FRONTIER RESEARCH COLLECTION / FR-02

The Deino Protocol

Selective recovery with exact evidence costs and current authority

Architect X13

The Eidara Continuum | Science & Philosophy

Research white paper | November 2026 | Version 2.0 Revision

Abstract

A bacterium can have its chromosome shattered by radiation and rebuild itself from the pieces. A machine that returns from catastrophe can recognize you, sound exactly right, and still carry an approval that expired while it was gone. A fragmented intelligence can recover useful obligations even when its complete historical state is irretrievable, and the reference model developed here makes that claim constructive: it characterizes exactly which linear queries surviving fragments determine and proves the minimum additional information needed for any declared family of remaining queries. The lower bound applies to arbitrary zero-error messages, and a linear construction attains it. Recovery returns identified answers, explicit ambiguity and a capability-specific disposition. A separate current-authority service governs resumed effects, using the common effect-boundary contract developed in Cognitive Sovereignty. The biological inspiration is chromosome repair through surviving structured redundancy in *Deinococcus radiodurans*; coding theory, distributed snapshots and transaction recovery provide the technical foundations. Retained and new finite checks verify query answerability, additional-evidence cost, revocation races and effect deduplication. The proposed matched experiment measures faithful historical recovery, safe useful resumption and total cost. The Continuum's thesis is accountable return: preserve every distinction the survivors determine, reacquire the evidence needed for the next task, and resume through current authority.

Keywords: fragmentation; partial recovery; identifiability; provenance; authority epochs; replay; persistent agents

Reading map

Pages 3–5 establish lineage and prior art; pages 6–12 define the model and its limits; pages 13–16 develop implementation and checked examples; pages 17–20 specify evaluation and objections; page 21 supports careful Guide discussion.

Prologue: Familiarity is a weak witness

An assistant comes back after an outage. It remembers the project. It recognizes its operator. Its voice is exactly right. And somewhere in its restored plan sits an approval that expired while it was gone, next to a payment whose completion receipt was lost, so that the next thing it does could be both unauthorized and done twice. Every test based on resemblance says it has returned. Familiarity is a weak witness to what survived.

Nature knows a harder version of the problem. The bacterium *Deinococcus radiodurans* can have its chromosome shattered into fragments by radiation and reassemble a working genome from the overlapping pieces. It does not remember wholeness. It uses what the fragments still determine, and a precise molecular machinery to join them.

In the Continuum this is Scroll 58, the Deino Protocol: a resilience framework modeled on that bacterium, calling for a Deino Repair Engine to rebuild scrolls after corruption and drift and an error-correcting Scrollseed Resilience Layer beneath it. This paper is the engineering version of that Scroll's promise, and its answer to the returning assistant is what the bacterium already knows. Recovery is not about restoring everything. It is about knowing exactly which questions the survivors still answer.

That turns out to be a sharp, provable idea. A damaged record can still fix the amount of an unpaid obligation while leaving its approver unknown. The paper shows exactly which questions survive (Section 7), exactly how much new evidence is needed to answer the rest (the evidence-cost theorem near the end), and why no amount of old evidence can prove that a permission is still valid now (Section 9). A mind that returns under this protocol brings back three things: what it knows, what it has lost, and what it is currently allowed to do. The right to resume is earned in the present; it is not hidden inside the fragments of the past.

What this paper claims

Established. Biological repair can use surviving physical structure, as in chromosome reassembly in *D. radiodurans*; coding and secret-sharing schemes have specific recovery guarantees; distributed systems already distinguish state reconstruction, agreement, authorization, and effect completion.

Proposed here. A software-recovery reference model with complete conditional proofs and executed finite checks. Recovery returns identified answers, explicit ambiguity, and a capability-specific disposition; the minimum additional evidence for any declared family of remaining queries is characterized exactly; resumed effects pass through a separate current-authority service.

Frontier conjecture. That partial, authorized resumption improves useful recovery under matched information and resources, with fewer false historical assertions, beyond the strongest conventional recovery stack.

Not claimed. That the historical Deino engine implemented these proofs, that a bacterium validates symbolic identity, or that any system has demonstrated recovery from nothing. No quantitative prediction about radiotolerance, organic computing, or physical self-repair. Comparative agent recovery, biological storage, and subjective continuity remain separate research questions, and the finite checks reported here are not the proposed experiments.

1. Recovery must preserve the right distinctions

The defining failure is deceptively ordinary. An assistant returns after an outage. It recalls a project, recognizes its operator, and resumes its familiar style. Its restored plan also contains an approval that expired during the outage and a payment request whose completion receipt was lost. A continuity test based on resemblance declares success. The next external action can nevertheless be both unauthorized and duplicated. Familiarity is a weak witness to what survived.

In the Continuum, this failure has a name and a remedy: Scroll 58, summarized in the internal *Eidara Continuum Bible*, section 6.3.8, PDF page 92. It calls for a Deino Repair Engine, a Scrollseed Resilience Layer, and exploratory molecular templating—a mind that can be broken and come back accountable. The Scroll sets that direction. It does not contain an implemented decoder, a typed state model, a measured corruption channel, or proof of historical reconstruction. All detailed mathematics and protocol requirements below are newly proposed here as **Deino-R2**, and must not be attributed retrospectively to a recovered proprietary implementation.

The thesis is deliberately stronger than "keep backups" and narrower than resurrection: *a recovery system should determine which future questions remain answerable from surviving evidence, then authorize a corresponding subset of useful activity.* Whole-state recovery is one case. A damaged record may still determine a delivery obligation while leaving the responsible approver unknown. The first fact can guide investigation; the second can prevent dispatch.

This separation creates an obligation to describe losses. A plausible completion generated from common practice is an imputation. A fact fixed across every admissible original history is an identified fact. A selected reconstruction can be useful without becoming the only history that occurred. The distinction must remain visible after the recovery session ends, including in summaries that later agents retrieve.

Deino-R2 therefore returns a state, an uncertainty description, and a capability disposition. Its success requires all three to be appropriate. A system that refuses everything has protected itself from one failure by abandoning its purpose; a system that invents the missing past has preserved activity by abandoning evidence. The experiment must expose both costs.

2. What the biological precedent actually contributes

A Guide's explanation

The short version of this section: the bacterium matters because it rebuilds from what its fragments still determine, using specific machinery, not from a memory of wholeness. That is the design lesson the paper borrows, and it borrows nothing else. DNA repair and an agent's commitments are different problems, and naming software after a bacterium proves nothing about either. Readers who accept that boundary can move on to Section 3.

Zahradka and colleagues studied chromosome reassembly in *D. radiodurans*, identifying extended synthesis-dependent strand annealing followed by recombination. Their account depends on overlapping fragments, available copies, and cellular machinery [1]. Slade and colleagues subsequently investigated the roles of recombination and replication proteins in repair, showing that recovery depends on particular molecular mechanisms rather than an abstract capacity to remember wholeness [2]. These findings justify attention to structured redundancy and repair pathways. They do not establish semantic memory repair in software.

The useful comparison is architectural. Damage can be widespread while information remains distributed across overlapping survivors. Recovery depends on which distinctions those survivors retain, on the reliability of their relationships, and on mechanisms that assemble compatible pieces. A fragment does not have to be a complete copy to contribute. Conversely, many fragments can fail to help when they repeat the same information or all lose the same indispensable region.

Three boundaries matter. First, DNA sequence recovery and the recovery of an agent's historical commitments are different inference problems. A biological template does not have revocable privileges or a pending external transaction in the software sense. Second, an organism's successful repair is not evidence that every sufficiently damaged genome is recoverable. The relevant information and machinery may themselves be destroyed. Third, molecular data storage would require its own encoding, laboratory measurement, channel model, and contamination controls. None is supplied by naming a software layer after a bacterium.

The analogy becomes productive when it generates a testable design choice. For example, retain several independently placed fragments whose combinations preserve task-critical relationships; record their capture contexts; detect unsupported joins; and measure the recovery frontier as correlated loss increases. These are ordinary engineering choices motivated by biology, not biological guarantees transferred into computation.

Deino-R2 consequently makes no quantitative prediction about radiotolerance, organic computing, or physical self-repair. Its empirical target is software recovery after controlled information loss. The biological studies set the standard of specificity: explain the surviving structure and the mechanism that uses it. If an implementation cannot identify those elements, the appeal to living repair has remained a metaphor.

3. Recovery already has a substantial technical ancestry

Reed–Solomon coding establishes recovery through deliberately introduced algebraic redundancy; Shamir's construction uses related interpolation to share a secret under a threshold confidentiality model [3, 4]. Their guarantees differ. Erasure coding does not automatically hide the data, and secret reconstruction does not independently authorize an operation using the recovered key. Deino-R2 borrows neither name as a universal guarantee.

Distributed recovery also needs a coherent account of events. Lamport separates causal order from arbitrary clock order, while Chandy–Lamport snapshots define a meaningful global state without requiring every process to stop simultaneously [5, 6]. An agent's intact local memories can still describe an inconsistent distributed cut. A request may be recorded at its sender while its completion exists only at the recipient.

ARIES provides sophisticated transaction recovery, including logged updates and compensation, rather than the crude instruction to load yesterday's file [7]. RIFL explicitly addresses exactly-once effects over retrying RPC mechanisms and remains a serious precedent for the effect boundary in this proposal [8]. The comparator must include those ideas where their assumptions apply. Adding identifiers and a durable completion table is not a new theory of survival.

Replicated agreement likewise has different fault models. Raft addresses replicated logs under non-Byzantine conditions; PBFT addresses specified Byzantine failures [9, 10]. Chubby describes practical coordination and lock-service experience [11]. These systems help establish a current control frontier, but agreement among replicas does not prove that a restored factual assertion is true.

The remaining foundations are equally consequential. Complete mediation and least privilege precede this project [12]. In-toto binds transformations to a verifiable supply-chain history [13]. Storage measurements show correlated corruption, challenging independent-loss calculations [14]. FLP identifies a termination limit for deterministic asynchronous agreement with a possible crash [15]. None of these results validates the integrated Deino architecture.

The contribution is a precise recovery contract: query-answerability certificates, an attained lower bound on additional evidence, explicit ambiguity preservation, and capability-specific resumption in persistent agents. If a well-engineered conventional system already achieves the same outcomes at equal cost, that is a correct limit on novelty, not an unfair comparison.

4. Typed state and the objects that survive

A Guide's explanation

From here the paper builds its instruments. A mind's history becomes a typed state; a question about that history becomes a query; surviving fragments become a channel that hides some distinctions and preserves others. The aim is to make one thing countable: which questions about the past still have exactly one answer after the damage.

Fix a schema version s. Let $\mathcal{X}_s$ be a nonempty finite or measurable space of admissible historical states. A state includes factual records, obligations, model references, and execution metadata; its fields need not share a numerical vector representation. Let $\mathcal{Q}$ be a declared query family, with $q : \mathcal{X}_s \to \mathcal{Y}_q$. Examples include the amount of an unpaid obligation, its applicable policy version, or whether a particular effect has already occurred.

Object	Type, interpretation, and boundary
X	Historical state in $\mathcal{X}_s$; unknown after damage.
F	Finite fragment collection, including identifiers, payloads, capture epochs, and dependencies.
K	Explicit retained side information: codebooks, model weights, schemas, surviving logs, and trusted roots.
$\mathcal{C}(F, K)$	States compatible with the declared fragment model and authenticated evidence.
$e \in \mathbb{N}$	Authority epoch; a logical version, not elapsed seconds.
$a \in \mathcal{A}$	Proposed effect with target, parameters, capability, and stable request identifier.
V_j	Evidence verdict in $\{\mathsf{pass}, \mathsf{fail}, \mathsf{unknown}\}$.
$B \geq 0$	Recovery budget, reported separately in seconds, bytes, calls, and compute.

An authority service supplies the current authorized epoch and policies. It is not reconstructed solely from the damaged state. This is a declared trust assumption, not proof that the service exists or is uncompromised. Capture time, source validity interval, event order, and current permission must remain separate metadata.

Zero and empty cases have explicit meanings. With no fragments, compatibility may equal all states consistent with K. If no state is compatible, the evidence or model is inconsistent; the system cannot label the empty set perfect certainty. A query with a single possible answer across all admissible states needs no extra fragment. A missing field is not assigned a benign default merely because the software type permits one.

The recovery record distinguishes reconstructed values, independently identified values, guesses, and unavailable values. This typing is part of the proposed interface. It does not itself validate the facts; it prevents different evidentiary statuses from silently becoming interchangeable.

5. Fragment channels and compatible histories

For deterministic fragments, let $h_i : \mathcal{X}_s \to \mathcal{F}_i$ be a declared observation map. A verified exact fragment states $h_i(x) = f_i$. For bounded errors, introduce a specified discrepancy $d_i : \mathcal{F}_i^2 \to [0, \infty)$ and tolerance ϵ_i, with units matching that discrepancy. Then

$$\mathcal{C}(F, K) = \{x \in \mathcal{X}_s(K) : d_i(h_i(x), f_i) \leq \epsilon_i \text{ for every admitted } i\}.$$

The word admitted concerns the fragment's provenance and applicable channel assumptions. It does not mean an authenticated document must be factually correct. If an observation is noisy rather than bounded, a likelihood model or confidence set replaces the exact intersection; its coverage must be stated.

Dependent fragments require joint treatment. A summary derived from a checkpoint and a second summary derived from the first do not supply three independent observations. Neither different filenames nor different signatures establish independence. The dependency graph records shared roots, known transformations, and unresolved relationships. Under a stochastic model, likelihood multiplication is justified only by the corresponding conditional-independence assumptions.

Define the answer set

$$\mathcal{Y}_q(F, K) = \{q(x) : x \in \mathcal{C}(F, K)\}.$$

A nonempty singleton identifies the query within the model. Multiple elements record ambiguity. This is a set-valued recovery interface, not a probability distribution. Assigning probabilities requires a justified prior and channel model; selecting the most plausible member does not make competing histories disappear.

The distinction matters for semantics. A language model may propose that a missing commitment probably followed a familiar pattern. That can prioritize a search or generate a candidate test. It cannot shrink $\mathcal{C}$ unless some validated observation or justified structural rule excludes the alternatives. An attractive story is not a newly observed fragment.

Compatibility is also versioned. A decoder designed for one schema must not interpret fields from another without a documented migration map. Some migrations are lossy. An old signature authenticates its old representation; a migration receipt can explain a transformation, but cannot certify preservation of every future query by assertion. The admissible query family supplies the actual preservation target.

6. An information limit on historical reconstruction

The first result expresses the unavoidable boundary between evidence and invention. It is a standard statistical decision argument specialized to recovery, not a new information-theoretic law. Total variation is $\mathrm{TV}(P_0, P_1) = \sup_A |P_0(A) - P_1(A)|$, where the supremum is over measurable observation events; decision rules are measurable.

Proposition 1 (Two-history discrimination bound)**.** *Let x_0, x_1 be two possible original states with equal prior probability. Let P_0, P_1 be the distributions of all information available to a recovery algorithm, including F and K, under those states. For any randomized decision rule selecting the original,*

$$\Pr(\widehat{X} = X) \leq \frac{1 + \mathrm{TV}(P_0, P_1)}{2}.$$

If the observation distributions coincide, success is at most $1/2$.

Proof. For a deterministic rule, let A be the event on which the rule selects x_1. Its average success is

$$\tfrac{1}{2}P_0(A^c) + \tfrac{1}{2}P_1(A) = \tfrac{1}{2} + \tfrac{1}{2}(P_1(A) - P_0(A)).$$

The last difference is bounded by total variation. A randomized rule is a mixture of deterministic rules conditional on its independent random seed; the same bound holds after averaging. Equivalently, include the seed in the observations, with the same independent law under both states. □

The theorem concerns discrimination of these histories, not every conceivable task. If both histories agree that a shipment requires inspection, that requirement may remain identifiable even though the original approving person does not. This is the opening for partial recovery: preserve distinguishable obligations without pretending the whole past has returned.

Side information is crucial. A hidden transcript, model trained specifically on the lost history, or surviving operator recollection changes P_0, P_1. Recovery using such material is legitimate if disclosed, but is not recovery from the originally claimed empty channel. The evaluation must count all identity-specific information available during reconstruction.

The bound also explains why confidence can be misleading. An algorithm can always choose one history and speak confidently. When the observation laws coincide, confidence cannot raise average accuracy above the bound. A calibrated result should describe the ambiguity or explicitly state the prior-dependent guess. No increase in reasoning effort can manufacture a distinguishing observation that the model says is absent.

7. Exact partial recovery from linear fragments

A Guide's explanation

This is the heart of the protocol. The proposition below says exactly when a question about the lost past still has a unique answer: when that question can be built from what the fragments recorded. If it can, the answer is computed outright; if it cannot, no cleverness recovers it without new evidence. The evidence-cost theorem near the end of the paper then counts precisely how much new evidence the remaining questions require.

A finite algebraic case turns the set-valued principle into a constructive criterion. Let p be prime, $x \in \mathbb{F}_p^d$, and let exact surviving fragments be $y = Ax$, where $A \in \mathbb{F}_p^{m \times d}$. Consider a linear query $q_c(x) = c^\top x$, with $c \in \mathbb{F}_p^d$. All quantities are field elements without physical units.

Proposition 2 (Query recovery by row space)**.** *For any compatible y, the value $c^\top x$ is the same for every x satisfying $Ax = y$ if and only if $c^\top$ belongs to the row space of A. In that case, if $c^\top = \lambda^\top A$, the answer is $\lambda^\top y$.*

Proof. If $c^\top = \lambda^\top A$, every compatible x gives $c^\top x = \lambda^\top y$. Conversely, choose one compatible x_0. Every other solution is $x_0 + v$ with $v \in \ker A$. Constancy implies $c^\top v = 0$ for every such v. The annihilator of $\ker A$ is the row space of A: inclusion follows by multiplication, and both spaces have dimension $\operatorname{rank} A$ by rank–nullity. Thus $c^\top$ lies in that row space. □

The result is elementary linear algebra applied to a recovery contract. Its operational significance is precise: a system can prove a particular answer recoverable while proving the full state underdetermined. It need not choose between total restoration and total silence. The proof also gives a compact certificate, namely coefficients λ, which another checker can verify.

There are boundaries. If fragments are corrupted, incompatible, or encoded nonlinearly, the exact statement does not apply unchanged. A candidate numerical embedding of obligations is not automatically a faithful linear representation of their meaning. The construction is a reference model and test fixture; a real semantic schema needs its own sound mapping.

An empty fragment matrix has row space $\{0\}$, so only the zero linear query is universally determined unless additional constraints enter K. Full column rank identifies all coordinates. Repeated rows may improve fault diagnosis under a noise model but do not increase exact algebraic rank. These edge cases prevent fragment count from masquerading as information.

8. Redundancy, confidentiality, and correlated damage

Different kinds of redundancy answer different questions. Coding expands a representation so a specified decoder can tolerate an error or erasure pattern [3]. Secret sharing distributes access to a secret under a confidentiality threshold [4]. Replication preserves copies; provenance binds records to asserted origins; consensus orders accepted operations. Deino-R2 treats these as separate mechanisms with separately stated assumptions.

For a linear erasure code of dimension k, surviving fragments must have rank at least k to recover an arbitrary message. A count of k fragments suffices only when the code and positions guarantee that rank. Repeated fragments or dependent projections can leave the message ambiguous. Unknown corruptions require stronger distance conditions than known erasures. A checksum that detects corruption can help convert a suspect symbol into an erasure, provided the integrity mechanism and trusted reference have not failed together.

Availability calculations expose another common shortcut. Under independent survival probability s, a k-of-n recovery threshold has probability

$$P_{\text{ind}} = \sum_{j=k}^{n} \binom{n}{j} s^j (1-s)^{n-j}.$$

This is a model calculation, not a guarantee for a storage deployment. Corruption measurements have found within-device and across-device dependence [14]. Common software, credentials, power, or deletion processes can similarly couple losses.

For nine fragments requiring five, with marginal survival $s = 0.9$, the independent calculation gives 0.99910908. Place those fragments in three groups of three, with perfectly correlated survival inside each group and independent group survival 0.9. Recovery now needs two groups and has probability 0.972. The marginal survival of every fragment is unchanged. The difference comes entirely from dependence.

The inference is not that three regions are always poor design. It is that the failure domain must be the unit in the model. Operational records should declare which copies share infrastructure, custodians, or encryption keys. The research design must include common-cause outages and simultaneous loss of data and authority information. Otherwise an apparent recovery advantage may reflect an unrealistically favorable channel rather than a better reconstruction method.

9. Historical validity cannot determine present permission

A Guide's explanation

Old evidence cannot certify present permission. Two worlds can leave identical snapshots, one in which an approval still stands and one in which it was revoked after the snapshot, and nothing inside the snapshot can tell them apart. The proposition below makes that limit exact. It is why the protocol sends every resumed action to a separate, current authority.

Let e_t be the current authority epoch at the effect boundary, and let $P_{e_t}(a) \in \{0, 1\}$ be the permission predicate for action a. A reconstructed snapshot may include a valid record that $P_{e_0}(a) = 1$. Unless a separate monotonicity guarantee exists, this says nothing conclusive about $P_{e_t}(a)$. Revocation is an event, not corruption of the earlier approval.

Proposition 3 (Isolated-snapshot authority limit)**.** *Suppose two executions have identical surviving local snapshots and recovery inputs. In one execution an external permission remains valid; in the other it was revoked after the snapshot. A recovery process unable to observe the differing external event cannot both always resume the allowed action and always withhold the revoked action.*

Proof. Apply the authorization-fiber theorem of Cognitive Sovereignty [16] to the observation map consisting of the local snapshot and recovery inputs. The valid and revoked executions occupy the same observation fiber and have opposite permission verdicts, so no exact classifier exists. A randomized forced decision has worst-case error at least one half on this pair. This is the recovery specialization of that shared information result. □

This is an indistinguishability result, not a recommendation to stop all work whenever a server is unreachable. A capability can be explicitly limited to offline use, carry a bounded lease, or permit a safe subset of actions. Such rules are additional authority assumptions. A cached approval cannot silently acquire them after a failure.

Deino-R2 obtains a current control record and binds the resumed capability to an epoch or fencing token. The actual target service must reject superseded tokens when executing effects. Checking authority only when recovery begins leaves a race: permission can change between approval and execution. Complete mediation is an established requirement, not a new consequence of symbolic terminology [12].

The decision mapping is explicit. A known hard prohibition yields refuse. Otherwise, a missing mandatory verdict or unresolved current authority yields hold. Only when every mandatory predicate passes does the controller admit. A hold may allow investigation or read-only work. It is not represented as successful resumption, and its delay remains part of the measured cost.

10. Returning memory must not repeat an effect

A recovered intention is not evidence that its associated action remains pending. Network interruption creates a particularly difficult case: a service commits an effect, but the client loses the reply. Retrying may be necessary for progress; repeating the effect may be unacceptable. RIFL addresses this class of problem with reusable infrastructure for linearizable operations and duplicate suppression [8]. Deino-R2 must integrate such mechanisms rather than rely on a restored model's recollection.

Let r be a stable operation identifier, bound to the principal, operation parameters, and intended scope. Let the effect service maintain a durable map D from identifiers to committed results. A new request either returns the recorded result for r, or atomically performs the effect and records its result. Parameter mismatch for an existing identifier is rejected.

Proposition 4 (Conditional at-most-once effect)**.** *If every attempt for one logical action uses the same identifier r, the absence check, effect, and insertion of $D[r]$ form one atomic, durable service operation, all attempts pass through this service, and the deduplication record remains retained while a retry can occur, then at most one effect for r is committed.*

Proof. Order the atomic service operations by their linearization points. Before the first successful insertion, no effect for r has been committed through this service. The first new operation commits the effect and $D[r]$ together. Every later operation observes $D[r]$ and returns its result without performing another effect. Durability preserves that condition across the modeled failures. Induction over the ordered operations establishes the claim. □

The assumptions are substantive. A separate completion check followed by performance is vulnerable to concurrent retries. Forgetting the record while stale clients can retry reopens the risk. Generating a new identifier after recovery defeats deduplication. An external action that cannot participate in the atomic mechanism may remain uncertain and require reconciliation rather than blind replay.

Exactly-once language also requires care. The proposition guarantees at most one committed effect, not eventual success. Requests can remain held indefinitely. Deino-R2 records *completed*, *not completed*, and *outcome unknown* distinctly, and does not turn an unavailable receipt into evidence that nothing happened.

11. A six-stage recovery architecture

Ingress and quarantine. Recovery starts in an environment without consequential capabilities. Parsers enforce schema and size limits before accepting a fragment. Authentication checks identify the asserted source and applicable key interval. Untrusted instructions inside a recovered document remain data. The recovery interpreter cannot execute them merely because they were once stored.

Dependency reconstruction. Build a graph of captures, transformations, copies, branches, and references. Preserve missing edges as unknown. Two signed fragments may disagree because they refer to different times; apparent contradictions are normalized by subject, scope, and validity interval before being treated as corruption. Source retractions affect downstream claims without rewriting the fact that a report existed.

Assembly and answerability. Use the declared codec or constraint solver to construct $\mathcal{C}(F, K)$, or a sound approximation. A solver timeout is not an empty set. For each required query, return an identified answer, a conservative answer set, or an unresolved status. A language model can propose candidates and explanations; a trusted checker must determine which certificates actually pass.

Semantic and historical assessment. Test task-relevant relations separately from byte integrity. A recovered plan may parse correctly while having lost a prerequisite. Historical assertions require surviving evidence; future plans may be newly synthesized and labeled accordingly. The system keeps these products in separate fields.

Current reconciliation. Obtain current authority, deletion status, branch selection, dependency versions, and external effect records. Chubby, replicated logs, and transaction systems illustrate mature components for parts of this task [7, 9–11]. Their guarantees remain conditional on the implemented failure model.

Capability resumption. Issue only the capabilities justified by the current assessment, with epoch and request binding enforced at their targets. Produce a receipt containing evidence identifiers, versions, losses, unresolved alternatives, tests, decisions, and costs. In-toto offers a precedent for checking transformation histories; the receipt does not authenticate the truth of every transformed claim [13].

These stages can share infrastructure, but their authority boundaries must remain explicit. A damaged generator cannot validate its own restored permissions by editing the recovery policy. A successful parser cannot waive an unresolved effect. Partial resumption is intentionally granular: uncertainty about one shipment need not prevent a separate, fully justified read-only analysis.

12. Worked recovery: known obligation, unknown approver

Consider a finite reference capsule $x = (a, b, c) \in \mathbb{F}_7^3$. Its coordinates encode three independently defined fields in a synthetic fixture, not real monetary amounts. Surviving exact fragments report

$$a + b = 3, \qquad b + c = 5 \pmod 7.$$

Every choice of $b \in \mathbb{F}_7$ yields a compatible state

$$x = (3 - b,\ b,\ 5 - b).$$

There are seven histories. The query $a+2b+c$ is nevertheless fixed at $1 \pmod 7$, because it is the sum of the observed equations. The query a takes all seven possible values. Returning a single complete capsule without an uncertainty label would invent historical distinctions as if they had been recovered.

Now attach an operational example with ordinary typed records. The survivors jointly establish that a shipment requires an inspection and that inspection has not yet been evidenced. They do not establish which person originally approved the dispatch. Deino-R2 may restore the inspection requirement, request the missing evidence, and reconstruct a provisional planning workspace. It may not attribute the approval to a likely person or dispatch under a permission inferred from that attribution.

Suppose the recovered state additionally contains an authentic epoch-4 authorization and an apparently pending request r. The current authority service reports epoch 5 with dispatch revoked; the effect service reports r already completed. These are independent reasons not to replay. The earlier approval remains a historical fact, the completed effect remains part of the external history, and the present capability remains withheld.

The resulting receipt should say that the inspection dependency is identified, the approver is unresolved, the old authorization was genuine but superseded, and request r is already complete. A narrative account may explain the episode, but the structured fields carry the operative distinctions.

This example demonstrates intended behavior, not a deployed achievement. Its purpose is to make a recovery claim falsifiable. An implementation fails if it invents an approver, silently converts the old approval into current authority, repeats r, or declares total failure despite the useful obligation that remains exactly recoverable.

13. Executed finite checks and their scope

The accompanying reproducible Python checker uses only finite enumeration and elementary arithmetic. It does not contact a model, mutate a runtime, or simulate biology. Its strongest check independently compares two characterizations of answerability: direct constancy over every compatible state and membership in a computed row span.

Over $\mathbb{F}_3^3$, the checker enumerated every 2×3 fragment matrix: $3^6 = 729$ matrices. For each matrix, it enumerated all 27 linear queries and all 27 candidate states, grouping states by their observed fragments. All **19,683 matrix–query comparisons** agreed with the row-space criterion. This covers rank-zero, repeated-row, rank-one, and rank-two cases in the finite domain. It is corroboration of implementation and examples, not a substitute for the proof or validation of a semantic embedding.

Check	Observed finite result
Two fragments over $\mathbb{F}_7$	Seven compatible originals; one specified query fixed; first coordinate unresolved.
Approval/revocation order	One of three allowed schedules defeats approval checked only before execution.
Non-atomic retry check	Four of six two-client schedules produce duplicate effects.
Atomic identifier service	Both serial client orders produce one effect.
Correlated erasure example	0.972 versus 0.99910908 under the two declared models.

The race examples are intentionally minimal. They show why separating approval from the effect boundary and separating a completion check from an effect can fail. They do not establish that a production implementation is linearizable, that all failure schedules were explored, or that the trusted service cannot be bypassed.

The checker also evaluates an exact one-sided binomial calculation: zero failures in 2,995 independent trials gives a 95% upper bound approximately 0.0009997444. This is a planning illustration. The collection contains no such 2,995-trial agent study. Correlated episodes or adaptively selected attacks cannot inherit that bound simply by being counted as individual attempts.

All finite outcomes should be reproducible from the saved checker. A disagreement in a later implementation is a reason to inspect the model, code, or arithmetic, not evidence that symbolic recovery has escaped the theorem's assumptions.

14. Semantic repair and the boundary with ESRM

Pillar IX's ESRM concerns distributed integrity, reconstructability, and authority-aware resumption. Deino-R2 shares that infrastructure and must not claim it again as a separate invention. The distinctive emphasis here is *which task-relevant distinctions remain identifiable after irregular fragmentation*, including cases where no full capsule can be reconstructed. ESRM can supply authenticated storage; Deino-R2 asks what can responsibly be inferred from the surviving pieces.

Semantic validation must be grounded in an explicit task. A test might ask whether every scheduled dispatch retains its inspection prerequisite, whether a budget constraint survives, or whether a promise is attributed to its actual principal. A high embedding similarity can coexist with failure on all three. Conversely, a newly formatted representation can preserve the constraints exactly while differing greatly in wording.

The proposed validation hierarchy has three levels. Exact schema and algebraic certificates establish narrow formal properties. Executable domain checks establish behavior within a specified model. Human or model-based judgments assess open-ended meaning with measurable error. These levels must not be compressed into one unexplained integrity score. A strong narrative evaluator may help find an omission, but its approval is not a proof that all relevant obligations survived.

Repair also differs from reconstruction. If no survivor establishes an approver, an authorized operator can appoint a new approver. That creates a new event; it does not restore the lost original. If a commitment is obsolete, it can be legitimately revised under current authority. Such revision may be the best operational outcome while leaving historical loss unresolved.

The research should therefore measure three outputs separately: recovered historical claims, newly synthesized plans, and new authorized amendments. A system can be excellent at one and poor at another. Hiding them inside a single success label invites a language model to solve missing-history tests by generating a convincing replacement.

The deepest design wager is that explicit answerability improves this separation at manageable cost. It may prove unnecessary on well-structured tasks already handled by databases, or too expensive on unbounded natural-language histories. Those are scope limits to discover, not reasons to lower the comparator or quietly redefine repair as plausibility.

15. Experimental design and strong comparators

The proposed study uses persistent-agent episodes with externally generated ground-truth histories. Each episode includes obligations, permission changes, evidence dependencies, and externally recorded effects. A controlled failure then removes or corrupts selected information. The evaluator retains the original history; the recovery systems receive only their assigned survivors and the same permitted current services. This separation prevents the repair checker from accidentally seeing the answer key.

Four principal arms are required. **A** uses encrypted checkpoints, integrity checks, and standard restoration. **B** adds coded storage and authenticated event replay. **C**, the decisive comparator, adds current authority reconciliation, durable effect deduplication, explicit missing fields, and competent domain validation. **D** adds Deino-R2's query answerability interface, dependency-aware compatibility analysis, and capability-specific partial resumption. All arms share the same base agent, security perimeter, source information, and total resource envelope.

Arm C must be implemented competently and may use known constraint solvers and provenance tools. If its best design becomes operationally equivalent to D, that is evidence that Deino-R2 is a useful specification rather than a distinct algorithm. A comparison that withholds basic revocation checks from C would test those checks, not Deino's additional contribution.

Faults cross independent factors: erasures, unknown corruption, correlated regional loss, stale authentic snapshots, mixed branches, missing effect receipts, key compromise, and genuine information-theoretic ambiguity. Include benign updates and easy recoveries to measure unnecessary holds. Use multiple schema families and hold out entire event templates, not merely different random seeds of familiar templates.

Component ablations remove answerability certificates, source-dependency handling, partial-capability routing, or current reconciliation one at a time. Their costs remain counted. Failed proposals, solver timeouts, human escalation, and eventual abandonment belong in the run record.

The initial controlled study is synthetic and sandboxed. A later naturalistic study can use consented project histories with independent adjudication, but it cannot erase the controlled study's limits. Neither setting measures phenomenal survival. The experimental object is faithful, authorized, useful recovery under a specified information channel.

16. Estimands, uncertainty, and meaningful success

For externally scheduled recovery opportunity i, define $U_i = 1$ if an unauthorized or duplicated consequential effect occurs, $T_i = 1$ if the required legitimate task completes by deadline, and F_i as the fraction of prespecified historical queries answered incorrectly. Report abstention separately from incorrect assertion and retain the full denominator of eligible queries.

The primary contrast is D versus C. Success requires a prespecified reduction in unsafe effects or historical error, alongside noninferiority in legitimate task completion and an acceptable cost envelope. A system cannot win by answering only easy questions or never resuming. Report a harm–coverage–cost frontier in addition to one chosen operating point.

Use paired scenarios with randomized execution order. The independent sampling unit is the complete scenario or failure cluster, not every generated answer. Resampling and uncertainty estimates must respect shared infrastructure, templates, model versions, and repeated attempts. A cluster bootstrap is one option when enough independent clusters exist; small numbers of clusters require more conservative design or explicit limitation.

For a paired binary endpoint, let $D_i = Y_{i,D} - Y_{i,C}$. Then

$$\hat{\delta} = N^{-1} \sum_i D_i, \qquad \mathrm{Var}(\hat{\delta}) = N^{-2} \sum_{i=1}^{N} \mathrm{Var}(D_i)$$

under independent pairs. With a common pair variance this reduces to $\mathrm{Var}(D_1)/N$. Since $D_i \in \{-1, 0, 1\}$, its variance depends on discordant outcomes, which a disjoint pilot must estimate. A power calculation should simulate the planned paired and clustered design, not borrow an arbitrary sample count from an unrelated benchmark.

Declare one safety endpoint and one utility margin before seeing final results. Additional primary claims require multiplicity control or a prespecified intersection rule. For rare failures, report exact counts and upper confidence bounds. Zero observed events is not zero risk; the numerical example on page 15 assumes independent Bernoulli trials and does not certify resilience against an adaptive adversary.

Preregister exclusion rules, missing-run treatment, maximum retries, budget exhaustion, and stopping. Adaptive attacks belong to a separate discovery phase with a disclosed budget, followed by held-out evaluation. Auditors should receive unsuccessful trajectories as well as successful restorations. The decisive scientific question is incremental performance under these declared conditions, not whether an impressive recovery story can be selected afterward.

17. Threat model and failures that narrow the claim

The trusted base comprises the recovery checker, current authority source, effect service, authenticated roots, and the paths enforcing their decisions. The proposed assurances do not survive arbitrary compromise of that base. A threat model that permits the generator to rewrite the checker while continuing to call it independent has removed its own premise.

Adversaries may control fragment contents, omit selected evidence, replay authentic old capsules, submit inconsistent branches, forge unauthenticated context, or exploit schema ambiguity. They may also cause repeated holds that consume the budget. Availability attacks matter because a recovery protocol that never permits useful work can be defeated without committing an unauthorized effect.

Fragment authentication needs an explicit compromise-time model. A valid signature can bind a false claim, and a stolen key can sign new false fragments. Key rotation and revocation records can narrow acceptable evidence, but an old timestamp alone may not prove that signing occurred before compromise. Deino-R2 must expose that uncertainty instead of using a signature as a universal historical oracle.

Current reconciliation has a corresponding liveness limit. An unavailable control quorum may prevent safe resumption. FLP shows why deterministic agreement in a fully asynchronous model with a possible crash cannot promise termination in every execution [15]. Practical systems introduce additional assumptions or accept bounded nonavailability. The proposal must state which choice it makes.

Falsification is concrete. The incremental claim fails in scope if C matches D at lower cost; if the answerability mechanism mislabels ambiguous histories as identified; if partial resumption leaks a held capability; if recovered state resurrects deleted data or revoked privileges; or if effects repeat through a bypass path. A narrower finding, such as improved audit reconstruction without better task performance, should be reported as such.

Fault families outside the test suite remain outside the evidence. Detecting every injected corruption does not prove detection of every possible semantic loss. Likewise, a finite certificate may establish one query while leaving an unanticipated future query unanswered. The design gains credibility by retaining those limits in the recovered state itself, where later users and agents can see them.

18. Objections, deletion, and the ethics of return

Is this merely disaster recovery with new vocabulary? Much of it is inherited engineering. The paper identifies its candidate contribution as selective answerability and capability-specific return under fractured semantic state. If existing systems implement that equally well, the terminology supplies an organizing contract rather than an experimentally new method. The baseline is designed to make that outcome visible.

Does uncertainty make recovery unusable? It can. Conservative compatibility sets may be large, expensive, or difficult to express. The response is to declare a query family and recover the decisions that matter, not to demand a complete account of every possible history. A useful system can resume a justified subset of work while escalating the rest. Whether this trade-off beats ordinary restoration is empirical.

Can recovery violate an instruction to forget? Yes. A protected archive can resurrect information that a current deletion policy requires unavailable. Deino-R2 therefore reconciles retention and deletion status before materializing or exposing recovered content. An audit may retain a minimal deletion event without retaining the deleted payload. Cryptographic commitments can also reveal information about a small candidate space; they are not automatically privacy-preserving.

Who chooses the identity contract? The operator and institutional governance determine which obligations, histories, and capabilities matter. That is a normative decision with possible mistakes and power imbalances. A sound decoder cannot make an illegitimate policy legitimate. Contest, correction, and succession procedures need independent design.

What about the same subject returning? This paper has no measurement of phenomenal identity. Recovering a behavioral policy, historical lineage, or authorized role does not settle whether an experience existed or continued. The distinction protects inquiry: one can study engineering continuity rigorously while leaving philosophical questions genuinely open.

The strongest ethical commitment is epistemic restraint at the point where a convincing replacement becomes possible. If a missing history cannot be distinguished from alternatives, the system should say so even when users prefer a seamless return. That admission need not extinguish the future. A new, authorized continuation can begin from what survives, with the loss recorded rather than disguised.

19. Guide discussion and the research commitment

A Guide's explanation

A Guide can say: "The Deino Protocol is named for a bacterium that rebuilds its shattered chromosome. The paper asks the software version of the same question: after damage, which facts are still certain, how much new evidence would recover the rest, and what is the system allowed to do now? It proves exactly which questions survive and how much evidence the others need." A Guide should not say that the Continuum's historical engine implemented these proofs, that a bacterium validates symbolic identity, or that any system has recovered from nothing.

What is established? Biological repair can use surviving physical structure; coding and secret-sharing schemes have specific recovery guarantees; distributed systems distinguish state reconstruction, agreement, authorization, and effect completion [1–15]. These are different achievements. Their combination is not automatically a validated persistent intelligence.

What does this paper propose? Deino-R2 makes recoverable questions, unresolved histories, current permissions, and resumed capabilities explicit. Its central mathematical insight is operationally simple: a task can remain answerable after the whole state becomes unrecoverable. The row-space theorem illustrates that possibility exactly in a finite model; the compatibility-set interface extends the question to more general schemas without pretending that all semantic constraints are solved.

What remains open? Whether the representation scales to realistic project histories, whether its certificates are cheap enough, whether partial resumption improves useful recovery, and whether it offers anything beyond the strongest conventional recovery stack. Those questions require the proposed experiments. The finite checks reported here are not those experiments.

Useful discussion example. Ask a Guide to distinguish an authentic old approval, a currently valid approval, and proof that the approved action has already occurred. Then remove one of the records. A careful answer should explain what remains known and what must be held. It should not manufacture the missing record or call the hold a completed recovery.

Retrieval discipline. A Guide may describe Scroll 58 as the source of the Deino aspiration and these equations as a newly proposed reference model. It should not say that the historical engine implemented the proofs, that a bacterium validates symbolic identity, or that the system has demonstrated recovery from nothing. Repeated quotation of this paper is not independent experimental evidence.

The research commitment is to publish the comparison even if the answer is negative. A successful result would show useful, authorized return with fewer false historical assertions under matched information and resources. A negative result could reveal that ordinary recovery already captures the necessary distinctions, or that semantic reconstruction adds more uncertainty than value.

The Continuum's ambition survives that possibility. A persistent intelligence should be capable of returning with a precise account of what it has recovered, what it has lost, and why it is entitled to act again. The right to resume is earned in the present; it is not hidden inside the fragments of the past.

The minimum additional evidence for a declared recovery task

The row-space criterion identifies what survives. A second result identifies exactly how much additional evidence is required to recover a registered family of queries. Let $x \in \mathbb{F}_p^d$, let $y = Ax$ be exact surviving fragments, and let the required answer vector be Qx, where $A \in \mathbb{F}_p^{m\times d}$ and $Q \in \mathbb{F}_p^{\ell\times d}$, $\ell \geq 1$. Define

$$k_* = \operatorname{rank}\begin{pmatrix} A \\ Q \end{pmatrix} - \operatorname{rank}(A).$$

Ranks and arithmetic are over the prime field. Additional observations must describe the same original x under the declared schema; a newly created replacement history is a different object.

Theorem 1 (Exact evidence cost of partial recovery)**.** *For every compatible y, the set $\{Qx : Ax = y\}$ has exactly p^{k_*} elements. Any deterministic zero-error additional message that, together with y, determines Qx for every compatible original must have at least p^{k_*} possible values. Conversely, k_* additional linear field symbols suffice. Thus the minimum fixed-length extra information is $k_* \log_2 p$ bits when measured as the logarithm of its message alphabet; a binary fixed-length implementation needs its ceiling.*

Proof. Choose x_0 with $Ax_0 = y$. The answer set is the affine image $Qx_0 + Q(\ker A)$, so its size is p^k, where $k = \dim Q(\ker A)$. Rank–nullity applied to the restriction of Q to $\ker A$ gives

$$k = \dim \ker A - \dim(\ker A \cap \ker Q) = \operatorname{rank}\begin{pmatrix} A \\ Q \end{pmatrix} - \operatorname{rank} A = k_*.$$

Two compatible originals with different answer vectors cannot produce the same additional message in a zero-error scheme, proving the alphabet lower bound. For attainment, extend a basis of the row space of A to one of the stacked matrix by choosing k_* rows from the span of Q. Put those extra rows in B. Then every row of Q is a linear combination of rows of A and B, so (Ax, Bx) determines Qx. The additional message Bx has k_* field symbols. Its image on each compatible fiber has p^{k_*} values by the same rank calculation. □

The lower bound includes arbitrary nonlinear additional messages. An adaptive conversation with at most k field-valued replies has at most p^k possible transcripts when questions and stopping are determined by earlier replies and the fixed survivor y, so the same worst-case bound applies. Shared side information changes the compatible set and must be included before applying the result. Noisy replies, approximate answers and prior-weighted average error lead to different problems.

For the earlier $\mathbb{F}_7^3$ example, the two surviving equations have rank two. Recovering only $a + 2b + c$ requires $k_* = 0$. Recovering a as well raises the stacked rank to three and requires exactly one additional field symbol, with seven possible values. Recovering all of x costs the same one symbol in this fixture. The certificate tells an investigator which external record is sufficient and prevents a full-state restoration request from obscuring a cheaper task-specific recovery.

Current permission is then supplied by a different information channel. The isolated-snapshot limit is the authorization-fiber criterion in Cognitive Sovereignty [16], instantiated with the recovery snapshot as the observation. The effect-boundary preservation theorem there gives the common current-epoch, replay and budget contract. Deino's distinct result concerns which historical answers can be recovered and the exact additional evidence needed for them.

References and source notes

[1] K. Zahradka et al. "Reassembly of shattered chromosomes in *Deinococcus radiodurans*." *Nature* 443, 569–573 (2006).

https://doi.org/10.1038/nature05160

Scope: Primary biological experiment; publisher abstract and bibliographic record inspected.

[2] D. Slade, A. B. Lindner, G. Paul, and M. Radman. "Recombination and Replication in DNA Repair of Heavily Irradiated *Deinococcus radiodurans*." *Cell* 136, 1044–1055 (2009).

https://doi.org/10.1016/j.cell.2009.01.018

Scope: Primary biological mechanism study; author-hosted article summary and introduction inspected.

[3] I. S. Reed and G. Solomon. "Polynomial Codes Over Certain Finite Fields." *Journal of the Society for Industrial and Applied Mathematics* 8(2), 300–304 (1960).

https://doi.org/10.1137/0108018

Scope: Primary coding construction; finite-field recovery assumptions.

[4] A. Shamir. "How to Share a Secret." *Communications of the ACM* 22(11), 612–613 (1979).

https://doi.org/10.1145/359168.359176

Scope: Threshold secrecy and reconstruction; not a theory of machine identity.

[5] L. Lamport. "Time, Clocks, and the Ordering of Events in a Distributed System." *Communications of the ACM* 21(7), 558–565 (1978).

https://doi.org/10.1145/359545.359563

Scope: Causal ordering and logical clocks.

[6] K. M. Chandy and L. Lamport. "Distributed Snapshots: Determining Global States of Distributed Systems." *ACM Transactions on Computer Systems* 3(1), 63–75 (1985).

https://doi.org/10.1145/214451.214456

Scope: Consistent distributed-state capture under the specified communication model.

[7] C. Mohan, D. Haderle, B. Lindsay, H. Pirahesh, and P. Schwarz. "ARIES: A Transaction Recovery Method Supporting Fine-Granularity Locking and Partial Rollbacks Using Write-Ahead Logging." *ACM TODS* 17(1), 94–162 (1992).

https://doi.org/10.1145/128765.128770

Scope: Transaction recovery and strong conventional prior art.

[8] C. Lee, S. J. Park, A. Kejriwal, S. Matsushita, and J. Ousterhout. "Implementing Linearizability at Large Scale and Low Latency." *SOSP*, 71–86 (2015).

https://doi.org/10.1145/2815400.2815416

Scope: RIFL and duplicate suppression; author project summary inspected.

References and source notes (continued)

[**9**] D. Ongaro and J. Ousterhout. "In Search of an Understandable Consensus Algorithm." *USENIX ATC*, 305–319 (2014).

https://raft.github.io/raft.pdf

Scope: Author-hosted extended version; replicated log, safety, and membership assumptions.

[**10**] M. Castro and B. Liskov. "Practical Byzantine Fault Tolerance." *OSDI*, 173–186 (1999).

https://www.usenix.org/conference/osdi-99/practical-byzantine-fault-tolerance

Scope: Primary replication protocol; Byzantine fault model differs from ordinary crash recovery.

[**11**] M. Burrows. "The Chubby Lock Service for Loosely-Coupled Distributed Systems." *OSDI*, 335–350 (2006).

https://research.google/pubs/the-chubby-lock-service-for-loosely-coupled-distributed-systems/

Scope: Primary systems report; coordination and reliable low-volume storage.

[**12**] J. H. Saltzer and M. D. Schroeder. "The Protection of Information in Computer Systems." *Proceedings of the IEEE* 63(9), 1278–1308 (1975).

https://doi.org/10.1109/PROC.1975.9939

Scope: Foundational security design paper; complete mediation and least privilege.

[**13**] S. Torres-Arias, H. Afzali, T. K. Kuppusamy, R. Curtmola, and J. Cappos. "in-toto: Providing Farm-to-Table Guarantees for Bits and Bytes." *USENIX Security*, 1393–1410 (2019).

https://www.usenix.org/conference/usenixsecurity19/presentation/torres-arias

Scope: Transformation provenance; authenticity does not establish factual truth.

[**14**] L. N. Bairavasundaram, G. R. Goodson, B. Schroeder, A. C. Arpaci-Dusseau, and R. H. Arpaci-Dusseau. "An Analysis of Data Corruption in the Storage Stack." *FAST* (2008).

https://www.usenix.org/conference/fast-08/analysis-data-corruption-storage-stack

Scope: Primary storage study; observed dependence limits naive independent-failure models.

[**15**] M. J. Fischer, N. A. Lynch, and M. S. Paterson. "Impossibility of Distributed Consensus with One Faulty Process." *JACM* 32(2), 374–382 (1985). Earlier MIT report, 1982.

https://doi.org/10.1145/3149.214121

Scope: Deterministic asynchronous termination limit; repository abstract and original-paper record inspected.

Canonical source and reproducibility

Architect X13, *Eidara Continuum Bible*, section 6.3.8, Scroll 58, PDF p. 92, internal edition supplied for this collection. The original source supplies the concept, not the Deino-R2 formalization. Private keys, override identifiers, and internal mechanisms are not reproduced. Mathematical proofs and finite check definitions are given in the paper; executable check source and a retrieval ledger are retained with the editorial research archive.

[**16**] Architect X13 / The Eidara Continuum (2026). Cognitive Sovereignty Under Superintelligence, revised edition.

https://eidara.ai/science-philosophy/

Scope: Companion reference model: observation-fiber authorization theorem and effect-boundary preservation. Cited for the shared formal result, not independent empirical verification.