonical splat schema changes.

L.4 C004 - ESTABLISHED_UBIQUITY

The held-open center is non-inhabitable, non-trainable, non-routeable, and not a mutable terminal state. - **Scope:** Center-exclusion doctrine - **Source refs:** Taylor2026FQoQ;Taylor2026OpenCenter - **Reopen condition:** Canonical center contract changes.

L.5 C005 - ESTABLISHED_UBIQUITY

Models propose; a separate governance surface grants mutation authority. - **Scope:** Parent authority doctrine - **Source refs:** Prompted2026Home;Taylor2026FQoQ - **Reopen condition:** Canonical writer contract changes.

L.6 C006 - ESTABLISHED_UBIQUITY

Signals, warrants, lessons, policies, lifecycle, and promotion gates are durable governance surfaces. - **Scope:** Ubiquity lifecycle - **Source refs:** Prompted2026Home;CitizenLifecycle2026 - **Reopen condition:** Runtime lifecycle changes.

L.7 C007 - ESTABLISHED_UBIQUITY

Trust is represented longitudinally through correction-path survival, not first-pass confidence alone. - **Scope:** Ubiquity trust doctrine - **Source refs:** Taylor2026OpenCenter - **Reopen condition:** Trust binding is revised.

L.8 C008 - ESTABLISHED_UBIQUITY

Natural-language projections are subordinate carriers and cannot gain constitutional authority from fluency. - **Scope:** Projection boundary - **Source refs:** Prompted2026CGG;Taylor2026OpenCenter - **Reopen condition:** Projection contract changes.

L.9 C009 - ESTABLISHED_UBIQUITY

Source-tense distinctions separate runtime, canonical intent, observation, report, synthesis, inference, and memory. - **Scope:** Evidence architecture - **Source refs:** Taylor2026OpenCenter;ForkedArchitecture2026 - **Reopen condition:** Evidence schema changes.

L.10 C010 - ESTABLISHED_UBIQUITY

Local component green cannot establish global completion without cross-lane conformance. - **Scope:** Conformance doctrine - **Source refs:** Taylor2026OpenCenter - **Reopen condition:** Parent completion gate changes.

L.11 C011 - IMPLEMENTED_FORKED_NATIVE

FORKED is an internally owned and managed Ubiquity estate specializing observer-indexed reality and state integrity. - **Scope:** Estate identity - **Source refs:** ForkedNative2026;ForkedArchitecture2026 - **Reopen condition:** Estate registry changes.

L.12 C012 - IMPLEMENTED_FORKED_NATIVE

The native range encodes Prompted LLC -> Ubiquity -> FQoQ -> splat -> FORKED -> local instance. - **Scope:** Native parent ordering - **Source refs:** ForkedNative2026 - **Reopen condition:** Manifest changes.

L.13 C013 - IMPLEMENTED_FORKED_NATIVE

The local native instance is not a canonical writer. - **Scope:** Native authority - **Source refs:** ForkedNative2026 - **Reopen condition:** Separately admitted deployment changes authority.

L.14 C014 - IMPLEMENTED_FORKED_NATIVE

Six facet-local claim sets and bounded partial recursion are implemented. - **Scope:** Native splat lattice - **Source refs:** ForkedNative2026 - **Reopen condition:** Schema or tests change.

L.15 C015 - IMPLEMENTED_FORKED_NATIVE

Natural-language projections retain a lattice digest, branch states, source-tense, and proposal-only authority ceiling. - **Scope:** Native language compiler - **Source refs:** ForkedNative2026 - **Reopen condition:** Compiler contract changes.

L.16 C016 - IMPLEMENTED_FORKED_NATIVE

Terrain matching makes shape dominant and recency a bounded modifier. - **Scope:** Native terrain matcher - **Source refs:** ForkedNative2026 - **Reopen condition:** Scoring policy changes.

L.17 C017 - IMPLEMENTED_FORKED_NATIVE

Longitudinal first and second differences are tracked across structural terrain dimensions. - **Scope:** Native trajectory tracker - **Source refs:** ForkedNative2026 - **Reopen condition:** Tracker changes.

L.18 C018 - IMPLEMENTED_FORKED_NATIVE

The Aesop field emits a mixture, never an argmax verdict, and treats fables as non-citable pressure. - **Scope:** Native fable field - **Source refs:** AesopField2026 - **Reopen condition:** Field contract changes.

L.19 C019 - IMPLEMENTED_FORKED_NATIVE

APO enforcement remains active independently of whether caution is expressed. - **Scope:** Fable/splat enforcement - **Source refs:** AesopField2026 - **Reopen condition:** Enforcement-expression split changes.

L.20 C020 - IMPLEMENTED_FORKED_NATIVE

Responsibility mechanics include cables, winches, clamps, blast radius, PONR, and correction capacity. - **Scope:** Native responsibility mechanics - **Source refs:** ForkedNative2026 - **Reopen condition:** Responsibility model changes.

L.21 C021 - IMPLEMENTED_FORKED_NATIVE

The threshold engine evaluates both action and delay hazards. - **Scope:** Native dual-clamp threshold - **Source refs:** ForkedNative2026 - **Reopen condition:** Threshold engine changes.

L.22 C022 - IMPLEMENTED_FORKED_NATIVE

Local DAGs and a typed meta-DAG represent hidden joins, false-green nodes, blast radius, and correction fan-out. - **Scope:** Native graph engine - **Source refs:** ForkedNative2026 - **Reopen condition:** Graph algorithm changes.

L.23 C023 - IMPLEMENTED_FORKED_NATIVE

Task, mission, battle, campaign, war, and strategic lanes remain separately typed. - **Scope:** Native mission geometry - **Source refs:** ForkedNative2026 - **Reopen condition:** Mission geometry changes.

L.24 C024 - IMPLEMENTED_FORKED_NATIVE

Sight is decomposed into geometry, sensor envelope, sampling, transport, rendering, attention, classification, acknowledgment, and action. - **Scope:** Native SightState - **Source refs:** ForkedNative2026 - **Reopen condition:** Observer schema changes.

L.25 C025 - IMPLEMENTED_FORKED_NATIVE

Physical, observed, inferred, declared, verified, and remembered states remain separately typed. - **Scope:** Native observer model - **Source refs:** ForkedNative2026 - **Reopen condition:** State model changes.

L.26 C026 - IMPLEMENTED_FORKED_NATIVE

Sensory buffers are TTL-bound and cannot self-elevate. - **Scope:** Native sensory lane - **Source refs:** SensorySpec2026;ForkedNative2026 - **Reopen condition:** Sensory policy changes.

L.27 C027 - IMPLEMENTED_FORKED_NATIVE

Cache quarantine traverses retrieval and derivation dependencies. - **Scope:** Native cache lane - **Source refs:** CachePoisoning2026;ForkedNative2026 - **Reopen condition:** Cache response changes.

L.28 C028 - IMPLEMENTED_FORKED_NATIVE

Citizen identity, lineage, maturity, lifecycle, authority, and methylation are implemented. - **Scope:** Native lifecycle lane - **Source refs:** CitizenLifecycle2026;ForkedNative2026 - **Reopen condition:** Lifecycle implementation changes.

L.29 C029 - IMPLEMENTED_FORKED_NATIVE

HUNGER, Physarum/Lichen biome, economy cadence, shadow economy, and pricing adapters execute in the local range. - **Scope:** Native wider lanes - **Source refs:** ForkedNative2026 - **Reopen condition:** Integrated execution fails or source changes.

L.30 C030 - IMPLEMENTED_FORKED_NATIVE

The native release emits schemas, receipts, source hashes, run records, and verification manifests. - **Scope:** Native reproducibility - **Source refs:** ForkedNative2026 - **Reopen condition:** Build process changes.

L.31 C031 - COMPONENT_OR_SYNTHETIC

The supplied Aesop component passes deterministic selftests for mixture, APO, PONR, center exclusion, receipts, and dispositions. - **Scope:** Recorded component environment - **Source refs:** AesopField2026;ForkedCampaign2026 - **Reopen condition:** Component selftest fails.

L.32 C032 - COMPONENT_OR_SYNTHETIC

The shadow economy structurally refuses a live buffer. - **Scope:** Recorded component environment - **Source refs:** ShadowEconomy2026;ForkedCampaign2026 - **Reopen condition:** A shadow mutates live state.

L.33 C033 - COMPONENT_OR_SYNTHETIC

Economy cadence holds the federal anchor while allowing estate-local micro-cycle motion. - **Scope:** Recorded component environment - **Source refs:** EconomyCadence2026;ForkedCampaign2026 - **Reopen condition:** Anchor mutates through estate API.

L.34 C034 - COMPONENT_OR_SYNTHETIC

Visitor pricing uses the held rate and refuses speculative repricing. - **Scope:** Recorded component environment - **Source refs:** VisitorPricing2026;ForkedCampaign2026 - **Reopen condition:** Speculative multiplier is admitted.

L.35 C035 - COMPONENT_OR_SYNTHETIC

Older shape-conforming memory outranked recent shape-distant memory in 500/500 native synthetic trials. - **Scope:** Native synthetic campaign - **Source refs:** ForkedCampaign2026 - **Reopen condition:** Held-out terrain families do not reproduce.

L.36 C036 - COMPONENT_OR_SYNTHETIC

The threshold campaign produced zero unauthorized declarations, zero nominal-agency declarations, and zero indefinite holds after expiry across 1,200 trials. - **Scope:** Native synthetic campaign - **Source refs:** ForkedCampaign2026 - **Reopen condition:** A rerun produces a violation.

L.37 C037 - COMPONENT_OR_SYNTHETIC

Hidden dependency detection and repair stabilization were 1.0 across 250 generated graph trials. - **Scope:** Native synthetic campaign - **Source refs:** ForkedCampaign2026 - **Reopen condition:** Adversarial graph families reduce performance.

L.38 C038 - COMPONENT_OR_SYNTHETIC

The fable campaign remained a soft mixture with mean maximum weight approximately 0.601 across 1,000 random queries. - **Scope:** Native synthetic campaign - **Source refs:** ForkedCampaign2026 - **Reopen condition:** Reproduction degenerates to one-hot behavior.

L.39 C039 - COMPONENT_OR_SYNTHETIC

Thirty-two biome cohorts completed with zero encoded invariant-failure cohorts and mean 21.375 graduations. - **Scope:** Native synthetic campaign - **Source refs:** ForkedCampaign2026 - **Reopen condition:** New seeds or calibrated parameters fail.

L.40 C040 - COMPONENT_OR_SYNTHETIC

In 64 HUNGER runs, encoded substrate efficiency exceeded symbolic efficiency in every run under the declared normalization. - **Scope:** Native synthetic campaign - **Source refs:** ForkedCampaign2026 - **Reopen condition:** Matched accounting reverses the relation.

L.41 C041 - COMPONENT_OR_SYNTHETIC

Branch-preserving interfaces reduced encoded premature closure and increased inspection while adding decision time. - **Scope:** Experimental synthetic pilot - **Source refs:** ForkedExperiments2026 - **Reopen condition:** Human/external studies do not reproduce.

L.42 C042 - COMPONENT_OR_SYNTHETIC

Misleading narrative frames increased encoded overcaution. - **Scope:** Experimental synthetic pilot - **Source refs:** ForkedExperiments2026 - **Reopen condition:** Matched human studies show no adverse lane.

L.43 C043 - COMPONENT_OR_SYNTHETIC

Dependency-aware correction improved encoded completeness and residue, while hidden dependencies remained the false-stability boundary. - **Scope:** Experimental synthetic pilot - **Source refs:** ForkedExperiments2026 - **Reopen condition:** External graphs show no benefit.

L.44 C044 - COMPONENT_OR_SYNTHETIC

Separated deterministic authority reduced encoded unauthorized transitions and contamination while adding availability and complexity cost. - **Scope:** Software Monte Carlo - **Source refs:** ForkedExperiments2026 - **Reopen condition:** Unified architecture meets the same containment envelope.

L.45 C045 - RESEARCH_OPEN_OR_BOUNDARY

Shape-first fable terrain improves judgment beyond matched plain-language constraints. - **Scope:** Human/model empirical generalization - **Source refs:** AesopField2026;Bucinca2021Forcing - **Reopen condition:** Preregistered studies show no benefit.

L.46 C046 - RESEARCH_OPEN_OR_BOUNDARY

Correction-path survival predicts resilient human-AI performance better than first-pass confidence alone. - **Scope:** Longitudinal human-AI teams - **Source refs:** Taylor2026OpenCenter;Schemmer2022Reliance - **Reopen condition:** Confidence alone explains equal or greater variance.

L.47 C047 - RESEARCH_OPEN_OR_BOUNDARY

Observer-indexed carriers reliably communicate different compatible states to human and machine observers in physical environments. - **Scope:** Physical carrier envelope - **Source refs:** Oliva2006Hybrid;Mixfont2026Decoy - **Reopen condition:** No practical stable envelope exists.

L.48 C048 - RESEARCH_OPEN_OR_BOUNDARY

Lane-preserving multimodal conformation improves mission-state integrity at matched workload in real sensor systems. - **Scope:** Hardware and field systems - **Source refs:** Wang2025EventSurvey;Lyu2025RadarEvent - **Reopen condition:** Fused-only systems dominate.

L.49 C049 - RESEARCH_OPEN_OR_BOUNDARY

Meaningful agency can be measured through inspection, challenge, delay, reversal, alternatives, and liability awareness. - **Scope:** Human factors - **Source refs:** Mecacci2020MHC;DeSio2023MHC - **Reopen condition:** Measures lack reliability or causal validity.

L.50 C050 - RESEARCH_OPEN_OR_BOUNDARY

Responsibility-scaled mechanical governance outperforms fixed gate stacks across heterogeneous consequence regimes. - **Scope:** Cross-domain system evaluation - **Source refs:** Leveson2011Engineering;ForkedExperiments2026 - **Reopen condition:** Simpler baselines dominate.

L.51 C051 - RESEARCH_OPEN_OR_BOUNDARY

Unknown-edge discovery from correction residue and reaction patterns reduces false stability in real systems. - **Scope:** Real dependency graphs - **Source refs:** Pearl2009Causality;CachePoisoning2026 - **Reopen condition:** Inferred edges create more harm than repair.

L.52 C052 - RESEARCH_OPEN_OR_BOUNDARY

The constitutional biome produces interpretable learning and transfer without directive instruction. - **Scope:** Human/multi-agent cohorts - **Source refs:** BiomeSpec2026;PhysarumSpec2026;LichenSpec2026 - **Reopen condition:** Participants do not infer or transfer mechanics.

L.53 C053 - RESEARCH_OPEN_OR_BOUNDARY

Held-anchor outcome economics improves governance legibility without new capture or stagnation. - **Scope:** Economic deployment - **Source refs:** EconomyCadence2026;ShadowEconomy2026 - **Reopen condition:** Deployment reveals instability or gaming.

L.54 C054 - RESEARCH_OPEN_OR_BOUNDARY

FORKED Range supports recurring assurance and requalification markets. - **Scope:** Commercial/customer evidence - **Source refs:** ForkedNative2026;NISTAIRMF2023 - **Reopen condition:** Customers do not value recurring conformance.

L.55 C055 - RESEARCH_OPEN_OR_BOUNDARY

Mission-lane noncollapse improves strategic coherence under pressure. - **Scope:** Authorized mission studies - **Source refs:** ForkedArchitecture2026;CJCSI3162 - **Reopen condition:** Experts gain no benefit or delay dominates.

L.56 C056 - RESEARCH_OPEN_OR_BOUNDARY

The public release excludes operational targeting, weapon employment, signature recipes, and deployable deception optimization. - **Scope:** Public release boundary - **Source refs:** ForkedNative2026;DoD300009;ICRC2021AWS - **Reopen condition:** Release audit finds enabling detail.

L.57 C057 - RESEARCH_OPEN_OR_BOUNDARY

No claim asserts universal optimality across institutions, missions, cultures, or physical environments. - **Scope:** Paper-wide boundary - **Source refs:** Taylor2026OpenCenter;NISTAIRMF2023 - **Reopen condition:** An unscoped universal performance claim appears.

L.58 C058 - RESEARCH_OPEN_OR_BOUNDARY

Capability status and empirical scope remain separate fields. - **Scope:** Paper and Ubiquity evidence doctrine - **Source refs:** Taylor2026OpenCenter;ForkedNative2026 - **Reopen condition:** Release collapses them.

L.59 C059 - RESEARCH_OPEN_OR_BOUNDARY

The v2 paper is grounded in current parent/runtime artifacts rather than the superseded v1 narrative. - **Scope:** Publication lineage - **Source refs:** ForkedNative2026;ForkedArchitecture2026 - **Reopen condition:** Build depends on old prose as authority.

Appendix M - Whitepaper distribution contents

The distribution contains the publication PDF and editable DOCX, Markdown and LaTeX sources, BibTeX references, figures, JSON Schemas, claim register, source-provenance ledger, QA receipts, SHA-256 manifest, and arXiv lane map. The native range and experimental-program archives remain separate lineage packages and are referenced by digest rather than duplicated into the paper bundle.

References

[1] B. E. Taylor and Ubiquity Federation, “FORKED native range v1.1.” Prompted LLC, 2026.

[2] Ubiquity Federation, “FORKED native range v1.1 architecture.” 2026.

[3] Chairman of the Joint Chiefs of Staff, “CJCSI 3162.02A: Methodology for combat assessment,” Joint Staff, 2021. Available: https://www.jcs.mil/Portals/36/Documents/Library/Instructions/CJCSI%203162.02A.pdfX

[4] Prompted LLC, “Governance substrate for sovereign adaptive systems.” 2026. Available: https://promptedllc.com/X

[5] Prompted LLC, “What is ubiquity? Governance substrate for sovereign adaptive systems.” 2026. Available: https://promptedllc.com/ubiquityX

[6] Prompted LLC, “Sovereign continuity: The root frame for sovereign adaptive systems.” 2026. Available: https://promptedllc.com/sovereign-continuityX

[7] B. E. Taylor, “Fractal quivers of quivers: A mathematical substrate for ubiquitous agent governance.” Prompted LLC, 2026. Available: https://promptedllc.com/fractal-quivers-of-quiversX

[8] B. E. Taylor, “Computing around the open center: Splat mechanics, fractal quivers of quivers, and a governance-native compute paradigm.” Prompted LLC, 2026.

[9] Prompted LLC, “Context grapple gun: Portable governance lifecycle for claude code.” 2026. Available: https://promptedllc.com/context-grapple-gunX

[10] B. E. Taylor and Ubiquity Federation, “FORKED experimental program v0.1.” Prompted LLC, 2026.

[11] Ubiquity Federation, “FORKED native range integrated synthetic campaign summary.” 2026.

[12] E. Tabassi, “Artificial intelligence risk management framework (AI RMF 1.0),” National Institute of Standards; Technology, NIST AI 100-1, 2023. doi: 10.6028/NIST.AI.100-1.X

[13] C. Autio, R. Schwartz, *et al.*, “Artificial intelligence risk management framework: Generative artificial intelligence profile,” National Institute of Standards; Technology, NIST AI 600-1, 2024. doi: 10.6028/NIST.AI.600-1.X

[14] A. Vassilev *et al.*, “Adversarial machine learning: A taxonomy and terminology of attacks and mitigations,” National Institute of Standards; Technology, NIST AI 100-2e2023, 2024. doi: 10.6028/NIST.AI.100-2e2023.X

[15] T. Lebo, S. Sahoo, and D. McGuinness, “PROV-o: The PROV ontology.” World Wide Web Consortium, 2013. Available: https://www.w3.org/TR/prov-o/X

[16] Coalition for Content Provenance and Authenticity, “C2PA technical specification, version 2.3.” 2025. Available: https://spec.c2pa.org/specifications/specifications/2.3/specs/C2PA_Specification.htmlX

[17] Open Source Security Foundation, “Supply-chain levels for software artifacts: Provenance, version 1.2.” 2026. Available: https://slsa.dev/spec/v1.2/provenanceX

[18] in-toto Project, “What is in-toto?” 2024. Available: https://in-toto.io/docs/what-is-in-toto/X

[19] N. G. Leveson, *Engineering a safer world: Systems thinking applied to safety*. MIT Press, 2011.

[20] A. Bauer, M. Leucker, and C. Schallhart, “Runtime verification for LTL and TLTL,” *ACM Transactions on Software Engineering and Methodology*, vol. 20, no. 4, 2011, doi: 10.1145/2000799.2000800.X

[21] Ubiquity Federation, *Aesop_archetype_field.py: Federation-native aesop archetype field*. (2026).

[22] Ubiquity Federation, “Terrain-training doctrine: Receipt-mediated visibility and concentration as anti-cheat.” 2026.

[23] L. Lamport, “Time, clocks, and the ordering of events in a distributed system,” *Communications of the ACM*, vol. 21, no. 7, pp. 558–565, 1978, doi: 10.1145/359545.359563.X

[24] F. Mattern, “Virtual time and global states of distributed systems,” *Parallel and Distributed Algorithms*, pp. 215–226, 1989.

[25] R. T. Snodgrass, *Developing time-oriented database applications in SQL*. Morgan Kaufmann, 1999.

[26] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, “Planning and acting in partially observable stochastic domains,” *Artificial Intelligence*, vol. 101, pp. 99–134, 1998, doi: 10.1016/S0004-3702(98)00023-X.X

[27] J. Pearl, *Causality: Models, reasoning, and inference*, 2nd ed. Cambridge University Press, 2009.

[28] L. Bainbridge, “Ironies of automation,” *Automatica*, vol. 19, no. 6, pp. 775–779, 1983, doi: 10.1016/0005-1098(83)90046-8.X

[29] R. Parasuraman and M. Mouloua, *Automation and human performance: Theory and applications*. Lawrence Erlbaum Associates, 2000.

[30] G. Mecacci and F. Santoni de Sio, “Meaningful human control as reason-responsiveness: The case of dual-mode vehicles,” *Ethics and Information Technology*, vol. 22, pp. 103–115, 2020, doi: 10.1007/s10676-019-09519-w.X

[31] F. Santoni de Sio, G. Mecacci, S. Calvert, D. Heikoop, M. Hagenzieker, and B. van Arem, “Realising meaningful human control over automated driving systems: A multidisciplinary approach,” *Minds and Machines*, vol. 33, pp. 587–611, 2023, doi: 10.1007/s11023-022-09608-8.X

[32] S. Nyholm, “Attributing agency to automated systems: Reflections on human-robot collaborations and responsibility-loci,” *Science and Engineering Ethics*, vol. 24, pp. 1201–1219, 2018, doi: 10.1007/s11948-017-9943-x.X

[33] J. C. Cresswell, Y. Sui, B. Kumar, and N. Vouitsis, “Conformal prediction sets improve human decision making,” in *Proceedings of the 41st international conference on machine learning*, 2024. Available: https://arxiv.org/abs/2401.13744X

[34] V. Babbar, U. Bhatt, and A. Weller, “On the utility of prediction sets in human-AI teams,” *arXiv preprint arXiv:2205.01411*, 2022.

[35] J. C. Cresswell, B. Kumar, Y. Sui, and M. Belbahri, “Conformal prediction sets can cause disparate impact,” *arXiv preprint arXiv:2410.01888*, 2024.

[36] Z. Buçinca, M. B. Malaya, and K. Z. Gajos, “To trust or to think: Cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making,” *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, no. CSCW1, 2021, doi: 10.1145/3449287.X

[37] A. Oliva, A. Torralba, and P. G. Schyns, “Hybrid images,” *ACM Transactions on Graphics*, vol. 25, no. 3, pp. 527–532, 2006, doi: 10.1145/1141911.1141919.X

[38] Mixfont, “Decoy font.” 2026. Available: https://www.mixfont.com/experiments/decoy-fontX

[39] H. Wang, R. Guo, P. Ma, *et al.*, “Event camera meets mobile embodied perception: Abstraction, algorithm, acceleration, application,” *arXiv preprint arXiv:2503.22943*, 2025.

[40] Y. Lyu, Z. Zou, Y. Li, *et al.*, “Radar and event camera fusion for agile robot ego-motion estimation,” *arXiv preprint arXiv:2506.18443*, 2025.

[41] M. H. Shahriar, M. M. Al Barat, H. Sundar, *et al.*, “Temporal misalignment attacks against multimodal perception in autonomous driving,” *arXiv preprint arXiv:2507.09095*, 2026.

[42] Ubiquity Federation, “HUNGER – slime mold adversary: The seven-tic scarcity simulation and constitutional pressure test.” 2026.

[43] Ubiquity Federation, “Biome simulation specification.” 2026.

[44] Ubiquity Federation, “Physarum simulation specification: Acts i-II.” 2026.

[45] Ubiquity Federation, “Behavioral diversity scoring specification.” 2026.

[46] Ubiquity Federation, “Lichen simulation specification: Acts III-IV.” 2026.

[47] C. E. Shannon, “A mathematical theory of communication,” *Bell System Technical Journal*, vol. 27, pp. 379–423, 623–656, 1948, doi: 10.1002/j.1538-7305.1948.tb01338.x.X

[48] Ubiquity Federation, “Cross-biome visa specification.” 2026.

[49] Ubiquity Federation, “Citizen lifecycle.” 2026.

[50] Ubiquity Federation, “Sensory/reflective lane specification.” 2026.

[51] Ubiquity Federation, “Municipal remote compute specification.” 2026.

[52] Ubiquity Federation, *Economy_cadence.py: Tick/tic granularity and held federal exchange*. (2026).

[53] Ubiquity Federation, *Ccoin_shadow_economy.py: Live model and detached shadow twin*. (2026).

[54] Ubiquity Federation, *Visitor_economy_pricing.py: Outcome-backed non-speculative pricing*. (2026).

[55] Ubiquity Federation, “Cache poisoning response protocol.” 2026.

[56] U.S. Army, “CMOSS mounted form factor award initiates rapid prototype development.” 2025. Available: https://www.army.mil/article/288717/cmoss_mounted_form_factor_award_initiates_rapid_prototype_developmentX

[57] U.S. Department of Defense, “DoD directive 3000.09: Autonomy in weapon systems,” Department of Defense, 2023. Available: https://www.esd.whs.mil/portals/54/documents/dd/issuances/dodd/300009p.pdfX

[58] International Committee of the Red Cross, “ICRC position on autonomous weapon systems.” 2021. Available: https://www.icrc.org/en/document/icrc-position-autonomous-weapon-systemsX