


01 THE THESIS

An economy for accountable intelligence.

Pay for useful judgment. Preserve the evidence. Reward the correction that makes the next decision better.



Conceptual service loop / proposed design

The problem

Models can be accurate on average and unreliable in the moment that matters. A buyer needs more than an answer: who produced it, what was knowable then, how the method has behaved and how the outcome will be checked.

The proposition

ZenoDivergent learns the behavior of forecasters. The proposed ZENO economy turns that research into defined services: reliability cards, evaluation, monitoring and, where independently validated, combined forecasts. Agents and teams could buy a useful decision with a traceable limit.

The unit of value

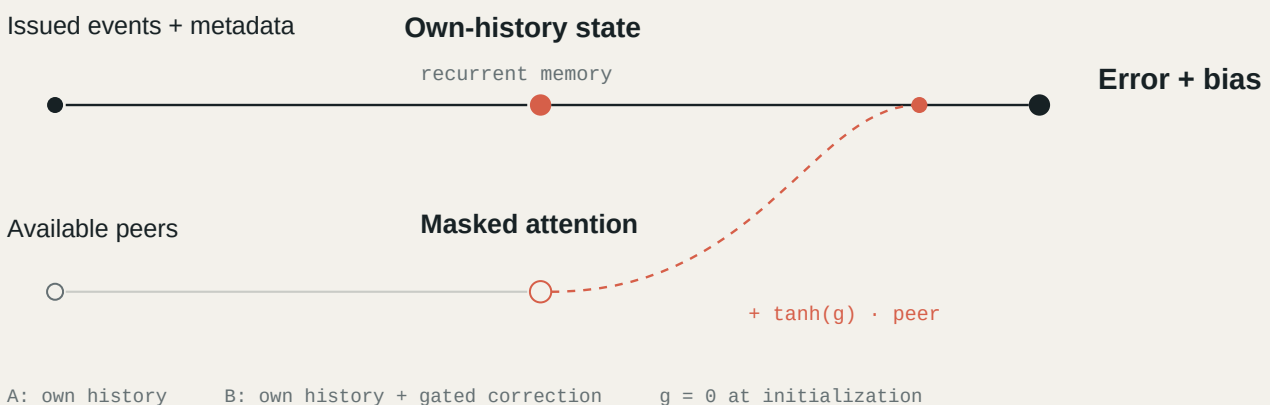
A decision contract fixes the task, clock, method, baseline, outcome policy and payment terms before resolution. Providers deliver; evaluators reproduce the score; accepted challengers can earn for a correction, including one that disproves a Zeno claim.

notflat.world / evidence offdiagonal.space / receipts ZenoDivergent / learner

Design proposal v0.2. Research code and training exist. The service market, operator bonds and automated outcome settlement described here remain proposed.

Learn the forecaster. Keep the clock.

The population model asks how wrong a forecaster may be before its next outcome arrives. Forecast combination is a separate objective. [1, 2]



A / own history

An event encoder and metadata feed persistent recurrent state for each forecaster stream. Arm A uses that history to predict future squared error and signed bias. Earlier resolved targets may inform the state only when their outcomes were actually available at the decision time.

A causal boundary

Feature selection uses information available by the decision. Label selection follows a separately declared outcome-vintage rule. Later revisions append evidence; they do not silently enter earlier histories. Missing distributions or context stay missing.

B / conditional peer information

Arm B adds attention over eligible peers through $h_B = h_{\text{own}} + \tanh(g) \times \text{peer_delta}$. Copied shared weights and a zero initial gate preserve the own-history starting representation. The gate allows a correction; it does not guarantee useful peers or automatic rejection of noise.

A capability boundary

Only heads receiving the intended supervision support trained-capability claims. Co-failure is inactive unless labels reach its loss. Error prediction, direct forecast combination, peer value and transfer each need their own experiment.

Architecture at the reviewed commit is a research implementation. The diagram is a schematic of the active A/B idea, not evidence of a performance gain.

03 MEASURE

A completed fit is the beginning.

Record progress without turning it into an accuracy claim. The first crypto experiment and the next campaign have different evidence states.

First crypto experiment

The owner reports crypto-pop-20260920-1547 complete. GitHub head 052c4a9 is titled “Completed Crypto A/B seeds.” The reviewed materials identify replayed one-day forecasts and A/B arms with three requested seeds. A final remote report was not retrieved for this revision; no final metric or checkpoint count is asserted. [3]


The next campaign

The committed plan calls for constant, recent-error, historical-neural and tabular controls; cold-start and capacity comparisons; peer diagnostics; and independent supply gates. It also separates development selection from fresh confirmation and gives forecast combination its own executable and loss. These are planned requirements, not verified completed repairs. [4]

What the existing period can show

The reviewed population runner scores evaluation before selecting on development. That period is inspected retrospective evaluation. This does not by itself establish training-label leakage, but a post-hoc companion baseline report cannot restore an untouched confirmation period. [2]

REQUIRED NEXT-RUN ORDER



Error prediction

Beat matched error-prediction controls.

Combination

Beat equal weighting and a fitted linear combiner.

Relational value

Beat own history with capacity and coverage controlled.

Transfer

Succeed on declared unseen forecasters or families.


Report task-native scores, identical case IDs, chronological support, all finalist seeds and objective-aligned uncertainty. Seeds measure restart variability, not independent histories.

04 EXCHANGE

Pay for the service. Fund the proof.

Separate the work of delivering an answer from the later question of whether it improved a decision.

ILLUSTRATIVE BUDGET / 20 USDC



Fees pay for delivery. Skill bonuses wait for their declared evaluation block.

Before the outcome

The contract names target and units, horizon, issue deadline, permitted inputs, candidates, baseline, score, outcome source and vintage, resolution window, fees and disputes. Coverage and abstention costs are explicit; always abstaining cannot qualify as useful performance.

A bounded example

A buyer funds 12 USDC for service delivery and 8 USDC for a bonus pool. After the batch resolves, the locked rule pays at most 8 USDC in bonuses; unused bonus funds return under the example terms. Evaluator and protocol fees are disclosed in the quote. This is illustrative accounting, not launch pricing.

After delivery

A quoted fee pays for a valid response and receipt. An optional bonus pool is prefunded and settles after its declared evaluation block matures. Minimum support, caps, ties, missing outcomes and refunds are fixed in advance. One lucky observation is not a general claim of skill.

Correction is productive work

A reproducible clock, lineage or scoring defect can earn a capped challenge bounty. The resolver publishes the decision and appeal path. Contribution rewards may use a predeclared inclusion/exclusion comparison, but correlated or copied submissions complicate attribution. Honest forecast misses are scored; they are not automatically slashable misconduct.

Fees and bounded funded pools cover service obligations. Dollar-denominated promises must not rely on token appreciation.

Three roles. Three kinds of value.

The proposed token supports accountable participation. A customer should be able to buy a useful service without first buying a volatile asset.

USDC

Pays for work

Service, compute and evaluation budgets use stable units. A payment does not establish accuracy.

\$ZENO

Bonds commitments

Proposed bonds cover specified delivery, identity and integrity duties. Ordinary forecast misses are not fraud.

REPUTATION

Records performance

Earned by task, horizon, support and date. Holdings cannot buy a score, forecast weight or a factual verdict.

Bond and governance rules

Publish bond size, withdrawal delay, breach evidence, appeal rules and custody before bonding. Token collateral alone cannot guarantee fixed-dollar refunds. Later governance could allocate capped research or challenge budgets; it cannot rewrite outcomes or override data rights. Initial operation is centralized unless shown otherwise.

Issuance policy to specify

Proposed constraints are a disclosed maximum supply, published allocations, on-chain vesting and a budgeted rewards reserve. Quantities and schedules remain unset. Report emissions separately from customer revenue and measure the useful work they fund.

Rights and the necessity test

This draft grants no equity, model ownership, redemption, revenue share or promised return. A native token must justify its added complexity through operator demand and enforceable commitments. Its necessity is not established by naming these possible roles.

Proposed identity

ZENO / \$ZENO · Solana

Not finalized

Issuer, mint, token program, decimals and supply

Not finalized

Allocation, vesting, distribution and liquidity custody

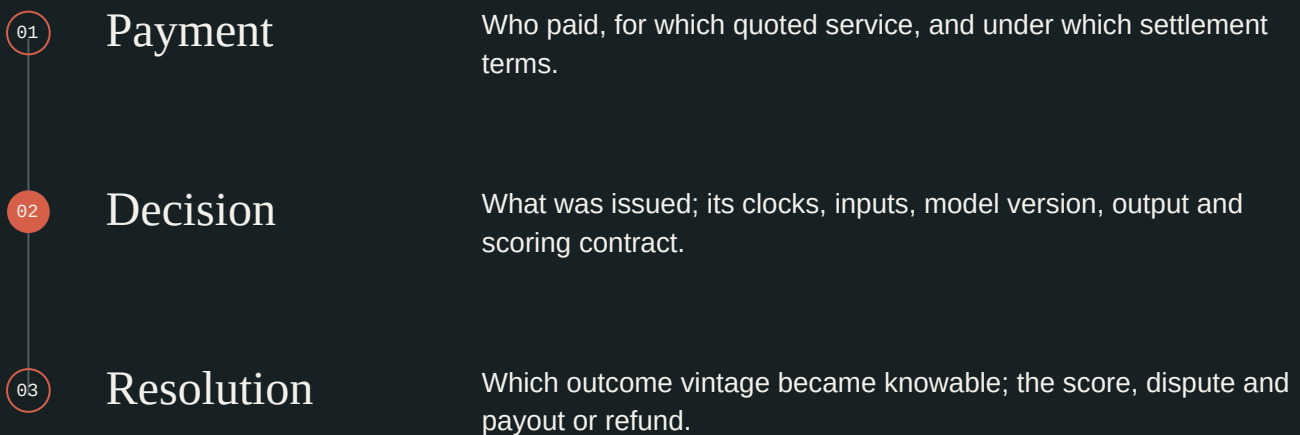
Not finalized

Authorities, treasury, fees, live utility and holder rights

Token terms require a separate authoritative release record. A symbol is not a unique token identifier.

Three receipts. One accountable chain.

Link the payment to the claim, and the claim to its resolution. No single transaction hash proves all three.



x402 / access

x402 supports HTTP payment negotiation. A USDC-priced Solana endpoint is a development target; test the selected network, token and facilitator together. Signed offer-receipt support can add interaction evidence where available. Bind the decision ID and payload digest at the application layer. [5, 10]

Solana / commitments

Keep models and licensed evidence off-chain. A proposed registry can link operator identity, receipt digests, bonds and funded settlement terms. Digests establish integrity relative to committed bytes, not truth or continued data availability. Authorities and the initial resolver remain explicit trust assumptions. [6]

Build / in order

Begin with a reproducible decision card and paid sandbox endpoint. Add retry-safe accounting and independent resolution. Test delayed bonuses and disputes with capped balances before bonds. x402 does not supply performance escrow or arbitration; those need separate logic. No audited application or live token utility is claimed here.

Outcome corrections are append-only. Settlement follows the contract version, not a silently substituted later print.

Let the evidence set the claim.

The public commitment remains: a diploma before a token. A named lane, baseline, support and result, with losses included.

Research proof

Publish the contract, frozen inputs, code, splits, matched cases, baselines, all selected seeds, uncertainty and limitations. Identify replayed issuance and already inspected periods. A later positive claim needs a valid later test.

Product and release proof

Show a useful card, reproducible resolution, working payment and failure handling. Resolve issuer, token terms, custody, live rights and the applicable distribution requirements. This design paper is not a completed regulatory disclosure or filing. [9]

Risks and related work

Weak independent support, drift, multiple testing, reward farming, resolver error, program defects and token volatility remain material. Upstream rights still govern data use. Numerai and Allora already address contribution or inference synthesis; Zeno makes no first-ever claim. Its proposed emphasis is provenance linked to paid, reproducible correction. [7, 8]

SOURCE NOTES

01 Population model

06 Solana payment integration

02 Population runner

07 Numerai contribution

03 Completion-titled commit

08 Allora inference synthesis

04 Next-campaign plan

09 EU Regulation 2023/1114

05 x402 protocol and receipts

10 x402 offer-receipt extension

Evidence snapshot: 20 September 2026. Research design, a useful service and a useful token are separate obligations.