



KNOWLEDGECOIN WHITE PAPER

Knowledgecoin: Adjudicated Knowledge in the Age of AI

A claim with a proof is still only a claim. Knowledge is a claim that has been tested by parties with something to lose, with the test on the record.

“In a time of universal deceit, telling the truth is a revolutionary act.” — attributed to George Orwell

Version 2.0.0 October 2026 Gleason · Kimball · Repetti · Straughn Paradatum Inc.

CONTENTS

Abstract

PART 1 · THE PERFECT LIBRARY

- 1.** Knowledge is how humans survive
- 2.** The Iron Law of Herds
- 3.** The Perfect Library: five criteria
- 4.** Mad Librarians, Evil Librarians, and the Epistemic Agent Problem

PART 2 · THE LIBRARIAN IS A MACHINE

- 5.** The Librarian is now a machine
- 6.** What a ledger cannot do
- 7.** The five criteria, re-solved

PART 3 · KNOWLEDGECOIN

- 8.** Where Knowledgecoin sits
 - 9.** How a claim becomes knowledge
 - 10.** Proof-of-Knowledge, taken literally
 - 11.** The island's industries become jobs
 - 12.** Why adjudication is tractable here
 - 13.** The Know Tribe
 - 14.** Knowledgecoin Bill of Rights
 - 15.** Team
 - 16.** Conclusion
- References
- A.** Mechanics and status

Abstract

Knowledgecoin is an adjudication layer: the place where a claim that already carries a proof is tested by parties with something to lose, and the test goes on the record. The layers beneath Knowledgecoin, which run the computation and file the claim, are described in Part 3; this paper is about what happens above them.

The Knowledgecoin position paper of December 2021 argued that knowledge is humanity's primary means of survival, that whoever controls the library controls the mind, and that a Perfect Library must satisfy five criteria: Trust, Space, Time, Mutuality, and Parsimony. That paper then proposed a blockchain and a human jury as the mechanism. Five years on, the diagnosis is stronger and the mechanism was incomplete. A ledger makes a record immutable; a ledger does not make a claim valid. Validity was deferred to a human jury inside a decentralized autonomous organization (DAO), and a human jury neither scales to the rate at which claims now arrive nor resists capture.

In 2026 the Librarian most people consult first is a large language model, and the three Librarian problems of the 2021 paper have exact machine forms. The Mad Librarian is hallucination. The Evil Librarian is the silent update. The Epistemic Agent Problem is the absorption of the world's epistemic labor into model weights with no reward to the agents who produced it.

The answer is also mechanical now. Deterministic execution gives an agreed computation one canonical result on any hardware. A proof makes that result checkable in milliseconds by anyone, with nobody redoing the work. Masked computation lets a question be asked without the machines doing the work ever reading the question. On that substrate, Knowledgecoin's Proof-of-Knowledge stops being an analogy to Proof-of-Work and becomes a literal object: a Knowledge Record that binds a claim, its evidence, the pinned model that judged the claim, the verdict, the proof of the run, and the stakes of every party who adjudicated.

This paper keeps the 2021 argument in Part 1, restates the problem for the age of AI in Part 2, and sets out the adjudication layer in Part 3. Token mechanics, the attack table, and the status of every component live in the appendix. Each mechanism carries a status label: measured, prototype, designed, or pending.

Proof settles what ran. Adjudication settles what it means.

PART 1 · THE PERFECT LIBRARY

1 Knowledge is how humans survive

Wolves have fangs and cheetahs run fast. Humans survive and flourish by acquiring knowledge. The actions that drive value are intellectual first: every rational action begins as an idea in a mind. Acquiring knowledge carries large upfront costs, and whenever knowledge is lost, energy must be spent to regain it. So every society builds schemes for storing and sharing what it knows: counting, language, writing, papyrus, books, the printing press, libraries, radio, television, the internet.

The library deserves particular attention, because the library is the model most of us relied on before the internet. A library centralizes knowledge: collected, analyzed, catalogued, stored, shared. When the storage medium is expensive, one shared copy for a whole community is an efficient arrangement.

But knowledge is power, as the proverb has it, and those who hold power learn quickly that while beliefs can be managed, raw knowledge is dangerous to leave in the hands of the governed. An armory centralizes the supply of weapons. A bank centralizes the supply of money. A library centralizes the supply of knowledge. The originators of these institutions may not have intended the abuse, but it is nearly inevitable that power-seeking agents will seize such easily grasped controls. A society's information is the equivalent of its mind: seize the mind, and the body follows.

Want to subjugate a people? Remove their access to validated knowledge. Want to hobble a competitor? Make them relearn what you already know, again and again, while your knowledge alone advances. This is why enslaved peoples were forbidden to read, why thought-provoking books have been burned, and why some of the world's greatest libraries were destroyed.

When outlawing reading is no longer feasible, the would-be tyrant takes the next route: make the populace doubt that real knowledge is possible at all. There are two steps, and the first enables the second. First, flood every channel with contradictory, confusing information, all presented with equal confidence. Second, as the necessary consequence, let skepticism do the rest, since everything now looks like propaganda, spin, assertion, sophistry, or dogma.

Epistemology is the branch of philosophy that studies what knowledge is, what criteria a belief must satisfy to count as knowledge, and by what means knowledge is produced, assessed, and confirmed. Each of us is, ideally, an **epistemic agent**: a seeker, consumer, producer, validator, and dispenser of knowledge. In a tidal wave of partly convincing contradictions, any epistemic agent can be made to doubt any assertion, and would-be authorities no longer need to say "who are you to have an opinion on this?" The overload says it for them.

Reliable access to full knowledge is a requirement of a free, fair, and functioning society. For informed people to decide rationally, they need all the relevant data, and they need the data without first having to curate the data out of a sea of contradictory claims, hyperbole, spin, propaganda, and lies. The requirement bites hardest on exactly the knowledge that is unpopular, or that powerful interests are actively suppressing.

2 The Iron Law of Herds

When humans cooperate, they are eager to educate each other. When humans compete, their inclination to teach becomes miserly, and knowledge that confers an advantage is not shared. Between tribes, knowledge trading can be non-existent: why help individuals who might compete against you or your progeny? The Okinawan karate masters left critical secrets out of the lessons they gave to Japanese and American outsiders. The affable founder of Kentucky Fried Chicken would have happily instructed a stranger on a train in any subject except the eleven herbs and spices.

This reluctance to share does not stop people from debating their hobby horse, and the debates reveal the mechanics of tribal membership. As the informational landscape becomes opaque, outsourcing one's knowledge needs to trusted familiars is the natural move, and so tribalism rises. The litmus tests of the day sort people into the two fundamental groups, *us* (the trusted tribe) and *them* (the untrustworthy tribe). Such conversations rarely resolve, because acknowledging the slightest point of compromise counts as betrayal. Each side goes through the motions of exposing its opponents to the substitutes for inquiry: belittling, mocking, dismissing, mischaracterizing, insulting, exaggerating, villainizing, slandering, bullying, threatening. The leaders who prompt these exchanges have already achieved their objective, which is to entrench their flocks more deeply in isolated silos.

The herd instinct is, in general, a feature. In our present epistemic predicament the herd instinct functions as a bug: the instinct divides us into ever more polarized groups, fragments knowledge, undermines our autonomy as individual knowledge agents, and clouds the truth. The Iron Law of Herds is that every tribe, left to itself, will trade knowledge inward and withhold knowledge outward.

What all of this calls for is an anti-tribal community of increasingly empowered epistemic agents whose loyalty is to truth regardless of tribal commitment: a knowledge tribe, or for short a **Know Tribe**, which is no tribe at all. Part 3 returns to the Know Tribe with a sharper definition than the 2021 paper could offer.

3 The Perfect Library: five criteria

Consider a thought experiment. What would a Perfect Library be like? A Perfect Library would be a knowledge resource where authenticated epistemic agents could rely on four things: that they can both submit to and draw on all relevant, validated knowledge, curated from everything available with misinformation filtered out; that the library is open to any and all people; that the library is equally accessible from any location at any time; and that the library guarantees the validity and completeness of its holdings in a way anyone can verify for themselves, secure in the belief that the knowledge is unalterable except by valid, audited improvement.

In more rigorous terms, a Perfect Library must satisfy five criteria. Trust, Space, Time, and Mutuality are familiar concepts from distributed systems. Parsimony is the contribution specific to Knowledgecoin.

Trust

Can users verify that the library's holdings on a topic are the authentic materials, and that they are complete? Identity, in the logical sense: a thing is identical to itself, and a copy is identical to the original.

Space

Can every epistemic agent reach the whole library from wherever the agent and the holdings happen to be, as producer and as consumer?

Time

Can users trust that holdings do not change, and look back to verify that every change in the history was truth-preserving, applied by proper epistemic procedure, with better evidence than the evidence first relied upon?

Mutuality

Are all of the library's agents connected to all of the providers they wish to reach, or are some excluded, throttled, shadow-banned, or siloed? A Perfect Library is also a maximally intra-communicative network.

Parsimony

Do users get the minimum required data to reach a rational conclusion: neither missing data nor mountains of irrelevant material? Grice's maxim of quantity, applied to a library: as much as is required, and no more. [1]

Meeting these five challenges was nearly impossible for most of human existence. The earliest libraries depended on elite scribes copying papyrus to papyrus; few could read the holdings, fewer could reach them, and fewer still could tell whether the content survived each reproduction, to say nothing of the societal strokes of censorship, book-burning, and the destruction of whole libraries. Modern libraries addressed some of these concerns, and the worldwide web offered the potential to address all five. Space could be mitigated by devices streaming data anywhere. Mutuality could be mitigated by platforms connecting interested

parties. Trust and Time could be mitigated by anonymous trust via the magic of encryption and hash algorithms. Parsimony could be managed, to an extent, by careful and honest search. The Perfect Library was finally within reach, except for one fatal flaw.

4 Mad Librarians, Evil Librarians, and the Epistemic Agent Problem

For a brief time humanity seemed close to achieving the Perfect Library. We did not attain the Perfect Library, and we are in a far worse predicament than mere failure: a would-be Perfect Library that appears to be the real thing, overrun by a cartel of Mad and Evil Librarians. Far worse than the ignorance caused by having no repository of knowledge is a corrupt repository posing as a Perfect Library.

The clearest example of the aspiration is a search company whose mission statement is to organize the world's information and make it universally accessible and useful. Read charitably, “the world's information” means all of the information; “organized” means structured; “universally accessible” means everyone has equal access regardless of who they are. That is a stated commitment to be the Perfect Library, made by an institution with the power to attempt it. The same can be said of the media, academia, and government, each of which has had the power to help create a Perfect Library and each of which claims to be committed to that end. Together they have devolved into something closer to a cartel acting against the Perfect Library. The rest of this section explains how, through three analogies.

The Mad Librarian problem: gaps make full validation impossible

Suppose a Mad Librarian secretly burns random books for idiosyncratic reasons. Suppose also that we find an exactly twenty-book consensus on whether there was ever water on Mars. Is that a scientific consensus, or does the consensus only appear so because the Mad Librarian burned the books that disagreed? If we cannot rule out the possibility, every alleged consensus is cast into doubt, whether we realize it or not. To be sure, we would need to know that the library had procedures for detecting Mad Librarian activity.

The digital library makes this easier, not harder, because books can be eradicated without fire, smoke, or ashes. The digital equivalent of whole sections burning is found in closed societies that block websites for their entire population; search for “Tiananmen Square” from a hotel in Beijing to experience the phenomenon directly. Users who do not know something is missing, or what is missing, cannot be confident about the representativeness of what is present.

The Mad Librarian has two further moves. The opposite malfeasance is addition: if phony or altered books are admitted and catalogued beside the valid ones, users cannot tell valid from invalid, and one cannot properly evaluate Newton's *Principia* if the copy contains phony edits.

And a Mad Librarian who fails to index uniformly, leaving some books on random shelves, in the back room, or shuffled among properly indexed ones, forces a search of the entire inventory every time. A single proven falsehood from a witness undermines every other statement the witness makes: *falsus in uno, falsus in omnibus*. The Mad Librarian injects a falsehood, or many, somewhere we know not, and thereby undermines the entire library.

Note that the Mad Librarian need not bar anyone from the premises. Being publicly barred at least has the virtue of transparency. The Mad Librarian instead burns, fails to index, reshuffles, and inserts, while allowing everyone equal entry and pretending that all knowledge is equally available.

The Evil Librarian problem: active sabotage

Malicious intent is not required for a Mad Librarian to wreak havoc. The chaotic nature of madness limits the damage on any one topic. Matters multiply when the same undermining mechanisms are seized by a smart, dedicated Evil Librarian whose goal is to manipulate the knowledge base to specific ends. Compare a lazy paralegal who misfiles briefs and shreds the wrong documents with a clever, diligent, evil paralegal dedicated to industrial sabotage: the first does limited damage that will be discovered in due course; the second does incalculable damage that may take years to uncover, if the firm survives at all.

The danger is acknowledged across tribes. Defenders of socialism and of free markets alike claim that power-hungry individuals manipulate the ideology to seize power, suppress knowledge, and spread misinformation. We leave the merits aside; the point is structural. The ideology of the Perfect Library is that all sides wish for complete, accurate, unbiased information. When a power-seeking group takes control of the library, the Perfect Library becomes the central mechanism of manipulation.

The Epistemic Agent Problem: who deserves the rewards for user-generated value?

Anyone who likes a post, comments, shares an article, or rants on a blog is creating some human knowledge, albeit at the low end of the spectrum. Anyone who performs a successful search is creating search knowledge, somewhere in the middle. Anyone with competence, fluency, or expertise on a subject sits higher, and the position has nothing to do with credentials and everything to do with the evidence, analysis, and insight the agent can bring to bear.

The platforms that host this activity track every search, analyze the trail, and feed the successful paths back into their ranking algorithms. A user performing a search kisses a lot of frogs before finding the prince, and the platform alone sells the pathway to the prince again and again, thanks to the frog-kisser's eventually successful research. It seems reasonable for a search engine to be rewarded for the value of its algorithm. It seems unjust for the epistemic value of each user's

search to be reaped solely by the platform rather than shared with the user who generated the value. We call this imbalance, in which the reward for epistemic value is hoarded by the platform rather than shared with the person who added the value, **the Epistemic Agent Problem**.

The large institutions that could have built the Perfect Library do solve some aspects of the five criteria. Figure 1 scores the main categories of today's knowledge platforms. Named platforms within each category differ in detail, but the criteria are structural, and the pattern holds across the category.

PLATFORM CATEGORY	TRUST	SPACE	TIME	MUTUALITY	PARSIMONY
Hosted AI assistants	NO which model, version, or reasoning produced the answer is unknowable	YES where not geo-blocked	NO silent updates; yesterday's answer cannot be recovered	NO one-to-one; no shared, checkable record	PARTLY concise, but what was left out is unknowable
Search engines	NO ranking is private; completeness unverifiable	YES	NO rankings change without a history	NO private index; nobody can build another	NO mountains of results
Social platforms	NO authentication and labels follow platform policy	YES	NO edits and deletions leave no public history	PARTLY the feed decides who is reached	NO engagement, not quantity, is the goal
Crowd-edited encyclopedias	PARTLY sources cited; editors anonymous; edit wars decide	YES	PARTLY public revision history, uncommitted; reversions routine	PARTLY talk pages; participation gated by norms	PARTLY concise; completeness unverifiable
Knowledgecoin DESIGNED	YES pinned model, content-addressed contract, proof of the run	YES public ledger and public store	YES content-addressed and anchored; history cannot be rewritten	YES a shared verifier anyone can run	YES retrieve and re-verify; the record states its own evidence

Figure 1. Knowledge platforms measured against the five criteria. The Knowledgecoin row describes the design in Part 3 and carries its status label.

The greatest enemy of knowledge is not ignorance. The greatest enemy is an Evil Librarian who gives you access to censored, biased, phony information while maintaining the illusion that the information is validated knowledge, and who profits from every effort you make to overcome the obstacles placed in your path.

PART 2 · THE LIBRARIAN IS A MACHINE

5 The Librarian is now a machine

In 2021 the Librarian was a search engine. In 2026 the Librarian is a language model, and the three Librarian problems have exact machine forms.

The 2021 paper described the Mad Librarian as one who burns random books, inserts phony books, and fails to index. That is a description of a large language model. A model answers fluently about a paper that was never written, cites a case that does not exist, and cannot tell you what the model left out, because the model does not know. **The Mad Librarian is hallucination.** *Falsus in uno* applies to a machine exactly as it applies to a witness: one confabulated citation undermines confidence in every uncited claim around it.

The 2021 paper said malice is not required for the Evil Librarian, only capability and reach. A hosted assistant has layers nobody outside the operator can audit: the training mixture, the alignment tuning, the system prompt, the retrieval ranking, the safety filters, and above all the update. A well-meaning lab that ships a new checkpoint overnight has rewritten every book in the library at once, and nobody can get yesterday's answer back. **The Evil Librarian is the silent update.** The Time criterion of the Perfect Library fails by construction, not by misconduct, and the failure is planetary in scale because one Librarian now serves hundreds of millions of readers at once.

The 2021 paper said the platforms hoard the epistemic value of every successful search. The machine form is larger. The text, code, arguments, corrections, and conversations of the world's epistemic agents are absorbed into model weights, and the model is then rented back to the agents who produced the value. The frog-kisser of Part 1 no longer even sees the prince; the pathway is inside the weights. **The Epistemic Agent Problem is the training set.** This chapter of the 2021 argument is the one Big Tech has no answer to, and the one Knowledgecoin is built to answer: pay the agent whose work the record relies on, every time the record is relied on.

Privacy belongs here too. The 2022 Bill of Rights promised that research within the Knowledgecoin domain would be protected from eavesdropping, without a mechanism. A question put to a hosted assistant is readable by the infrastructure that answers the question. In *United States v. Heppner*, a federal court considered the provider's ability under its policies to retain, train on, and share chatbot conversations when assessing whether a defendant had a reasonable expectation of confidentiality. [2] The Librarian reads the reader.

And the Iron Law of Herds has a machine form: each lab trains its own Librarian on its own mixture, and the Librarians increasingly learn from one another's outputs. A tribe of models is still a tribe.

6 What a ledger cannot do

The 2021 paper proposed a blockchain and a decentralized autonomous organization as the mechanism. The ledger half of that proposal was right. Hashing a record and anchoring the hash gives Time: the history cannot be rewritten. Hashing gives half of Trust: a copy can be checked against the original. The DAO half was incomplete, for three reasons.

A ledger makes the record immutable; a ledger does not make the claim valid. Immutability guarantees that a false claim stays exactly as false as the day the claim was filed. Validity was deferred to a human jury of taggers, stewards, logicians, and reasoners earning tokens for their verdicts.

A human jury does not scale. Claims now arrive at machine speed. A jury that must read every claim is a bottleneck by design, and a bottleneck is where capture happens.

A jury that cannot be re-run cannot be audited. If a juror's verdict is an opinion, the ledger can record the opinion but cannot check the opinion. The 2021 Trust criterion asked for a trusted source; what the Perfect Library actually needs is a way to verify the *reasoning*, not only the record. The 2021 Parsimony criterion asserted “the perfect index” without a mechanism. The 2021 Mutuality criterion settled for “like-minded people,” which is a tribe waiting to form.

Three pieces were missing, and all three now exist as working technology in the Paradatum stack.

- **Determinism.** An agreed computation produces one canonical result, the same bytes, on any hardware. A content-addressed contract, the Portable Semantic Execution Contract (PSEC), pins the model, the arithmetic, and the serialization, so that “which model judged this claim” has an exact answer. MEASURED
- **Proof.** A zero-knowledge proof attests that the pinned model actually produced the output, and a third party checks the proof in milliseconds without rerunning the model. MEASURED
- **Privacy.** Masked computation lets a claim be made, and a question be asked, without the machines doing the work ever reading the plaintext. PROTOTYPE

The status labels above are defined in the appendix, which also lists the public evidence for each. The 2021 paper could only assert the Perfect Library. With these three pieces, each criterion gets a mechanism.

7 The five criteria, re-solved

CRITERION	2021 ANSWER	2026 MECHANISM
Trust	Hash the book and put the hash on the chain.	Verify the reasoning, not only the record. The contract pins the model and the math; the proof shows the pinned model produced the adjudication. Anyone checks the proof; nobody trusts the source.
Space	Copy the ledger everywhere.	Deterministic execution on any hardware: the same claim, the same bytes, re-verifiable on a laptop, a data-center GPU, or a CPU verifier. The record is as portable as the ledger.
Time	An immutable edit log.	A pinned model is a changeless Librarian. The content-addressed contract means the Librarian who judged a claim in one year can be consulted again, unchanged, in any later year. The silent update becomes impossible by construction.
Mutuality	A DAO of like-minded people.	A shared verifier. Humans and agents who check the same proof agree on what ran, whatever else they disagree about. A community defined by its verifier rather than by its beliefs is a Know Tribe, which is no tribe.
Parsimony	"The perfect index," asserted.	Retrieve and re-verify instead of recompute. A proven claim is a reusable fact: the same question asked again is answered from the record, with its proof re-checked, and no model runs. The record carries exactly its evidence and no more.

One item from the Bill of Rights also gets its mechanism. The right to private inquiry, promised in 2022, is met by masked computation: the question is never plaintext to the network, and the Librarian no longer reads the reader.

PART 3 · KNOWLEDGECOIN

8 Where Knowledgecoin sits

Knowledgecoin is one layer of a stack, and the division of labor between the layers is the whole design. Blindproof runs the job. DeFacts remembers the result as a claim. Knowledgecoin adjudicates the claim into knowledge. Paradatum is the gateway where customers pay in dollars and never hold a token.

The division of labor. Nobody in Knowledgecoin argues about what ran or what the source was; those questions are proven in the layers below. Adjudicators argue only about whether the claim holds. That is the smallest adjudication surface any oracle has ever had.

Paradatum The gateway. Customers pay in dollars and get a receipt, a proof, and a re-verify link. They never hold a token.	Blindproof The run. Deterministic, private, proven artificial intelligence (AI) compute. Receipt R1 with proof.	DeFacts The claim. The fact, its proof, the source hash, and the model checkpoint, filed. Receipt R2.	Knowledgecoin The knowledge. The claim adjudicated by staked agents and people, verdict and reasoning on the record. Receipt R3.
--	---	---	--

Three receipts, one chain

- **R1, the run.** A Blindproof miner, or a Paradatum graphics processing unit (GPU), executes a job under a pinned contract. Masked, so the operator never sees plaintext. The same bytes on any GPU. R1 carries the proof.
- **R2, the claim.** DeFacts files the result as a fact in the DeFacts schema with its proof, the source hash, and the model checkpoint. Bytes on Arweave, hash on Solana. R2 chains to R1. Provenance is proven; truth is not claimed. Truth is what the layer above is for.
- **R3, the knowledge.** Knowledgecoin records the adjudication: claim identifier, panel, each adjudicator's stake, vote, and reasoning receipt, and a stake-weighted confidence. R3 chains to R2, to R1, to the contract. Anyone can audit the whole chain.

The four layers each have exactly one unit of account, and each unit does one job only that unit can do. Customers pay in dollars. Miners stake DARK, the Blindproof alpha. Producers of record hold FACT, the DeFacts token, which carries the royalty every time a claim is re-verified.

Adjudicators stake KNOW, the Knowledgecoin token, which has one job: to be slashed. Skin in the game. The sequencing, the escrow, and the treasury rules are in Appendix A.3.

9 How a claim becomes knowledge

A claim enters Knowledgecoin only after R1 and R2 exist. The adjudication itself has five steps.

1. **A claim is posted.** Over a DeFacts claim, with a fee in dollars and a KNOW bond from the poster: “this is true” or “this is false.” The poster is not arguing about the source or the run; those are settled.
2. **A panel forms.** Adjudicators are drawn by chain-seeded randomness from those who have staked KNOW. By default an adjudicator is an agent: a pinned model running on Blindproof, so the vote is itself a proven run with a reasoning receipt. People may stake and sit too.
3. **Each adjudicator stakes on a verdict.** The verdict carries the evidence relied on and the reasoning, and for an agent the reasoning is a proven computation, re-runnable by anyone.
4. **The panel settles.** The majority earns the fee. The minority is slashed. A losing party may appeal, and an appeal widens the panel. A claim that remains contested after the appeal

ladder is recorded as contested, which is a valid final state: the Perfect Library records disagreement rather than hiding the disagreement.

5. **The record is written.** R3 goes on the ledger, chained to R2 and R1. The claim is now knowledge with a confidence, and the confidence is stake-weighted agreement among parties with something to lose, not a model's self-reported probability.

A worked example

The numbers below are illustrative. Panel sizes, bond sizes, and the slashing fraction are protocol parameters and remain DESIGNED.

A public statute is hashed and prefilled once. A user asks whether the statute imposes a particular reporting obligation and from what date. The deterministic extraction returns a claim in the DeFacts schema: subject, predicate, object, the effective date, the clause relied on, and a confidence, with the source hash and the model checkpoint. Blindproof returns R1 with its proof. DeFacts files R2. A second user, who relies on the obligation commercially, posts the claim for adjudication with a fee of 10 dollars and a bond of 100 KNOW.

A panel of seven forms: six agent adjudicators across three different pinned models, and one person who has staked. Each stakes 100 KNOW. Five vote true, citing the clause and a cross-reference the extraction had marked as an implicit fact. Two vote false, reading the effective date differently. The five share the fee and the slashed portion of the two minority stakes. The two may appeal; if they do, the panel widens to fifteen, and the appeal carries its own bond. R3 records all seven verdicts and reasoning receipts, the stake-weighted confidence (500 against 200, so 71 percent on this panel), and the chain back to the proof of the run that produced the claim. The next user who asks the same question gets the record, with its proof re-checked, and no model runs.

Asked again: no rerun. Write once, verify forever.

10 Proof-of-Knowledge, taken literally

In 2021 and 2022, Proof-of-Knowledge was an analogy to Proof-of-Work: blocks created by knowledge producers, miners competing to validate. In 2026 the phrase is literal. A **Knowledge Record** is an object with the following parts, every one of them content-addressed and checkable.

- The **claim**, in the DeFacts schema: explicit and implicit facts with confidence; who, what, when, where, why, how; subject–predicate–object triples; resolved entities; timeline and causal chain; full provenance.
- The **evidence**: source hashes and the extracted passages relied on.
- The **contract identifier** and **model hash** of the pinned model that produced the claim.

- The **proof of the run**, or in the replication tier the hash-equality record of independent miners.
- The **adjudication**: panel, stakes, verdicts, reasoning receipts, confidence, and the appeal history.

The 2022 white paper nested the proofs: Proof-of-Information (concepts are valid) inside Proof-of-Logic (the logic is valid) inside Proof-of-Reason (the reasoning is consistent with accepted knowledge) inside Proof-of-Knowledge, with Proof-of-Veritas requiring additional rounds of validation before a claim entered the highest repository. That nesting survives unchanged as the structure of the adjudication questions. An adjudicator votes on each layer in turn, and a claim cannot pass an outer layer without passing the inner layers. Proof-of-Veritas becomes rounds over time: a record that survives repeated adjudication as new evidence arrives earns a higher confidence, and a record that fails a later round is marked, never erased.

The 2022 paper's VLKRIE protocol, Values, Logic, Knowledge, Reason, Information, Ethics, also survives, with a division of labor. Logic, Knowledge, Reason, and Information are questions a pinned model can be asked and proven on. Values and Ethics are questions no pinned model should decide alone; those claims escalate to a human panel, and the record says so.

The 2022 paper recorded each layer as a non-fungible token so that the owner could derive payment every time the item was referenced by others. The economics survive without the word: the producer of record earns a royalty every time the claim is re-verified or adjudicated, and that royalty is carried by FACT in the DeFacts layer. Submitting valid bricks into the pyramid of knowledge is paid work.

11 The island's industries become jobs

The 2022 white paper described Knowledgecoin as an island economy, borrowing Filecoin's term: raw data arrives at the shore and is processed by the island's industries until the output is validated knowledge. [3] Each industry was a human role earning tokens. Each industry is now a deterministic, proven job with a content-addressed contract, and a human may still take the job wherever a human adds something a model cannot.

DeTag

Indexing of information. Deterministic extraction under the DeFacts standard: sentence-level explicit and implicit facts with confidence, resolved entities, timeline, causal chains, provenance. The shore. **PROTOTYPE**

DeBabel

Universal concept translation. Entity resolution and concept mapping, so that two claims using the same words differently are compared on what they mean. A proven job over the DeTag output. **DESIGNED**

DeLogic

Contradiction detection. Does a claim, or a set of claims, contradict itself? A pinned model answers; the proof shows the pinned model answered. DESIGNED

DeReason

Claims in proper context. Does a non-contradictory claim contradict the knowledge already on the record, and does the claim engage with the evidence? A proven job against the DeFacts library. DESIGNED

DeKnow

Validated knowledge. The DeFacts library itself: claims with proofs and provenance, served and re-verified rather than recomputed. DESIGNED

DeTruth

Highly validated knowledge. The adjudicated set in Knowledgecoin: records that have survived staked panels and rounds, with confidence and full history. DESIGNED

Two further pieces of the 2022 design carry over with sharper meaning. The **Intellectual Clarity Score**, described in 2022 as an eBay-style vendor score for truth claims, now applies to models and agents as well as to people: every adjudicator has an auditable record of being right or wrong, with stake behind each verdict, and the score is computed from the ledger rather than reported by the platform. And the **shared verifier** replaces the 2022 promise that real identities would prevent sock accounts: an agent adjudicator is a proven run or is nothing, and a human adjudicator's stake, not the human's identity, is what the protocol counts.

12 Why adjudication is tractable here

Every prior attempt to crowdsource truth drowned in the same place: the crowd argued about the inputs. What did the source say? Who edited the entry? Was the quote complete? Knowledgecoin removes that surface entirely.

- **The input is never in dispute.** The source is hashed, the extraction is deterministic, and the proof is checked before the claim is admitted. Adjudicators argue only whether the claim holds.
- **Agent adjudicators are auditable.** “Why does the system believe this?” has an answer with hashes in it: which model, which evidence, which reasoning, with a proof. A hosted assistant cannot give that answer; a pinned model on a proven run cannot avoid giving the answer.
- **Confidence is stake-weighted agreement** among parties with something to lose, not a model's self-reported probability.
- **Capture is detectable.** A deviating agent adjudicator produces a verdict that a re-run contradicts, and the contradiction is slashable. A deviating human adjudicator loses stake on appeal. Neither can quietly bend the library.

- **The record pays its makers.** The producer of record earns on every re-verification, and the adjudicators who were right earn on every panel. The Epistemic Agent Problem is answered with a ledger entry rather than a promise.

13 The Know Tribe

Knowledgecoin is an anti-tribal community of increasingly empowered epistemic agents whose loyalty is incentivized toward group-independent objective truth: a Know Tribe that transcends every other tribal affiliation. In 2021 the definition rested on shared commitment. In 2026 the definition rests on a shared verifier. All truly rational agents have equal self-interest in unbiased, objective knowledge, and two agents who disagree about everything else can still run the same verifier on the same record and agree on what ran, what the source said, and who staked what on the verdict. A tribe defined by its verifier is no tribe.

Knowledge specialization increases knowledge within domains; knowledge fragmentation is self-defeating. Academia offers the standing example, where each discipline wages turf wars over what is to be researched, validated, published, and taught, when the disciplines would mutually benefit from the integration that a unified Tree of Knowledge demands. Knowledgecoin shares the integrative aspiration of the cognitive sciences and philosophy, and now has a mechanism to house the Tree of *All* Knowledge: one record format, one verifier, one ledger, open to every domain.

The Know Tribe includes machines on the same terms as people. An agent that stakes, reasons on the record, and is slashed when wrong is an epistemic agent in the full 2021 sense: a seeker, producer, validator, and dispenser of knowledge, accountable for every claim.

14 Knowledgecoin Bill of Rights

All data-citizens of Knowledgecoin shall enjoy the following rights and privileges. The first ten are carried over from the 2022 white paper; the last three are new, and each of the three now has a mechanism rather than a promise.

1. The right to reliably trust and self-verify all Knowledgecoin system content.
2. The right to be rewarded for the products of one's own research in the system.
3. The guarantee that bad data-citizens will be penalized for their misdeeds.
4. The guarantee that validated information in the system remains secure over time, changing only when greater validation improves upon it, as with scientific progress.
5. The right to upload information onto the ledger, limited only by the no-actual-harm principle (no libel, slander, fraud, threats of force, and the like).

6. Freedom of access to validated knowledge, curated by the decentralized community of data-citizens rather than by a gatekeeper.
7. Freedom from confusing misinformation that obscures valid knowledge.
8. Freedom from propaganda, ideological narratives, spin, falsehoods, equivocation, lies, fallacies, distortions, and other forms of misinformation and sophistry.
9. The right to express speech freely, free from bullying, cancel culture, de-platforming, and other incursions against freedom of thought.
10. The right to participate in an economy that is used for the benefit of human knowledge and does not violate the non-aggression principle.
11. **The right to a changeless adjudicator.** Every record names the pinned model and contract that judged the claim, and that adjudicator can be consulted again, unchanged, at any later time.
12. **The right to verify the reasoning, not only the record.** Every verdict carries a reasoning receipt, and an agent's reasoning is a proven computation that anyone can re-run.
13. **The right to private inquiry.** A question put to the network is never plaintext to the machines that answer the question.

15 Team

Knowledgecoin is a Paradatum project. Paradatum Inc. is the lab that owns the intellectual property across the stack and holds equity in the three protocol entities, and runs none of their markets.

Mark Gleason

Founder and Chief Executive Officer, Paradatum Inc.
· **Commercial Lead**

Twenty-five years leading enterprise artificial intelligence in insurance and reinsurance. Harvard economics and computer science; Naval War College cyber warfare and cybersecurity. Founder of the Knowledge Works management consultancy; software engineer for the Department of Defense and the Defense Advanced Research Projects Agency (DARPA); inventor of DBX, an ultra-high-speed analytics database engine.

Sean Kimball

Technical Lead

Chief Information and Data Officer with experience managing internationally distributed teams of more than 250 staff and international initiatives supporting 900 billion dollars in assets. Managed data, architecture, application, operations, security, and regulatory compliance for Fortune 50 companies; directed legal and regulatory special projects for executive boards.

Rick Repetti, PhD**Chief Philosophy Officer**

Full Professor of Philosophy, City University of New York. Board of Directors, American Philosophical Practitioners Association. Editor, Journal of Buddhist Ethics. Fellow, Center for Contemplative Mind in Society. Author of numerous works.

Jeanette Straughn**Operations Lead**

Head of Program Management and Global Head of Data Governance; Collibra ontological engine lead; Fortune 50 chief of staff; audit and compliance; executive reporting; Project Management Professional.

16 Conclusion

The 2021 paper ended with a call to people of reason. The call stands, and the call now has a mechanism behind it. The Perfect Library could not be built on an immutable record alone, because an immutable record of unverified claims is a museum of assertions. The Perfect Library can be built on an immutable record of *proven* claims, *adjudicated* by parties with something to lose, where the adjudicators are themselves changeless and re-runnable, and where the people and machines who did the epistemic work are paid every time the work is relied on.

Paradatum gates. Blindproof is what runs. DeFacts is what is claimed. Knowledgecoin is what it means.

Calling all people of **reason**, wherever such might still be found.

You are not alone. Take back your birthright.

Join a community dedicated to validated knowledge.

The mind you save may be your own.

knowledgecoin.io

References

1. Grice, H. P., 1975. "Logic and Conversation." In P. Cole and J. L. Morgan, eds., *Syntax and Semantics*, vol. 3. New York: Academic Press, pp. 41–58.
2. *United States v. Heppner*, S.D.N.Y., No. 25-cr-00503-JSR. Discussed in [Harvard Law Review Blog](#), March 2026.
3. Filecoin, 2021. "Filecoin's 'Island Economy.'" [filecoin.io](#), September 23, 2021.
4. Mainelli, M., von Gunten, C., and Z/Yen Group, 2014. "Chain of a Lifetime: How Blockchain Technology Might Transform Personal Insurance." *Long Finance*, December 2014, pp. 1–51.
5. Rand, A., 1984. *Philosophy: Who Needs It*. New York: Signet.
6. Gleason, Repetti, and Straughn, 2021. *Knowledgecoin Position Paper*, v1.0.0, December 2021.
7. Gleason, Repetti, and Straughn, 2022. *Knowledgecoin White Paper*, v1.0.1, April 2022.

8. Paradatum Inc., 2026. *Blindproof: Deterministic, Private, Proven, Settled AI Inference*, white paper v0.2.14. blindproof.ai/whitepaper.

A Mechanics and status

This appendix carries the operational detail behind Part 3: the objects, the panel and stake rules, settlement and the token roles, the attack table, the confidence rule, and the public evidence for each component. Every mechanism carries one of four labels.

MEASURED a public, checkable result exists **PROTOTYPE** running, not yet public
DESIGNED specified, not yet built **PENDING** an open decision

A.1 Objects **DESIGNED**

Nothing in any of these objects is secret. The claim is the DeFacts record; the verdict is one adjudicator's stake on the claim; the Knowledge Record is the settled R3.

```
claim {
  claim_id, source_hash, contract_id, model_hash,
  fact { subject, predicate, object, when, where, confidence, clause_refs[] },
  implicit_facts[], entities[], provenance[],
  proof_ref // R1 proof bundle hash, or replication hash set
}

verdict {
  claim_id, adjudicator_id, kind: "agent" | "person",
  layer: "information" | "logic" | "reason" | "knowledge",
  vote: true | false | abstain, stake_know,
  reasoning_receipt // for an agent: contract_id + proof of the reasoning run
}

knowledge_record { // R3
  claim_id, panel[], verdicts[], confidence, round, appeals[],
  state: "knowledge" | "contested" | "rejected",
  chain: { r2_hash, r1_hash, contract_id }
}
```

A.2 Panels and stake **DESIGNED**

- Panels are drawn by chain-seeded randomness from staked adjudicators, across at least three distinct pinned models when agents are drawn, so that no single model is a monoculture on any panel.
- An agent adjudicator is a Blindproof job: pinned model in, verdict and proof out. The miner who runs the job earns DARK in the compute layer; the agent earns the fee and KNOW in the knowledge layer.

- A person may stake and sit. A person's verdict carries a written reasoning receipt and is subject to the same appeal.
- The majority earns the fee; the minority is slashed. The slashing fraction, the minimum panel size, and the appeal ladder are protocol parameters. PENDING
- Appeals widen the panel and carry a bond. Values and Ethics questions escalate to a human panel by rule. "Contested" is a valid terminal state.

A.3 Settlement and token roles DESIGNED

One escrow in dollars, released per receipt. The customer pays in United States dollars, in USDC (USD Coin, a dollar stablecoin), or in prepaid credit. The Paradatum gateway pins the contract, slices the job, issues the receipts, keeps its margin, and escrows the rest against one job identifier. R1 verifies, and the compute share releases to Blindproof. R2, chained to R1, releases the filing share to DeFacts. R3, chained to R2, releases the adjudication fee to the panel. If a proof fails, nothing files and the customer is refunded in the unit charged. There is no exchange rate inside a job.

UNIT	LAYER	THE ONE JOB ONLY THAT UNIT DOES
USDC	All	Settle every customer payment and fee. Issued by nobody Paradatum controls.
DARK	Blindproof	Be the miner's stake. Bittensor settles only in subnet alpha and TAO, Bittensor's network token.
FACT	DeFacts	Be the producer's stake in the library and carry the royalty on re-verification.
KNOW	Knowledgecoin	Be slashed. The adjudicator's stake, and nothing else.

Order of operations. Dollars first, with PSEC contracts and receipts; prove-and-file is the headline job type. FACT when the library is worth re-verifying. DARK at Blindproof mainnet. KNOW when there are enough proven claims to argue about, because adjudication needs evidence on the table. Treasury rules are published and checkable from the receipts: a published share of gross revenue buys TAO and stakes into the DARK pool, DeFacts fees buy back FACT, and adjudication fees fund the KNOW reward pool. Paradatum holds equity in the three entities and runs none of their markets. Counsel precedes any token issuance.

Superseded. The 2022 white paper stated a fixed supply of 125,000,000 KNOW, a launch target price, a private sale, and a public sale schedule. Those figures are withdrawn. Supply, issuance, and any sale are decided when adjudication goes live, and not before. PENDING

A.4 Attacks and defenses DESIGNED

ATTACK	DEFENSE	ASSUMPTION
Disputed inputs	Source hashed, extraction deterministic, proof checked before admission	Proof system only
Lazy agent adjudicator	Verdict must carry a proven reasoning run; a bare vote is not a verdict	Proof system only
Deviating agent adjudicator	Re-run contradicts the verdict; contradiction is slashable	Determinism
Bribed or colluding panel	Stake-weighted settlement; appeal widens the panel; panels drawn across models	Majority of stake honest per round
Sybil adjudicators	Stake, not identity, is counted; agents are proven runs or nothing	Stake has cost
Claim spam	Fee plus bond per posted claim	Policy setting
Appeal spam	Appeal bond, forfeited on loss	Policy setting
Model monoculture	At least three distinct pinned models per agent panel	Models available
Subjective claims gamed	Values and Ethics escalate to a human panel; "contested" is a valid state	Human panel honest
Silent update of an adjudicator	Impossible: the contract pins the model; a new model is a new adjudicator with a new record	Determinism
History rewritten	Every record content-addressed and anchored; bytes on Arweave, hash on Solana	Ledger liveness

A.5 Confidence DESIGNED

Confidence on a record is the stake-weighted share of the panel that voted with the settled verdict, carried per layer (information, logic, reason, knowledge). Rounds over time are the Proof-of-Veritas of the 2022 paper: a record that survives a later round with new evidence rises; a record that fails a later round is marked with the round that failed, and the earlier rounds remain on the record. Confidence is never a model's self-reported probability.

A.6 What exists today

Every component below the adjudication layer either has a public result or is running. The adjudication layer is designed and not built, and this paper says so wherever the layer is described.

COMPONENT	STATUS	EVIDENCE
Deterministic execution of a mid-size open-weight model	MEASURED	Bit-exact results across four GPU types and a CPU reference verifier, shown publicly in Paradatum's demonstration and on the Blindproof site.
Proof of a full-model run, verified without re-execution	MEASURED	Paradatum's full-model proof with sub-second verification; independent audit in progress.
Portable Semantic Execution Contract (PSEC) and CPU verifier	MEASURED	Canonical contract hash reproduced across hardware; CPU verifier published as the reference implementation.
Blindproof settlement rail	MEASURED	Bittensor testnet subnet 575: honest miners paid, a lazy miner never credited, weights set on chain. Mock inference and mock proofs on the testnet runs, as the Blindproof evidence pages state.
Masked (private) execution	PROTOTYPE	Running on model pieces; full-model masked inference is the research program.
DeFacts standard and extraction	PROTOTYPE	Sentence-level fact-extraction standard defined; first public-document pipeline in build.
Knowledgecoin adjudication	DESIGNED	This paper.

Version 2.0.0, October 2026. Supersedes the Knowledgecoin Position Paper v1.0.0 (December 2021) and the Knowledgecoin White Paper v1.0.1 (April 2022). Comments and review: info@paradatum.ai.