

Synthetic Diagnostic Package

Software Decision Intelligence Demonstration

Repository under analysis · **REPO X**

This package is provided **exclusively for demonstration purposes**. Repo X is a synthetic enterprise repository model used solely to demonstrate the RepoTruth methodology.

No customer repository. No customer data. No customer conclusions.

Read-only

Snapshot-based

No production access

No credentials

Human review terminal

Observation ≠ Execution

Confidence ≠ Certainty

Interpretation ≠ Recommendation

Evidence ≠ Truth

Priority ≠ Decision

Unknown ≠ Low Risk

How to read this package · evidence flows toward a human decision, never past it



Artifact RT-SDP-REPOX-0001 · Diagnostic Package (Package ≠ Verdict)

Constitutions RT-DIO · RT-SC · RT-CC · RT-UPS · RT-HRW · RT-ERM · RT-DPA · RT-PRC · validated by RT-CE

Status Generated for demonstration · awaiting human review (terminal)

A reconstructed reality model, projected for the people who decide.

Executive Takeaways

The reconstructed reality of this repository, in thirty seconds. Illustrative only, no scores, no grades.

The reconstructed repository reality shows

- ✓ Strong governance concentration
- ✓ Observable boundary maturity
- ✓ Preserved continuity mechanisms
- ✓ Visible historical sediment
- ✓ Localized ownership uncertainty

Primary human attention is warranted where

- Ownership cannot be attributed
- Rationale is absent
- Dependency relationships remain incomplete

This package does

- ✓ Reconstructs repository reality
- ✓ Surfaces unknowns
- ✓ Prioritizes investigation
- ✓ Prepares human review

This package does not

- ✗ Execute code
- ✗ Access production
- ✗ Certify quality
- ✗ Approve architecture
- ✗ Make decisions

The Observatory

The Observatory is the **interactive interface to reconstructed repository reality**. This PDF is its executive projection; the Observatory itself is live, explorable, and drillable to the evidence beneath.



An explorable projection of repository reality

Each region is a reconstructed characteristic of the system. Colour is categorical reality, what kind of region this is, and chromatic prominence marks investigation interest, never severity. Amber marks the unknown.

Observatory ≠ Authority · the projection decides nothing.

What can be explored

- Topology
- Morphology
- Dependencies
- Historical layers
- Unknowns
- Investigation corridors
- Evidence traces
- Structural signals

The Observatory is the primary reality interface; the PDF is an executive projection of it.

What This Package Is, and Is Not

Stated before the evidence begins, so there is no misunderstanding.

This package is

- ✓ A repository reality reconstruction
- ✓ Structural intelligence
- ✓ Investigation preparation
- ✓ Software decision intelligence

This package is not

- ✗ A code review
- ✗ A quality audit
- ✗ A security certification
- ✗ A remediation plan
- ✗ A decision engine

Structural Signals

The most prominent observable characteristics of the reconstructed reality. Bands describe what is present, they do not judge it. Signal $\neq$ Score.

SIGNAL	OBSERVED	WHAT THE RECONSTRUCTED REALITY SHOWS
Authority Concentration	HIGH	Inbound dependency concentrates on the constitutional core; well-observed.
Ownership Certainty	LOW	Ownership is locally unattributable, among the highest investigation interest.
Historical Sediment	DENSE	Distinct legacy layers are visible; the order of accretion is unverifiable.
Legacy Continuity	MODERATE	Continuity mechanisms preserved deliberately; rationale recorded as unknown.
Boundary Clarity	MODERATE	Most region boundaries are explicit; two are inferred, not stated.
Unknown Density	DENSE	Knowledge thins around legacy custody and rationale, the densest fog.
Governance Visibility	HIGH	Gates and audit families are explicit in-snapshot; low residual interest.
Structural Coupling	HIGH	Coupling concentrates inward toward the core; surfaces orbit at lower coupling.

Band text describes the observed level. Band colour marks **investigation interest** (teal = look here first; amber = unknown), never severity, risk, or quality. Signals describe reality; they do not judge it.

Questions This Package Helps Humans Answer

The reconstructed reality is built so people can answer these. RepoTruth surfaces the evidence; humans answer.

Where should modernization investigation begin?

Which boundaries deserve earlier human review?

Where is knowledge concentrated?

Which regions remain poorly understood?

Where does inherited complexity concentrate?

Where should AI-adoption investigation begin?

Which historical regions deserve clarification?

Which unknowns would most improve understanding if resolved?

RepoTruth surfaces evidence. Humans answer.

Organizational Decisions This Package Supports

The reconstructed reality model prepares the decisions below. It does not make them.

AI adoption preparation

M&A technical due diligence

Platform evolution

Modernization preparation

Legacy investigation

Human knowledge preservation

Architectural transformation

Governance review

RepoTruth prepares decisions. RepoTruth does not make decisions.

Executive Interpretation

A high-level reading of what the evidence shows — and, with equal weight, what it does not. This is interpretation, not recommendation; it orders attention, it does not order action.

Reporting principle: we report what the evidence shows, including what is uncomfortable, and we point to where it most deserves attention. We do not soften it, and we never fabricate what we could not see.

What the evidence indicates

Repo X presents as a system shaped by **deliberate constitutional evolution**: an inner core of invariant-bearing logic, surrounded by surfaces that preserve compatibility with earlier behavior. The strongest, most directly observed structures cluster in the governance and intelligence regions; the thinnest knowledge concentrates around legacy continuity and the rationale behind preserved duplicates.

Investigation priorities (where to look first)

Earliest human attention is warranted at the **Mirror Facade** and **Legacy Compatibility** surfaces, where structure is observable but intent and ownership are not. Priority reflects *investigation urgency*, never severity.

Architectural themes

Boundary maturation, governance emergence, compatibility preservation, and intelligence layering recur across regions. These are described, not judged.

Evidence boundaries

Observation is bounded to a read-only snapshot. Runtime behavior, deployment topology, and live data flows are outside the evidence surface and are recorded as unknowns, not assumptions.

Confidence boundaries

Confidence is highest where structure is directly readable and lowest where it depends on history or absent rationale. No observation is asserted more firmly than its evidence allows.

The unknown landscape (read this section as evidence, not as gaps)

The most decision-relevant findings in this package are often the **preserved unknowns**: ownership that the snapshot cannot attribute, rationale never written down, and historical layering whose origin is unverifiable. An unknown here is a first-class output — it tells a reviewer precisely what to ask next.

Reads as: a calm map of a maturing, compatibility-preserving system.

Not as: a quality grade, a maturity score, a risk acceptance, or an instruction to act.

Repository Reality Atlas

The repository as it actually is — structure observed from the snapshot, not architecture as documented or intended.

Topology

A core-and-surface topology: a dense, highly-interconnected constitutional core, ringed by lower-coupling compatibility and presentation surfaces. Dependency concentrates inward; knowledge thins outward.

Major architectural regions

Six observed regions (mapped in Section 10): Constitutional Core, Legacy, Compatibility, Governance, Intelligence, Observatory. Boundaries between them are mostly explicit; two are inferred.

Bounded domains

Domain boundaries align with the governance and intelligence regions. The legacy region spans more than one domain, recorded as a boundary unknown rather than resolved.

Module families

Module families recur by responsibility: gate/audit families (governance), generation/analysis families (intelligence), and facade/mirror families (compatibility). Naming is consistent enough to be observed, not consistent enough to be relied upon.

Structural relationships

The core is consumed by every outer region; the Observatory consumes the Intelligence region read-only and originates nothing of its own. The Legacy region is the only region observed to be consumed by, and to consume, multiple others — the densest point of structural ambiguity in the snapshot.

Atlas reflects observed structure at snapshot time. Relationships not present in the snapshot are absent here by design, not by conclusion.

Architectural Region Map

Six observed regions, each with what was seen, how confident the observation is, and what remains unknown.

Constitutional Core DIRECTLY OBSERVED

Invariant-bearing logic and shared boundaries; highest inbound coupling. Structure directly readable. Unknown: which invariants are enforced at runtime vs. by convention.

Legacy Regions UNKNOWN-HEAVY

Older pathways and preserved duplicates retained for continuity. Present and reachable. Unknown: ownership, rationale for retention, and whether any path is dormant.

Compatibility Surfaces STRONGLY DERIVED

Mirror facades that hold an older contract stable while internals change. Pattern observed across several families. Unknown: the full set of consumers depending on each facade.

Governance Surfaces DIRECTLY OBSERVED


Operator gates and audit/ledger families. Decision points are explicit and routed through review. Unknown: completeness of audit coverage outside the snapshot.

Intelligence Surfaces DIRECTLY OBSERVED

Generation and analysis families that produce evidence and priorities. Deterministic, read-mostly shape. Unknown: external inputs not captured by the snapshot.

Observatory Surfaces STRONGLY DERIVED

A visualization layer over the Intelligence region. Originates no logic; renders evidence. Unknown: none material structurally; projection fidelity is review-bounded.



Key Evidence Catalog

Representative evidence. Each item is an observation with a reference, a confidence level, and a note on how the surrounding unknowns affect it. Evidence $\neq$ Truth.

REF	EVIDENCE SUMMARY	CONFIDENCE	UNKNOWN IMPACT
EV-001	Constitutional core is the most depended-upon region; all surfaces reference it.	DIRECTLY OBSERVED	Low — structure is explicit.
EV-014	Mirror facades hold a stable external contract over changing internals.	STRONGLY DERIVED	Moderate — consumer set not fully enumerable.
EV-022	Preserved duplicate pathways exist in the legacy region.	MODERATELY DERIVED	High — rationale and ownership unknown.
EV-031	Operator gates route state-changing actions through explicit review.	DIRECTLY OBSERVED	Low — gates present in structure.
EV-047	Observatory consumes intelligence outputs read-only; originates none.	STRONGLY DERIVED	Low — directionality observed.
EV-058	Legacy region participates in more than one bounded domain.	MODERATELY DERIVED	High — boundary intent unwritten.

Catalog is representative, not exhaustive. References are traceable to the full evidence set in the generated package.

Review Priority Register

Regions that warrant **earlier** human investigation. Order reflects *investigation urgency* — how much a reviewer's attention would change what is known. It never reflects severity, risk, or importance of the underlying code. Priority $\neq$ Decision.

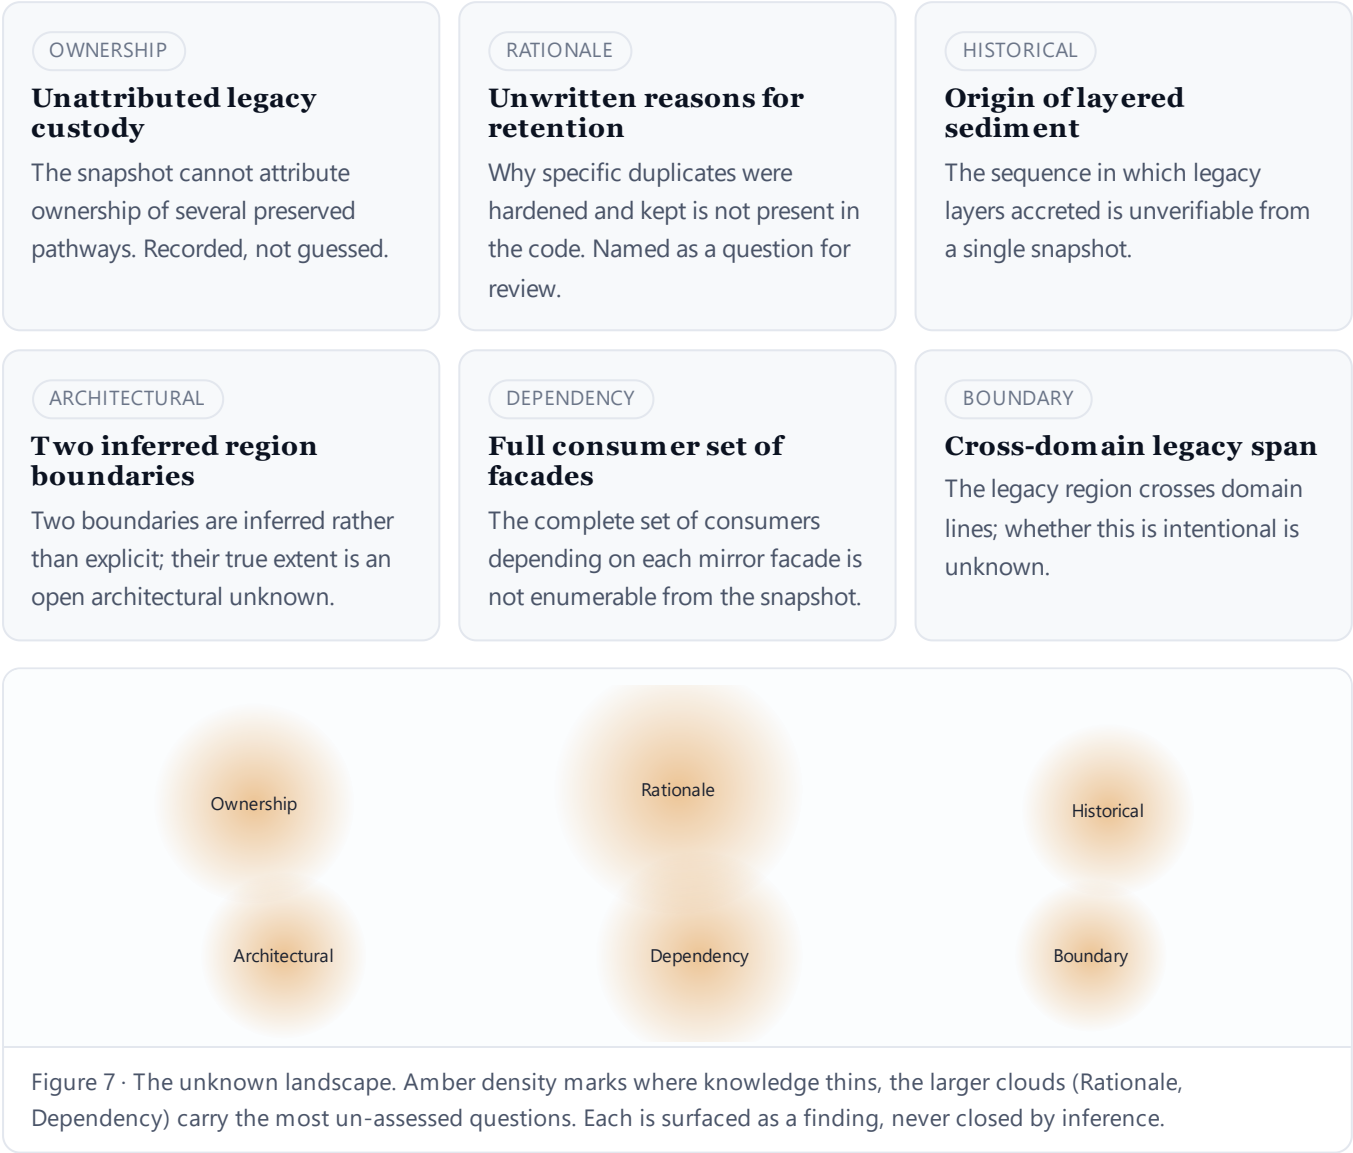
- P1** Mirror Facade surfaces — structure is clear, but the set of dependents is not; investigating early resolves the most downstream uncertainty.
- P2** Legacy continuity pathways — presence is observed, intent is not; review converts unknowns into decisions a human can own.
- P3** Cross-domain legacy boundary — a boundary the snapshot cannot attribute; earliest human reading sharpens the domain map.
- P4** Audit coverage edges — governance is explicit in-snapshot; coverage outside the snapshot is unverifiable and worth a deliberate look.



Why urgency, not severity: RepoTruth orders *where understanding is thinnest and most consequential to investigate*. It does not claim any region is good, bad, risky, or safe.

Unknown Inventory

Unknowns are first-class outputs, surfaced as prominently as any finding. An unknown is not a failure, not noise, and **not low risk** — it is an un-assessed question, and often the most useful page in the package. Unknown $\neq$ Low Risk.



Each unknown is preserved with its category and who could resolve it (evidence, or a human). None is closed by inference. This is the discipline RepoTruth exists to enforce.

Confidence Matrix

How firmly each major observation is held, and why. Confidence is an epistemic level, not a quality score. There is no black-box confidence — every level is explainable. Confidence $\neq$ Certainty.


- 01 DIRECTLY OBSERVED** Read from the snapshot with no inference — core structure, governance gates, intelligence shape.
- 02 STRONGLY DERIVED** From multiple consistent observations — mirror-facade pattern, observatory directionality.
- 03 MODERATELY DERIVED** From limited but consistent evidence — preserved duplicates, cross-domain legacy span.
- 04 WEAKLY DERIVED** From sparse or partial evidence, with gaps named.
- 05 HISTORICALLY INFERRED** From repository history where present and readable.
- 06 INDETERMINATE** Evidence supports more than one reading; surfaced as ambiguity, not resolved by guess.
- 07 UNVERIFIABLE** A question a static snapshot structurally cannot answer.
- 08 UNKNOWN** No evidence either way — recorded explicitly, never closed by inference.

No observation in this package is asserted above the confidence its evidence permits. Where the matrix says “Unknown,” the package says so plainly.

Concern Candidate Register

Candidates for human review — explicitly **not** confirmed problems. A Concern Candidate is a place the evidence suggests a human should look, paired with the confidence and the unknowns around it. It is never a defect, a verdict, or an instruction.

CC-01 • Facade dependent-set opacity Evidence: EV-014. Confidence: Strongly Derived. Unknown interaction: consumer set not enumerable. Human review: required to confirm whether opacity is intentional encapsulation or unmapped coupling.	CC-02 • Retained duplicate rationale Evidence: EV-022. Confidence: Moderately Derived. Unknown interaction: rationale + ownership both unknown. Human review: required; only a human owner can supply intent.
CC-03 • Cross-domain legacy boundary Evidence: EV-058. Confidence: Moderately Derived. Unknown interaction: boundary intent unwritten. Human review: required to decide whether the span is by design.	CC-04 • Audit coverage at snapshot edge Evidence: EV-031. Confidence: Directly Observed (in-snapshot). Unknown interaction: out-of-snapshot coverage unverifiable. Human review: required to confirm completeness beyond the snapshot.



Each candidate carries its own evidence, confidence, and unknowns, and each ends the same way: with a human. RepoTruth raises the candidate; it does not adjudicate it.

Investigation Corridors

Suggested **paths** for human investigation — a way of looking, never a way of acting. Corridors are for human review only; they prescribe nothing.

Legacy Compatibility Corridor

Follow the preserved pathways from the legacy region outward to each compatibility surface that depends on them. Purpose: convert continuity unknowns into owned decisions.

Constitutional Continuity Corridor

Trace the invariant-bearing core to the surfaces that consume it. Purpose: see where the constitution is enforced vs. assumed.

Governance Evolution Corridor

Walk the gate and audit families across regions. Purpose: read how decision points emerged and where their coverage ends.

Mirror Facade Corridor

Enter each mirror facade and map its true dependents. Purpose: resolve the highest-priority dependency opacity first.

A corridor is a lens, not a lever. It tells your people where to walk and what to ask — your people decide what, if anything, to do.

Constitutional Evolution Summary

The observable constitutional evolution of Repo X, described — not evaluated.

Boundary maturation

Region boundaries are observed to have become more explicit over the system's layers; two remain inferred. Described as a trajectory, not graded.

Governance emergence

Explicit gates and audit families are present where state changes. Their emergence is observable in structure; their sufficiency is a matter for human review.

Compatibility preservation

Mirror facades and continuity mechanisms recur, holding older contracts stable. A consistent pattern, neither endorsed nor faulted here.

Intelligence layering

A read-mostly intelligence region sits above the core, with an observatory layer above it that renders rather than originates.

No statement in this section asserts that any evolution is good, complete, correct, or finished. It records what is observable.

Compatibility Preservation Analysis

Where most tools see “duplication” or “dead code,” RepoTruth reads **intent under uncertainty**: mechanisms that preserve continuity while a system changes beneath them. This is a RepoTruth differentiator — naming a pattern instead of flagging a defect.

Mirror facades

Surfaces that mirror an older interface while internals evolve. Observed across several module families; their dependent sets are the primary unknown.

Compatibility surfaces

Boundaries that hold an external contract stable. Present and reachable; completeness of coverage is review-bounded.

Continuity mechanisms

Pathways retained specifically to keep prior behavior available. Their rationale is an explicit unknown, not an assumed redundancy.

Legacy preservation patterns

Hardened duplicates retained for continuity rather than removed. RepoTruth records the pattern and the open questions; it does not label them waste.

A scanner would mark these for deletion. An intelligence layer asks first: *who depends on this, and why was it kept?* — and leaves the answer to a human.

Legacy Sediment Analysis

Legacy as **sediment** — layers laid down over time, preserved rather than erased. Surfaced as evidence and unknowns, never as debt to be judged.

Legacy regions

Older pathways remain present and reachable. Their boundaries are partly inferred (see Section 13).

Preserved pathways

Specific routes retained for continuity; observed, with ownership and rationale recorded as unknown.

Historical layering

Distinct layers are visible in structure; their order of accretion is unverifiable from one snapshot.

Continuity evidence

The compatibility surfaces (Section 18) are the observable evidence that these layers are kept deliberately, not by neglect — a reading, offered for confirmation.

Human Review Roadmap


What humans may wish to investigate next — **investigation candidates only**. This roadmap prescribes no actions and accepts no risk on anyone's behalf.

- Map the dependent set of each mirror facade (resolves CC-01, P1).
- Source the rationale and owner for retained duplicates (resolves CC-02, P2).
- Confirm whether the cross-domain legacy boundary is intentional (resolves CC-03, P3).
- Verify audit coverage beyond the snapshot edge (resolves CC-04, P4).

Each line is a question to investigate, paired with the candidate and priority it would resolve. None is a directive. What to do — and whether to do anything — is a human decision.

Observatory Projection Gallery

Example projections from the Observatory — the **Evidence Visualization Layer**. These render the evidence behind the package; they assign no grades and decide nothing. The Observatory shows; your people interpret.



Dependency Galaxy

Inbound concentration on the constitutional core; surfaces orbit at lower coupling. The densest node is the most depended-upon region (EV-001).

Confidence: Directly Observed structure; coupling counts read from the snapshot.



Unknown Fog

Where knowledge thins — densest over legacy continuity and retained-duplicate rationale. Fog marks questions, not faults.

Confidence: amber density maps Unknown / Unverifiable observations; it implies no risk level.



Investigation Corridors

Suggested human pathways through the system (Section 16), rendered as routes, not actions. Endpoints mark where a corridor opens onto an open question.


Confidence: corridors derive from observed structure; the destination remains a human decision.



Confidence Landscape

Distribution of observations by confidence level — tallest at Directly Observed, tapering through derived levels, with an explicit Unknown band (amber) on the right.

Confidence: this is a distribution of epistemic levels, not a quality or risk score.



Architectural Topology

Core-and-surface layout: the constitutional core mediates between legacy/compatibility (left) and governance/intelligence/observatory (right) regions.

Confidence: Directly Observed core; two outer boundaries are inferred (Section 13).

Projections are illustrative of the live Observatory. In a generated package they are interactive and drillable to the evidence beneath. Observatory ≠ Authority.

Constitutional Manifest

Every package declares the boundaries it was generated under. Trust is easier to extend when the limits are written down.

Constitutions applied

RT-DIO (ontology) · RT-SC (scores = urgency) · RT-CC (confidence = epistemic) · RT-UPS (unknown preservation) · RT-HRW (human-terminal review) · RT-ERM (report morphology) · RT-DPA (package architecture) · RT-PRC (production boundaries).
Conformance validated by RT-CE.

Unsupported surfaces

Runtime behavior, live data flows, deployment topology, and any out-of-snapshot system are outside the evidence surface and recorded as unknowns, not inferred.

Confidence boundary

No observation is asserted above its evidence. Eight explainable levels, ending in an explicit Unknown.

Unknown boundary

Unknowns are preserved as first-class outputs and never closed by inference. Unknown $\neq$ Low Risk.

Review boundary

Every output — evidence, priority, candidate, corridor — terminates at human review. RepoTruth is non-terminal by design.

Refusal duty

Where evidence is insufficient, the package says so rather than fabricating a conclusion. Absence is reported, not filled.

Final Boundary Statement

RepoTruth provides Software Decision Intelligence.

RepoTruth does not decide.

Human review remains mandatory.

No recommendation contained in this package should be inferred.

This is a synthetic demonstration generated to demonstrate the RepoTruth Software Decision Intelligence model. No customer repository. No customer conclusions. No customer data. Evidence, confidence, priorities, and preserved unknowns are offered for human review — the only place a decision is made.

Attestation & Signatures

This package closes with two marks: RepoTruth's attestation of *method*, and the space for the human review that is its only valid conclusion.



Generated by RepoTruth · Software Decision Intelligence

Artifact	RT-SDP-REPOX-0001 · Diagnostic Package (Package ≠ Verdict)
Generated	Synthetic demonstration · 2026 · Methodology v1
Constitutions	RT-DIO · RT-SC · RT-CC · RT-UPS · RT-HRW · RT-ERM · RT-DPA · RT-PRC
Conformance	Validated by RT-CE (Tier-0, unwaivable)

RepoTruth

Generation Engine · attestation of method and provenance

This attests that the package was produced under the named constitutional standards and passed Conformance Engine validation. It is an attestation of **method and provenance**, not an approval, certification, sign-off, or verdict. RepoTruth signs how the package was made; it does not sign what should be decided.

Human review · the decision belongs here

Human review is terminal. The conclusion drawn from this package is recorded by the people who sign below, never by RepoTruth. Priority ≠ Decision; Interpretation ≠ Recommendation.

_____	_____	_____
REVIEWED BY	ROLE	DATE
_____	_____	_____
REVIEWED BY	ROLE	DATE