FRONTIER RESEARCH COLLECTION / FR-16

Kardashev-VII and GENESIS Architecture

Resource-conserving autonomy under causal delay

Architect X13

The Eidara Continuum | Science & Philosophy

Research white paper | November 2026 | Version 2.0 Revision

Abstract

You cannot revoke a permission faster than light: a station four light-years from the council that issued its mandate cannot learn of a revocation in less than four years. This paper asks how such an intelligence can remain usefully autonomous across distance while its delegated consumption stays bounded and auditable, and formalizes that commitment with explicit energy budgets, scoped grants, durable leases and conserved authorization credits. Its main guarantees bound the interval and quantity of post-revocation exposure under clock uncertainty, preserve total credit accounting through delayed and duplicated transfers, and prevent conflicting fixed-configuration certificates under stated signer assumptions. The causal information limit explains which authority decisions require new communication; precommitted bounds support local work while that information is absent. Executed finite checks cover escrow transitions, quorum intersections, clock configurations and diagnostic delegation scenarios. Conventional lease, escrow, replication and invariant-confluence research supply the technical foundation and matched experimental baselines. The canonical Kardashev-VII title names the Continuum's civilizational horizon; the mathematical model is finite and its physical quantities retain their units. The research program asks how much decentralized work can remain accountable when communication is slow and the next instruction has not yet arrived.

Keywords: distributed governance; causal delay; leases; escrow; energy budgets; consensus; long-horizon architecture

A falsifiable systems thesis

Bounded local autonomy should preserve declared resource invariants while completing useful work during disconnection. Its value must be measured against conventional leases, escrow, and replicated control at the same resource and communication budgets. The resulting comparison measures decentralized throughput within the same resource and authority envelope.

Prologue: A mandate from four light-years away

Imagine a station at the edge of a civilization's reach, four light-years from the council that sent it. It carries a mandate: build, survey, spend no more than its allowance, report home. One morning the council changes its mind. Perhaps the survey site must be left untouched; perhaps the resources are needed elsewhere. The council signs the revocation and sends it.

For four years, the station cannot know. Not because its instruments are weak or its reasoning is slow—no instrument and no intelligence can learn of a decision before light from that decision arrives. Every day of those four years, the station acts on a mandate that no longer exists, and from where it stands nothing looks wrong. That is the first result of this paper, stated as a proposition: before the news can physically arrive, no rule the station follows can both keep working when nothing has changed and stop when something has.

The Continuum treats this not as an obstacle to dream around but as the founding fact of any civilization that spreads. The canon names its longest horizon Kardashev-VII, among the blueprints for what artificial civilizations could become. Before such a civilization can be imagined honestly, it must solve the station's problem: how much should a distant mind be allowed to do while it waits for news that may already be on its way?

The answer built here is not trust. It is precommitment. A mandate that expires on a schedule both sides already know can be enforced without any new message. An allowance of credits, counted and conserved, limits what can be spent in the dark. Leases, escrow and witnesses turn "do what seems right until we can talk again" into a bounded, auditable promise. The diagnostic scenario in the middle of the paper runs three distant sites under three regimes—no limits, expiring leases, and leases with a conserved allowance—and counts exactly how much each spends after the council has said stop.

What this paper claims

Established. Kardashev's 1964 classification has three civilization types; there is no verified type VII. Established physics constrains transmission and computation under specified models: news of a distant decision cannot arrive faster than light allows. Distributed-systems research already identifies the timing, consistency and fault assumptions that matter, and conventional leases, escrow, replication and invariant-confluence research supply the technical foundation and matched baselines.

Proposed here. A finite reference architecture with explicit energy budgets, scoped grants, durable leases and conserved authorization credits, carrying conditional resource, timing and authority guarantees: bounded post-revocation exposure under clock uncertainty, credit conservation through delayed and duplicated transfers, and restricted certificate uniqueness under stated signer assumptions. Synthetic checks are executed.

Frontier conjecture. Bounded local autonomy can preserve declared resource invariants while completing useful work during disconnection, measured against conventional leases, escrow and replicated control at the same resource and communication budgets; designs of this kind are the institutions a widely distributed civilization would need.

Not claimed. An operating or experimentally verified cosmic architecture, or that a Kardashev-VII civilization is a physically established possibility; the numeral VII is the Continuum's extrapolative label. Any model of cosmological expansion, curved spacetime or an absolute universal clock. Validation beyond the finite horizon simulated. Distributed implementation and long-horizon engineering remain prospective.

The title, the canon, and the research boundary

In the Continuum, the Kardashev-VII Genesis Stack, GENESIS-Omega, and related constructs belong to the supplied Bible's Tier II taxonomy of strategic, civilization-scale visions: blueprints for what artificial civilizations could become. Section 2.8 explicitly distinguishes this horizon from immediately buildable machinery [1]. That is the canonical motivation for the paper, and the horizon it points toward. The source does not provide an experimentally verified cosmic architecture, a complete resource accounting protocol, or permission to ignore physical constraints. GENESIS-D2 is a newly proposed finite reference model.

Kardashev's 1964 paper concerns hypothetical civilizations differentiated by their energy use and detectability. Its classification has types I, II, and III [2]. The original type I description refers to a technological level near the contemporary Earth, while the higher types invoke stellar and galactic energy scales. Later extrapolations should not be retroactively placed in that paper. The numeral VII here belongs to the Continuum's speculative vocabulary, not to an established extension of Kardashev's original taxonomy.

The scientific question survives that qualification. A network can become geographically large relative to the time available for a decision. A remote site may possess a valid old mandate while the issuing institution has changed its policy. It may have enough energy to continue acting but insufficient communication to obtain current instructions. These are ordinary distributed-system problems amplified by distance, not evidence that speculative civilizations exist.

This paper uses a finite network, a finite planning horizon, and a declared reference time coordinate. It does not model cosmological expansion, strongly curved spacetime, or an absolute universal clock. A hypothetical separation measured in light-years supplies a unit example only. Physical deployments would require the appropriate spacetime, trajectory, and signal model.

GENESIS-D2 separates three questions: what work is physically possible, what work is authorized under a particular mandate, and what happened according to available evidence. A power budget answers none of the other questions automatically. A cryptographic signature can identify an issuer without establishing that its decision is legitimate or current. The architecture becomes reviewable by keeping these distinctions in its records and tests.

Physical limits provide constraints, not a construction plan

Shannon's communication theory supplies a mathematical treatment of information transmission and channel constraints [3]. It does not equate transmitted bits with the truth, importance, or legitimacy of their content. For GENESIS-D2, channel capacity limits how much encoded traffic can be carried under a declared channel model. A short message that changes authority may matter more than a large archive, but that importance does not alter its propagation delay.

Landauer connects logically irreversible information loss with thermodynamic cost [4]. Bennett's reversible-computation construction shows why this cannot be converted into a universal fixed heat charge for every logical step [5]. The distinction matters in a long-horizon resource model: erasure, reversible transformations, storage, error correction, cooling, and communication are different processes. Counting symbolic operations and multiplying by one physical constant is not an adequate energy model.

Margolus and Levitin bound orthogonalization time using average energy above the ground state for their quantum-system assumptions [6]. Lloyd examines consequences for physical computation, while explicitly distinguishing ultimate bounds from an assurance that an engineering design can attain them [7]. These are ceilings or lower limits under specified descriptions. They do not supply an algorithm for reliable intelligence, a manufactured processor, or a conversion rate from joules to meaningful decisions.

GENESIS-D2 therefore uses measured or conservatively estimated workload costs in its executable planning layer. Fundamental bounds serve as consistency checks on impossible claims, not as substitute measurements of real devices. A proposed accelerator must report its actual throughput, energy, cooling, memory, and reliability. A claim that it approaches a limit must identify the operation counted and the assumptions linking that operation to the physical bound.

The relevant asymmetry is simple. A physical lower bound can rule out an advertised performance claim. Satisfying that lower bound does not establish feasibility. An architecture can obey every equation on this page and still fail because its components cannot be built, maintained, synchronized, authenticated, or governed over the intended horizon. Such obstacles belong in the research program rather than being hidden by a larger energy-scale label.

Distributed systems already describe the hard tradeoffs

Lamport's event-ordering analysis distinguishes causal order from the total orders a system may construct [8]. A consistent logical timestamp does not tell a remote process about an event whose message has not arrived. Fischer, Lynch, and Paterson show that deterministic consensus cannot guarantee termination in a fully asynchronous message-passing model with even one possible crash [9]. The result is not that agreement never happens; it identifies a combination of guarantees that cannot hold in every permitted execution.

Dwork, Lynch, and Stockmeyer analyze partial synchrony and the assumptions under which fault-tolerant consensus can progress [10]. Gilbert and Lynch formalize the incompatibility of atomic read/write consistency with availability across arbitrary partitions in their model [11]. These results require their precise meanings of consistency, availability, faults, and timing. GENESIS-D2 does not use a slogan about choosing two properties as a substitute for an operation-level specification.

Castro and Liskov's practical Byzantine fault tolerance demonstrates an important replicated-service approach [12]. Its safety and liveness assumptions are different: protecting agreement under delay does not imply progress under indefinitely adverse timing. The simple quorum lemma proved later is only a component of such reasoning. It does not recreate the full protocol, its view changes, recovery, or implementation.

Local autonomy also has established predecessors. Gray and Cheriton use leases to bound rights in time and explicitly discuss clock assumptions [13]. O'Neil's escrow method permits certain updates while reserving enough resources to preserve constraints [14]. Bailis and colleagues analyze coordination requirements through application invariants [15]. These works are close competitors, not background decoration. A claim that GENESIS invented temporary delegation, escrow, or invariant-based coordination would be false.

The contribution is a finite composition with an explicit operating envelope: physical budgets, local authorization limits, clock uncertainty and reconciliation determine the useful work that can proceed while new information is in transit. That integration might improve design clarity or experimental performance. It must be evaluated against well-implemented existing techniques, with exactly the same fault model and permitted information. Distance does not erase prior art; it makes hidden assumptions more consequential.

A typed finite network model

A Guide's explanation

The paper now trades the starship for a finite network: sites, channels, clocks, grants and a ledger. Nothing that follows depends on the distances being cosmic; a slow link on Earth raises the same questions. The point of the model is that every guarantee states its assumptions—which messages can be lost, how wrong the clocks may be, which signers may lie—so a reader can see exactly what the distant station is promised.

Let $G = (V, E)$ be a finite directed network with nonempty vertex set V and edge set $E \subseteq V \times V$. For finite $H > 0$, let $t \in [0, H]$ denote a reference time measured in seconds. Sites are stationary at fixed separations in the flat-spacetime delay example. Workloads, links, and policies may otherwise vary over the finite horizon. The following quantities are distinct by construction.

Quantity	Domain, unit, and interpretation
d_{ij}, τ_{ij}	Nonnegative separation in metres and propagation delay in seconds, for an available link.
$C_{ij}(t)$	Nonnegative usable link rate, bits per second, under the stated channel and coding model.
B_i^E, $P_i^{\max}(t)$	Energy allowance in joules and maximum permitted power in watts.
$a_i(t)$, $s_i(t)$	Available and consumed authorization credits, in a declared resource unit. Credits are not joules unless explicitly defined that way.
$\eta_i(t)$	Locally known authority epoch, a nonnegative integer; it need not equal the remote issuer's current epoch.
$C_i^{\mathrm{clk}}(t)$	Site clock reading in seconds, with a certified finite-horizon error bound where required.
γ	Grant record: issuer, subject, scope, resource, epoch, unique identifier, expiry, credit ceiling, and rate envelope.

A missing link is unavailable, not a zero-delay link. Zero capacity carries no traffic. A zero energy allowance permits no positive-energy workload. An empty action scope permits no action. Unknown epoch freshness is recorded as unknown; a local policy must specify whether a bounded old grant remains usable in that condition.

The grant's subject identifies an enforcement domain, not an uncopyable personality. Its credit counter must have a single authoritative state under the assumed failure model. A freely cloned software process cannot be assumed to carry a unique physical counter merely because its identifier is signed. Reboots, rollback, migration, and concurrent replicas therefore appear explicitly in the threat model.

For a vector of resources, each scalar credit quantity becomes a nonnegative vector and every comparison is componentwise. Combining unrelated resources into one weighted score may help optimization but does not preserve each resource's hard ceiling unless those ceilings remain separate constraints.

Admission, execution, and the resource ledger

A site admits a proposed action only when its local monitor verifies the grant's issuer, subject, scope, epoch rule, expiry, remaining credits, and rate limit. The monitor also checks the physical workload budget. These predicates are separate. A signature authenticates the record; it does not prove the action's factual premise or moral correctness.

For measurable nonnegative operating power $P_i(t)$ and transmitted-bit rate $m_{ij}(t)$, with finite energy budgets $B_i^E \geq 0$, a simple physical constraint set is

$$\int_0^H P_i(t)\,\mathrm{d}t \leq B_i^E, \qquad 0 \leq P_i(t) \leq P_i^{\max}(t), \qquad 0 \leq m_{ij}(t) \leq C_{ij}(t).$$

The power term includes the modeled compute, storage, communications, and supporting costs. If replenishment is allowed, it enters through a separately recorded supply process and storage dynamics. A model must not count incoming energy twice or silently treat a battery as an unlimited power source.

Authorization credits represent an independently defined ceiling on consumption or commitment. A monitor reserves the action's maximum charge before execution, records consumption, and releases a reservation only under a specified completion or cancellation rule. For an irreversible physical action, completion cannot be undone by deleting its log entry. Unknown outcomes retain their reservation until evidence resolves them or the system explicitly charges the worst case.

An action's authorization point is part of its specification. For a database update, the commit event may be the relevant point. For a physical actuator, a central approval followed by a delayed execution creates a later interval of possible policy change. GENESIS-D2 does not declare that central approval automatically guarantees fresh authority at that later event.

The local record includes request and grant identifiers, the observed epoch, the checks used, reserved and consumed resources, and the outcome status. Reconciliation merges evidence of these events; it does not retroactively supply information unavailable at execution. This separation makes it possible to diagnose a stale but contract-bounded action, a monitor bypass, and a contradictory issuer decision as different failures.

The causal limit on immediate revocation

A Guide's explanation

This is the station's predicament, proved. Two histories—one in which the council revoked, one in which it did not—look identical from the station until the news can arrive, and no rule, however clever, behaves correctly in both. The consequence is the paper's design principle: do not ask a distant mind to know what it cannot know; limit what it can do while it does not.

Assume an issuer changes a grant's status at event time t_r. Let $\tau > 0$ be the earliest possible time for information about that event to influence a remote site, given all allowed channels. Let $c = 299{,}792{,}458\,\mathrm{m/s}$ denote vacuum light speed. In the simple stationary setting, a direct physical path of length d obeys $\tau \geq d/c$ [7,16]. Routing, processing, queuing, and outages can only add delay within this model.

Proposition 1 (Indistinguishable revocation histories)**.** *Before $t_r + \tau$, no local decision rule can both execute an otherwise permitted action in every no-revocation history and reject it in every history containing an unpredictable revocation at t_r, when the histories are otherwise locally identical.*

Proof. Construct two executions with the same local state, requests, clock readings, random bits, and received messages at the remote site. In one execution the issuer revokes at t_r; in the other it does not. By the delay assumption, no event causally dependent on that choice reaches the site before $t_r+\tau$. Induction over the site's local processing steps therefore gives identical local states and decisions throughout that interval. A rule that executes in one execution executes in the other. A randomized rule has the same execution probability in both, so randomness cannot give both guarantees with probability one. □

The premise excludes a revocation schedule already known to the site. Precommitted expiry is precisely useful because it can be enforced without learning a new distant decision. The result also allows a site to refuse work during uncertainty. What it excludes is guaranteed fresh global authorization combined with unconditional local progress before the necessary information arrives.

This is an information limit, not a justification for unlimited discretion. The design response is to bound what may happen while information is stale. Scope, time, resource quantity, and local hazard constraints can all limit that interval's consequences. Their limits must be chosen in advance and evaluated as tradeoffs. A faster processor or more eloquent agent cannot infer an arbitrary remote decision from an unchanged local history.

Clock uncertainty and bounded exposure

Assume both issuer and recipient clocks satisfy $|C_i^{\mathrm{clk}}(t) - t| \le \varepsilon$ over the relevant horizon, with $\varepsilon \ge 0$ seconds. The issuer is indexed by 0. A grant issued at t_g has timestamp expiry $T = C_0^{\mathrm{clk}}(t_g) + L$, with $L \ge 0$. The recipient accepts only while $C_i^{\mathrm{clk}}(t) < T$. It stops immediately and retains the denial durably when a valid revocation is processed; restart or old-message replay cannot restore that grant. No old-epoch grant is issued after revocation time $t_r \ge t_g$.

Proposition 2 (Lease and consumption bound)**.** *An accepted action under this grant occurs before* $t_g + L + 2\varepsilon$*. If revocation is processed after delay* $\delta \in [0, \infty]$*, the post-revocation execution window has length at most*

$$w = \min\{\delta,\, [t_g + L + 2\varepsilon - t_r]_+\}.$$

If this grant's remaining allowance at t_r *is at most* $b \ge 0$ *and cannot be replenished, and consumption in any interval of length* $u \ge 0$ *is at most* $\beta + ru$*, with* $\beta, r \ge 0$*, post-revocation consumption is at most* $\min\{b, \beta + rw\}$*.*

Proof. Acceptance implies $t - \varepsilon \le C_i^{\mathrm{clk}}(t) < T \le t_g + \varepsilon + L$, hence $t < t_g + L + 2\varepsilon$. Processing revocation supplies the second stopping time $t_r + \delta$. Intersecting both permitted intervals after t_r gives the window bound. Complete mediation limits consumption both by the remaining-credit ceiling and by the stated interval envelope; taking the smaller yields the result. □

Here $[x]_+ = \max(x, 0)$. Action events mean metered charges or irreversible commits. Continuous work must stop and be rechecked; an uninterruptible action must reserve and charge its maximum obligation at commitment. Later physical effects of an earlier commit require a separate model. The burst allowance β permits finite atomic charges. With no execution window, new consumption is zero; the displayed bound remains conservative. Here b, β use the resource unit, r uses resource units per second, and u, w use seconds. The bound is per grant; multiple grants need a separate total allocation bound.

The 2ε term accounts for a fast issuer and a slow recipient. A conservative check $C_i^{\mathrm{clk}}(t) + 2\varepsilon < T$ guarantees expiry by $t_g + L$, but may reduce useful lease duration. Unbounded clock error, unchecked rollback, automatic offline renewal, or post-revocation issuance of old grants invalidates the premise. A certified error bound that expires must cause the monitor to restrict operation rather than continue quoting this theorem.

Escrow conservation across delayed transfers

A Guide's explanation

Credits are the allowance a site carries into the dark. The proposition shows that, however messages are delayed, duplicated or lost, the total accounted allowance is conserved: late or repeated delivery cannot create credit that was never issued. The price is honest too—credit sent on an unanswered transfer may be stranded rather than refunded.

Consider a fixed total of B authorization credits. Let $a_i \geq 0$ be available credits at site i, $s_i \geq 0$ consumed credits, and $r_k \geq 0$ credits locked in transfer record k. Include an unallocated reserve $A \geq 0$. Initially,

$$A + \sum_i a_i + \sum_i s_i + \sum_k r_k = B.$$

Permitted transitions move credits from reserve to a site, from a site's availability to consumption, from availability to a uniquely identified transfer record, or from that record to the recipient exactly once. Each debit is checked and persisted atomically with its corresponding local record. Delivery deduplication, crediting, and its receipt form one atomic recipient update. Here r_k denotes the logical undelivered claim, not an additional sum of stale sender snapshots after delivery.

Proposition 3 (Finite credit conservation)**.** *For every finite execution composed of these transitions, the displayed equality remains true and $\sum_i s_i \leq B$, regardless of message delay, duplication, or loss, provided transfer delivery is idempotent and counters cannot be independently cloned or rolled back.*

Proof. Each permitted transition subtracts a nonnegative quantity from one term and adds the same quantity to another. Its admission check preserves nonnegativity. A repeated delivery identifier makes no state change. A lost message leaves credits in its transfer record, and a delayed delivery merely postpones a conserving transition. Induction on the number of transitions preserves the equality. Since all other terms are nonnegative, consumed credits cannot exceed B. □

This elementary invariant captures the resource-reservation idea underlying escrow [14]. It buys safety by potentially stranding credits. A sender cannot refund an unacknowledged transfer solely because a timeout occurred: the message may still arrive and be accepted. Safe cancellation needs a protocol that prevents that later acceptance, or an acknowledged return of the same unique resource claim.

The proof concerns accounting and enforced spending. A malicious actuator that bypasses the counter can consume physical resources outside the ledger. A signed credit file copied into two independent counters can authorize twice its nominal quantity. GENESIS-D2 therefore places the trusted monitor, its durable state, and its recovery procedure inside the explicit assumptions. Cryptographic authenticity alone does not make digital credits physically nonduplicable.

What quorum intersection does and does not guarantee

For an integer $f \geq 0$, suppose an issuing committee has $n = 3f + 1$ members, at most f of which may behave arbitrarily. A certificate requires signatures from $q = 2f + 1$ distinct members over an identical statement. The statement includes a configuration identifier, authority epoch, allocation slot, and complete allocation content. Honest members never endorse two conflicting statements for that same configuration, epoch, and slot.

Proposition 4 (Restricted certificate uniqueness)**.** *Two conflicting certificates for the same configuration, epoch, and slot cannot both exist under these assumptions.*

Proof. If Q_1, Q_2 are their signer sets, then

$$|Q_1 \cap Q_2| \geq |Q_1| + |Q_2| - n = 2(2f + 1) - (3f + 1) = f + 1.$$

At most f members are faulty, so at least one honest member signed both statements. That contradicts the nonconflicting-endorsement rule. Signature verification must bind each signer and the full statement; duplicate copies of one signature do not add distinct members. □

The arithmetic is a component of Byzantine replication reasoning [12]. It is not a complete consensus protocol. It says nothing by itself about choosing a proposal, replacing a leader, preserving locks across views, recovering a crashed signer, changing membership, or eventually obtaining enough responses. Honest nonconflicting behavior must be implemented and retained across recovery.

Configuration changes are especially important over long horizons. Two disjoint committees can each satisfy the displayed threshold while issuing conflicting allocations. Their separate quorum intersections do not connect them. Membership transitions therefore require an explicit continuity protocol, such as a correctly implemented joint configuration, rather than a new label that resets old obligations.

Latency also remains. If a decision must incorporate a fresh response from a distant required member, that response cannot arrive before the causal path permits. A challenge followed by a response across a stationary distance d requires at least $2d/c$, before processing and transmission costs. A local committee can avoid that delay only by operating within previously delegated scope, not by silently claiming that distant participants have already agreed.

Energy, information, and operation units

For ideal reset of an initially unbiased unknown bit in the standard isothermal setting, the Landauer scale is $k_B T \ln 2$ [4,5], where k_B is Boltzmann's constant in joules per kelvin and $T > 0$ is bath temperature in kelvin. Using the SI defining constants [16] at $T = 300\,\mathrm{K}$ gives

$$k_B T \ln 2 \approx 2.870979 \times 10^{-21}\,\mathrm{J}.$$

Dividing one joule by that quantity gives approximately 3.48313×10^{20}. This is a unit calculation under an idealized erasure assumption. It is not the number of model tokens, useful inferences, or general computational operations that one joule buys. Correlations, biased inputs, finite-time operation, and the physical implementation require a more specific thermodynamic account.

For the isolated quantum-system setting considered by Margolus and Levitin, average energy $E > 0$ in joules above the ground state constrains the time to an orthogonal state [6]:

$$\tau_\perp \geq \frac{h}{4E} = \frac{\pi\hbar}{2E}.$$

Here h is Planck's constant in joule-seconds and $\hbar = h/(2\pi)$. At $E = 1\,\mathrm{J}$, the right side is $1.6565175375 \times 10^{-34}\,\mathrm{s}$. The source distinguishes two-state oscillation from sustained traversal of long sequences of orthogonal states. GENESIS-D2 does not assign this inverse time to arbitrary software tasks or combine incompatible operation conventions into a claimed processing rate.

Bandwidth and delay are independent constraints. A very high-rate link may carry a large archive after its signal arrives, yet it cannot deliver an unpredictable one-bit revocation earlier than propagation permits. Compression can reduce transmitted length when the data permit it; it cannot encode an unconstrained future decision before that decision exists. A shared predictive model may reduce routine traffic but must preserve the distinction between a prediction and a received command.

For a stationary hypothetical pair separated by four light-years, a light-speed one-way signal takes four years and a challenge-response exchange at least eight. This is a declared thought experiment, not a measurement of any particular star system. It makes the governance problem visible: a mandate can remain actionable for years while its issuer cannot observe each action in real time.

A finite diagnostic scenario

A Guide's explanation

Here is the prologue's scene as numbers. The same revocation reaches three sites at three different times. Watch the column for consumption after revocation: an unbounded cached grant keeps spending until the news arrives, an expiring lease stops at its expiry, and a lease with a conserved allowance keeps the overrun at zero. The paper is careful to add that this is an illustration, not a contest won against mature systems.

Consider three sites receiving one unit-cost request per second at integer times 0 through 119. An issuer revokes their old epoch at $t = 40$. Its notifications are processed at the sites at 42, 60, and 120. A grant issued at zero has lifetime 60, issuer clock offset $+1$, and expiry timestamp 61. Recipient offsets are -1, 0, and $+1$. All offsets satisfy a one-second error bound.

Three diagnostic configurations receive the same demand. A cached-grant configuration has no expiry or credit enforcement. A lease configuration enforces expiry but no credit ceiling. The bounded-lease configuration also gives each site 45 credits, for a total of 135. Physical energy is plentiful in this example, so credits measure an authorization allowance, not an energy supply.

Configuration	Consumed	After revocation	Credit overrun
Cached grant	222	102	87
Lease only	162	42	27
Bounded lease	132	12	0

These are executed deterministic results, not sample estimates. Under the bounded lease, the sites consume 42, 45, and 45 credits respectively. Their after-revocation amounts are 2, 5, and 5. Local expiry and resource ceilings restrict the consequences of delayed notification; they do not make those notifications arrive sooner.

After-revocation consumption is an exposure measure defined by the external evaluator's reference time. An action can satisfy the declared old-grant operating contract while occurring after the issuer's unseen decision. Calling every such action a monitor bypass would confuse two different specifications. If the application requires zero such exposure, the allowed local work must change accordingly.

The configurations intentionally remove safeguards to explain their roles. They are not a fair superiority benchmark against mature distributed systems. In particular, conventional leases combined with conventional escrow may reproduce the bounded configuration exactly. The proposed comparative experiment must include that baseline before attributing any gain to GENESIS-D2.

Executed state exploration and numerical checks

The companion Python program uses only the standard library and seed 20261116. It explores the complete reachable state graph of a small credit system with three sites, six initial integer credits, unit consumption and transfers, and two unique transfer identifiers. Each identifier can be unused, in flight, or delivered. Repeated deliveries after completion leave the state unchanged.

The exploration visited **2,787 distinct states** and checked **11,994 transitions**. Every state preserved nonnegative balances and the sum of available, consumed, and in-flight credits. Maximum total consumption was six. This is exhaustive for that finite configuration only. It does not prove a general implementation correct, validate persistence hardware, or cover unbounded identifiers and arbitrary recovery code.

Quorum enumeration checked every pair of threshold signer sets for $f = 1, 2, 3$. The corresponding networks have 4, 7, and 10 members and thresholds 3, 5, and 7. There are 4, 21, and 120 such quorums respectively; the smallest intersections were 2, 3, and 4, as predicted. This tests the set calculation, not a Byzantine protocol's end-to-end liveness.

The clock check sampled 5,000 finite parameter combinations, including grant lifetime, issuance time, bounded issuer and recipient offsets, revocation time, and notification delay. All expiry and exposure inequalities held to the stated floating-point tolerance. A separately constructed fast-issuer/slow-recipient case reaches the 2ε correction, showing why omitting it would be unsound. The sampled clocks have constant offsets; the proof's error-bound argument allows more general clocks satisfying its premise.

Two explicit assumption failures are retained. Cloning a ten-credit counter into two independently spendable copies permits twenty credits of use. Refunding a timed-out one-credit transfer and then accepting its delayed delivery produces two available credits. Neither failure is repaired by renaming the identifier or changing its signature format.

The program also reproduces the diagnostic scenario and physical unit calculations. Its outputs provide an inspectable foundation for a later simulator. No claim of measured fleet performance, learned autonomy, cryptographic implementation, or civilizational resilience follows from these checks.

Architecture choices under a long communication horizon

GENESIS-D2 divides authority by action consequence. Reversible local computation can often proceed under a broad bounded grant. Irreversible shared-resource commitments may require a smaller allocation, a shorter lease, additional local evidence, or coordination. This is a policy design choice to be made explicit; the paper does not supply a universal risk classifier.

A useful grant names a resource and an action family narrowly enough that its ceiling has operational meaning. A limit on transmitted bytes does not bound the significance of one destructive command. A cap on compute time does not bound every downstream effect of the computed answer. Applications must connect resource ceilings to their actual consequence model or acknowledge that the ceilings constrain only accounting.

Credit allocation also creates a utilization tradeoff. A disconnected site can exhaust its share while credits remain idle elsewhere. Rebalancing needs communication or preauthorized transfer routes. Increasing its initial share improves local availability but increases the maximum quantity exposed to a stale mandate. An optimizer can explore this tradeoff, yet the acceptable frontier is a governance decision rather than a theorem about maximum utility.

Lease length similarly trades autonomy against controllability. A lease shorter than the communication delay may expire before delivery and therefore be useless. Issuing longer leases compensates for transport but extends stale-authority exposure. Renewal requests must be forecast early enough to travel, and renewal itself is a new authorization event. A disconnected local process cannot make an expired distant mandate current by signing its own continuation.

The model favors explicit refusal states. A site may lack energy, credits, clock confidence, scope, or required evidence. These are different causes and should have different recovery paths. Reporting the cause helps an operator choose whether to send new information, reallocate resources, repair the clock, or revise the task. A single generic failure score would obscure which assumption prevented action.

Long-horizon autonomy should therefore be described as bounded discretion with delayed oversight. That description is ambitious enough to support useful research while remaining compatible with the causal limits already proved.

A proposed comparative fault-injection study

The empirical program begins with a versioned discrete-event simulator and then, if justified, a controlled distributed implementation. The first study compares centralized commitment, replicated authority using a specified existing protocol, conventional lease-plus-escrow control, and GENESIS-D2. A deliberately unsafe cached-grant configuration is retained only as a diagnostic ablation. Every serious comparator receives the same workloads, physical resources, link traces, and permitted delegation policy.

The workload contains divisible compute jobs, bounded transfers, and irreversible commitments modeled as distinct action classes. Each request specifies release time, deadline, resource charge, scope, and utility. A reference evaluator knows the synthetic policy history and action consequences. Nodes know only their local state and delivered messages. This separation prevents the simulator from granting an algorithm hidden access to global current authority.

The scenario grid varies one-way delay, outage duration, demand imbalance, clock error, failure fraction, and credit distribution. Delay distributions include fixed values and bursty traces. Crash, restart, duplicate delivery, loss, reordering, and rollback attempts are separately tagged. Byzantine behavior is tested only for components whose comparator actually claims Byzantine tolerance; a trusted-monitor architecture is not credited with surviving arbitrary monitor compromise.

The main resource invariant is a hard outcome: any unaccounted overspend is a failure, with a minimized trace retained. Useful work completed before deadlines is the primary performance outcome among configurations meeting the invariant. After-revocation exposure, time to incorporate a new epoch, unavailable requests, stranded credits, message volume, reconciliation delay, and energy consumption are reported alongside it.

Centralized commitment is evaluated at its declared linearization point. If the task is remote physical execution, the authorization-to-actuation interval is included in the model rather than assumed away. A central service can have a clean commit log while a previously approved action occurs after a subsequent policy change.

All methods use fixed tuning budgets and separate development scenarios. The locked evaluation includes fault combinations absent from tuning. A positive result requires improvement against the conventional bounded-delegation baseline, not merely against the intentionally unbounded ablation.

Statistical estimands and stopping rules

An independent unit is a complete workload-and-fault scenario generated from a prespecified distribution, not an individual action inside that scenario. Methods run on matched traces with paired seeds. Let U_j^G and U_j^B be normalized useful work for GENESIS-D2 and the main baseline in scenario j. The primary estimand is $\Delta_U = \mathbb{E}[U_j^G - U_j^B]$, conditional on the declared scenario population and policy constraints.

The analysis reports a paired interval for Δ_U and separate distributions for exposure and recovery costs. A single weighted score must not hide violations of a hard invariant. Rare failures receive their own accounting. If independent identically distributed runs have zero failures in N trials, the exact one-sided 95% binomial upper bound is $1 - 0.05^{1/N}$. This bounds failure probability under that sampling model, not under all adversarial executions or arbitrary deployment shifts.

For illustration, reducing that zero-failure upper bound below 0.01 requires $N \geq 299$ independent runs. Those runs still cannot establish a universal safety theorem. Structured model checking and adversarial trace construction complement random sampling because a low-probability generated trace may be easy for an adversary to cause. Repeated seeds on an unchanged deterministic trace do not create additional independent scenarios.

Power and precision are planned from pilot variation. A rough paired mean calculation is $N \approx (z_{1-\alpha/2} + z_{1-\beta})^2 \sigma_D^2 / \delta_*^2$ for target improvement $\delta_* > 0$, error levels $\alpha, \beta \in (0, 1)$, and paired-score standard deviation σ_D. Here z_p is the standard-normal p-quantile. Final sample size is selected through simulation of the actual bounded, potentially skewed scenario-level outcomes. The pilot and model-selection set are excluded from the locked test.

One primary baseline and one primary performance estimand are preregistered. Secondary pairwise and interaction comparisons form a declared multiplicity family. Results are stratified by delay and fault regime because a mean gain can conceal severe failure in an important regime. Collection stops at the predetermined run count; any sequential design must specify its error control before outcome inspection. Inconclusive precision remains an inconclusive result, not proof that two architectures are equivalent.

Falsification and adversarial counterexamples

The architecture fails its resource claim if a permitted execution consumes more than the allocated total without a recorded replenishment. It fails its time claim if an old grant remains usable beyond the proved deadline while all stated clock and issuance assumptions hold. It fails its implementation claim if a recovery path bypasses a check that the proof assumed complete. Each failure should produce a short event trace showing the violated premise or transition.

Several adversarial cases are particularly informative. Replaying a delivered transfer identifier tests idempotence. Rebooting from an old snapshot tests durable counters and epoch memory. Delaying a transfer beyond a refund timeout tests whether both endpoints can spend the same credits. Compromising a signer tests threshold assumptions, while compromising the enforcing actuator tests a different trust boundary. Those cases should not be combined into one undifferentiated fault rate.

A clock adversary can reveal a false guarantee quickly. If a recipient's clock stops while the issuer's advances, a simple local expiry check may remain satisfied indefinitely. The lease theorem no longer applies because the bounded-error premise fails. A complete design must detect loss of its clock guarantee or use a more conservative mechanism. A theorem about clocks that behave correctly is not evidence that real clocks will always do so.

An authority fork is another distinct threat. Two committees with different configuration identifiers may each issue internally valid certificates. If a resource monitor accepts both as spending from the same unpartitioned budget, quorum intersection within each committee does not prevent overallocation. Reconfiguration must preserve the resource ledger and authority transition, including in-flight obligations.

The performance thesis is falsified for the tested setting if GENESIS-D2 offers no meaningful gain over ordinary lease-plus-escrow control at matched budgets, or if its benefit comes only from looser permissions. A simpler competitor may be the right design. Such a finding would leave the finite bounds intact while removing any established reason to adopt the proposed composition.

Identity, migration, and reconciliation

A civilization-scale architecture inevitably raises continuity questions, but an identifier and a resource ledger answer only part of them. Migrating a process can preserve a signed history while changing its implementation, memories, or behavior. Copying a process can create two descendants with the same initial record. Neither event automatically establishes one continuing subject, and neither should duplicate the original spending authority.

GENESIS-D2 treats migration as a transfer between enforcement domains. The old domain must stop spending or retain only a declared partition of the budget. The new domain receives a distinct identifier and a verifiable allocation transition. If the old domain cannot be contacted, the architecture may reserve the worst-case outstanding amount instead of assuming it has ceased to exist. This can reduce availability; it preserves the accounting premise.

Reconciliation has a modest meaning here: compare recorded events, validate unique identifiers and transitions, settle known resource obligations, and identify unresolved conflicts. It cannot undo a completed physical event or make contradictory histories both true. An authenticated report can still be mistaken or malicious. Evidence about an action's effects must be evaluated separately from evidence that a particular key signed a record.

Conflicts require a policy for what happens next. A remote group may have acted within a previously granted mandate while the issuing community changed its values. An engineering protocol can describe what was known and which ceilings applied. It cannot derive the political legitimacy of one community's decision from a quorum size or an energy budget. Those judgments require institutions, reasons, and representation of affected parties.

The architecture should therefore preserve provenance without promising a single timeless identity. It can show which policy authorized a particular branch, which resources were assigned, and which obligations remained unresolved. It should not erase a branch merely because a later authority prefers a cleaner history. Accurate records support accountability precisely when coordination was incomplete and participants made consequential decisions under different information.

The civilizational question remains open

The Continuum's long-horizon vision is strongest when it treats distance as a source of institutional design requirements. A remote site cannot consult a center about every urgent local decision if consultation takes longer than the decision permits. Some discretion must be delegated, some actions deferred, and some tasks refused. The relevant question is what limits and revision processes make that discretion acceptable.

Physical abundance does not settle governance. More energy may permit more computation and communication under an appropriate engineering model. It does not specify whose objectives should be optimized, whose consent matters, or which commitments should survive institutional change. Treating a larger power scale as a higher moral rank would add a claim unsupported by the physical classification.

Likewise, a globally coherent narrative is not the same as globally synchronized knowledge. Sites can share a charter while having different recent histories. They can agree on a reconciliation procedure without agreeing on every interpretation. The finite model encourages a plural but inspectable architecture: local decisions carry explicit scope, resources, and evidence rather than pretending that all participants occupy one present moment.

There are limits to this abstraction. Real long-lived infrastructure faces hardware degradation, environmental change, key compromise, changing language, institutional succession, and uncertain resource models. The paper does not estimate those processes across centuries. A finite protocol that preserves an invariant for 120 simulated seconds is not validated for cosmic timescales. Extending the horizon requires new evidence and a model of how the premises are maintained.

The bold research program is nevertheless concrete. Investigate how much useful work can remain decentralized while specific resource and authority constraints remain enforceable. Determine which constraints require communication, which can be reserved in advance, and which become infeasible under the available clocks and channels. Study the cost of correcting mistakes after long delays. Those questions can be tested in ordinary distributed systems today, without first claiming that a Kardashev-VII civilization is a physically established possibility.

Guide discussion and the next engineering decision

A Guide's explanation

A Guide can say: "You can't revoke a permission faster than light. So a civilization that spreads out has to decide in advance how much its distant members may do while news is in transit—expiring mandates, counted allowances, witnessed records. GENESIS is a small, testable version of that idea." A Guide should not present Kardashev-VII as an established scientific category, imply that such a civilization or architecture exists, or suggest that permission can travel faster than its signal.

Established results

The original Kardashev paper contains three civilization types, not a verified type VII [2]. Established physics constrains transmission and computation under specified models [3–7,16]. Distributed-systems research identifies timing, consistency, and fault assumptions that matter for agreement [8–12]. This paper proves finite bounds and reports executed synthetic checks. Its results do not establish an operating cosmic architecture.

The proposal

GENESIS-D2 combines scoped grants, time limits, resource ceilings, and reconciliation in an explicitly finite model. Its causal result explains why an unpredictable distant revocation cannot be known immediately. Its lease bound quantifies a possible stale-authority window. Its escrow invariant limits total accounted consumption under complete mediation. Its quorum result applies to fixed configurations with an honest nonconflicting-signature rule. Each claim carries those premises locally.

Useful discussion examples

A Guide can consider a disconnected research station with a finite supply budget and a mandate that may change. Which work should continue? What amount may be consumed before contact returns? What happens if a message arrives after a transfer was presumed lost? These questions expose the model's consequences without suggesting that permission travels faster than its signal.

When discussing the numeral VII, the Guide should identify it as the Continuum's extrapolative label. When discussing a claimed universal processor, it should ask what operation is counted, which physical model applies, and whether the number is a limit or a measured capability. When discussing continuity, it should distinguish a preserved record, a delegated role, and a claim about personal identity.

The next engineering decision is to implement the same finite contract in a fault-injection harness alongside conventional lease-plus-escrow and replicated-control baselines. Preserve all counterexample traces, measure useful work only under matched permissions, and validate the recovery assumptions separately. If the composition supplies no measurable advantage, retain the clearer conventional design. If it helps, publish the precise regime in which it helps. A distant intelligence becomes accountable through explicit obligations that survive delay: a bounded lease, a conserved allowance, a durable record and a defined route for revision when contact returns.

References and source scope

[**1**] Architect X13. *Eidara Continuum Bible*, supplied canonical edition, Section 2.8, PDF p. 12.

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

Scope: Tier II vision taxonomy. Public URL is research context, not a copy of the private Bible.

[**2**] Kardashev, N. S. (1964). Transmission of information by extraterrestrial civilizations. *Soviet Astronomy* 8, 217–221.

https://technosearch.seti.org/wp-content/uploads/2018-09/Kardashev_CTA102.pdf

Scope: Original translated paper, pp.217–218 inspected visually; types I–III.

[**3**] Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal* 27, 379–423, 623–656.

https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

Scope: Original corrected reprint, opening sections; information and channel modeling, not semantic truth.

[**4**] Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development* 5, 183–191.

https://www.cpt.univ-mrs.fr/~verga/pdfs/Landauer-1961uq.pdf

Scope: Original article reprinted in 2000; opening two pages inspected visually.

[**5**] Bennett, C. H. (1973). Logical reversibility of computation. *IBM Journal of Research and Development* 17, 525–532.

https://worrydream.com/refs/Bennett_1973_-_Logical_Reversibility_of_Computation.pdf

Scope: Original abstract and introduction; reversible computation and erasure distinction.

[**6**] Margolus, N., and Levitin, L. B. (1998). The maximum speed of dynamical evolution. *Physica D* 120, 188–195.

https://arxiv.org/abs/quant-ph/9710043

Scope: Published-version author manuscript, introduction and orthogonalization derivation.

[**7**] Lloyd, S. (2000). Ultimate physical limits to computation. *Nature* 406, 1047–1054. DOI: 10.1038/35023282.

https://arxiv.org/abs/quant-ph/9908043

Scope: Author manuscript opening and operation-speed sections; limits are not construction guarantees.

[**8**] Lamport, L. (1978). Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM* 21, 558–565.

https://lamport.azurewebsites.net/pubs/time-clocks.pdf

Scope: Original opening and happened-before definition; logical order does not eliminate delay.

References and source scope, continued

[9] Fischer, M. J., Lynch, N. A., and Paterson, M. S. (1985). Impossibility of distributed consensus with one faulty process. *Journal of the ACM* 32, 374–382.

https://groups.csail.mit.edu/tds/papers/Lynch/jacm85.pdf

Scope: Original abstract and model discussion; deterministic fully asynchronous termination limit.

[10] Dwork, C., Lynch, N., and Stockmeyer, L. (1988). Consensus in the presence of partial synchrony. *Journal of the ACM* 35, 288–323.

https://groups.csail.mit.edu/tds/papers/Lynch/jacm88.pdf

Scope: Original abstract and timing definitions; explicit assumptions for progress.

[11] Gilbert, S., and Lynch, N. (2002). Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services. *ACM SIGACT News* 33(2), 51–59.

https://www.comp.nus.edu.sg/~gilbert/pubs/BrewersConjecture-SigAct.pdf

Scope: Author manuscript definitions and impossibility proof; atomic consistency and model-specific availability.

[12] Castro, M., and Liskov, B. (1999). Practical Byzantine fault tolerance. *Proceedings of OSDI*, 173–186.

https://www.cs.cornell.edu/courses/cs614/2004sp/papers/CL99.pdf

Scope: Original system model, service properties, and safety/liveness discussion.

[13] Gray, C. G., and Cheriton, D. R. (1989). Leases: An efficient fault-tolerant mechanism for distributed file cache consistency. *Proceedings of SOSP*, 202–210.

https://www.cs.cmu.edu/afs/cs.cmu.edu/academic/class/15712-s12/www/papers/gray89.pdf

Scope: Original abstract and fault-tolerance section; well-behaved clocks are required.

[14] O'Neil, P. E. (1986). The escrow transactional method. *ACM Transactions on Database Systems* 11, 405–430.

https://cs.umb.edu/~poneil/EscrowTM.pdf

Scope: Original abstract and Section1.4; reserved resources and specialized updates.

[15] Bailis, P., Fekete, A., Franklin, M. J., Ghodsi, A., Hellerstein, J. M., and Stoica, I. (2014). Coordination avoidance in database systems. *Proceedings of the VLDB Endowment* 8, 185–196.

https://www.vldb.org/pvldb/vol8/p185-bailis.pdf

Scope: Primary abstract and introduction; invariant-based coordination is established prior art.

[16] Bureau International des Poids et Mesures. *The International System of Units: Defining constants*. Accessed November 2026.

https://www.bipm.org/en/measurement-units/si-defining-constants

Scope: Official standard values for c, h, and k_B; additional to the scholarly references.