

AI & Blockchain: Decentralized Compute, Agents, and Verification


When Machine Learning Meets Distributed Ledgers

Prof. Dr. Jörg Osterrieder

BSc Blockchain, Crypto Economy & NFTs

Spring 2026

[Cartoon] When AI Met Blockchain



The convergence nobody asked for — but everybody needs.

Two technologies walk into a bar.

One can predict the future but nobody knows how. The other records everything but cannot think.

AI is **powerful but opaque**. It makes decisions no human can audit, runs on servers no user controls, and learns from data nobody can verify.

Blockchain is **transparent but dumb**. It records everything immutably but cannot analyze, predict, or adapt.

What if they worked together?

Blockchain could make AI verifiable. AI could make blockchain intelligent. Or they could just cancel each other out.

Today we find out which.

AI is powerful but opaque. Blockchain is transparent but dumb. This lecture asks: can opposites combine?

By the end of this lecture, you will be able to:

1. **Describe** the convergence thesis — why AI and blockchain complement each other's weaknesses [*Understand*]
2. **Explain** how decentralized compute networks replace centralized cloud providers for AI workloads [*Understand*]
3. **Calculate** the cost comparison of decentralized vs. centralized GPU compute for a training job [*Apply*]
4. **Compare** centralized AI services with decentralized alternatives on transparency, cost, and censorship resistance [*Analyze*]
5. **Evaluate** whether a proposed AI-blockchain project creates real value or is just “buzzword stacking” [*Evaluate*]

Prerequisites: Basic blockchain concepts (Blockchain Foundations lecture).

These objectives span Bloom levels from Understand to Evaluate.

The AI-Blockchain Convergence

What this lecture covers:

AI and blockchain are usually studied separately. This lecture explores what happens when they meet.

Three pillars:

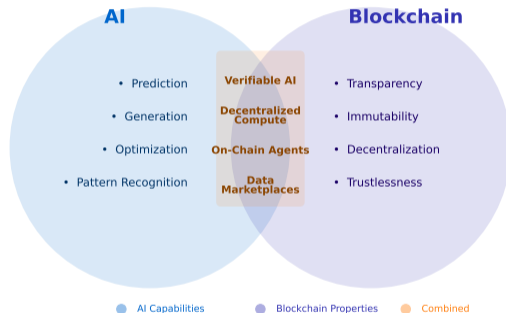
1. **Decentralized compute** — break the cloud monopoly on AI training and inference
2. **AI agents** — autonomous software that lives on-chain and acts without human intervention
3. **Verification** — prove that an AI model produced a specific output without revealing its weights

Each pillar addresses a real limitation. Together, they define a new design space.

Three pillars — compute, agents, verification — define the AI-blockchain design space.

The AI × Blockchain Convergence

Where machine learning meets distributed ledgers



- **What you see:** The three convergence pillars and their intersections
- **Key pattern:** Most real projects sit at one pillar; few span two; almost none span all three
- **Takeaway:** The field is early — most value is still in single-pillar solutions

A Thought Experiment

You use ChatGPT every day. It knows your business plans, your code, your health questions, your relationship problems.

But OpenAI can change the model without notice. They can raise prices tomorrow. They can censor outputs based on policy decisions you never agreed to. They can sell aggregated data about your usage patterns.

You have no alternative because the model weights are proprietary. Switching costs are enormous — your conversation history, your fine-tuned workflows, your muscle memory for prompting — all locked in.

What if AI ran on open, decentralized infrastructure?

What if you could choose which model to run, verify that it was not censored, and pay per query with no subscription lock-in?

This is not hypothetical. Projects like Akash, Bittensor, and Ritual are building it today.

You trust a single company with your most sensitive queries. Decentralized AI asks: what if you did not have to?

Definition: Convergence Thesis

AI and blockchain have complementary weaknesses: AI is powerful but opaque, centralized, and unverifiable. Blockchain is transparent, decentralized, and trustless — but cannot perform complex computation. Their combination addresses each other's limitations.

Dimension	AI Strength	Blockchain Strength
Computation	Complex pattern recognition	Simple deterministic logic
Transparency	Opaque (black box)	Fully auditable
Trust model	"Trust the company"	"Trust the math"
Data handling	Learns from massive datasets	Records immutable transactions
Centralization	Concentrated (few providers)	Distributed (thousands of nodes)

Where one is weak, the other is strong. The question is whether combining them creates value — or just complexity.

Opposite strengths do not automatically combine. Integration must solve a real problem, not just sound clever.

Three needs blockchain can address:

1. Data provenance


Prove that training data is authentic and unaltered. A model trained on verified data is more trustworthy than one trained on scraped web pages.

2. Decentralized compute

Break the cloud monopoly. AWS, Google Cloud, and Azure control 65%+ of cloud GPU capacity. Blockchain-based marketplaces let anyone rent idle GPUs.

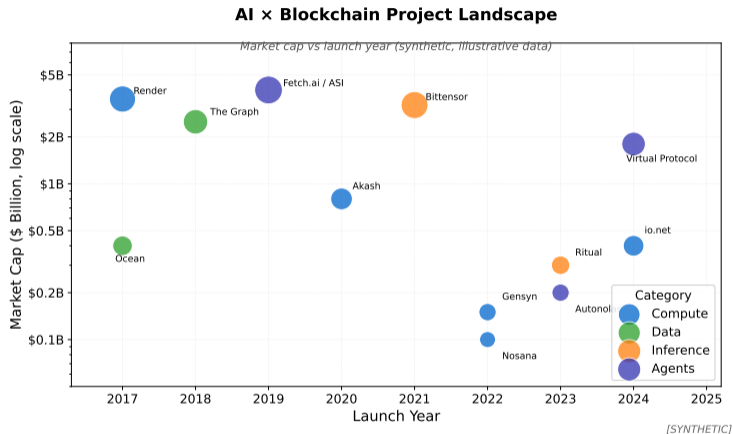
3. Verifiability

Prove that inference was correct without revealing the model weights. Zero-knowledge proofs make this possible — in theory.



- **What you see:** Cloud compute market share dominated by three providers
- **Key pattern:** AWS, Google, and Azure together control the vast majority of AI compute capacity
- **Takeaway:** Decentralized compute targets this concentration — but must match reliability

AI needs provenance, compute, and verifiability. Blockchain offers all three — at the cost of speed and simplicity.



- **What you see:** AI-blockchain projects mapped by category — compute, agents, data, verification — with market cap sizing
- **Key pattern:** A cluster of launches in 2024–2025 signals rapid growth; market cap is concentrated in a handful of projects
- **Takeaway:** The landscape is wide but shallow — many projects, few with proven traction

Over 200 AI-blockchain projects exist. Most launched after 2023. Few have sustainable revenue.

Three needs AI can address:

1. Smart contract optimization

AI predicts gas costs, suggests optimal transaction timing, and identifies inefficient contract patterns before deployment.

2. Fraud detection

Machine learning identifies suspicious on-chain patterns — wash trading, rug pull signatures, Sybil attacks — faster than manual analysis.

3. Autonomous agents

AI bots execute complex on-chain strategies: portfolio rebalancing, yield farming optimization, MEV extraction, and DAO governance participation.

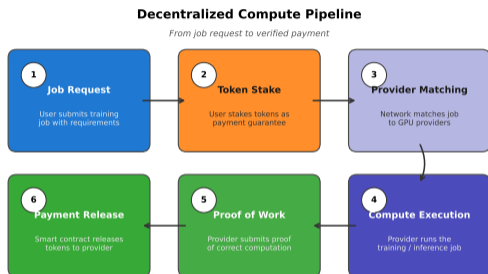
Real-World Example

MEV (Maximal Extractable Value) bots already use machine learning to find arbitrage opportunities in DeFi. These bots scan pending transactions, predict price impacts, and insert profitable trades — generating over \$600M in extracted value in 2023 alone.

They are the first successful AI agents on blockchain — profitable, autonomous, and controversial.

Blockchain gives AI something to act on. AI gives blockchain something to think with.

Blockchain needs intelligence. AI needs a trustless execution environment. The fit is conceptually clean.



The 6-step pipeline:

1. **Request:** User posts a compute job (model training or inference) with budget
2. **Stake:** Provider locks tokens as collateral (slashed if they cheat)
3. **Match:** Protocol matches job to provider based on hardware specs and price
4. **Compute:** Provider runs the workload on their GPU
5. **Prove:** Provider submits proof of correct execution (optimistic or ZK)
6. **Pay:** Smart contract releases payment; stake is returned

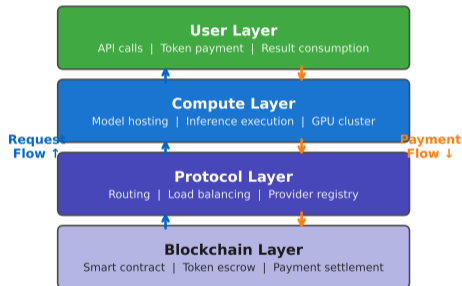
Key projects: Akash Network, Render Network, io.net, Gensyn

- **What you see:** A 6-step pipeline from compute request to payment
- **Key pattern:** Staking and proof steps add trust without a central authority
- **Takeaway:** The protocol replaces AWS's reputation with cryptographic guarantees

GPU marketplaces replace corporate trust with economic incentives: providers stake tokens they lose if they cheat.

Token-Gated AI Inference Architecture

Pay per query with verifiable results



How it works:

1. User sends query + payment to smart contract
2. Contract escrows tokens
3. Provider runs inference, returns result
4. Verification check (optimistic or ZK proof)
5. Payment released if valid; stake slashed if not

Why it matters: No lock-in, no censorship, pay per query, switch providers freely. *Projects: Bittensor, Ritual, Allora*

- **What you see:** Token flow from user to provider via smart contract escrow
- **Key pattern:** Payment released only after validation — no trust required
- **Takeaway:** AI becomes a metered utility, like electricity

Tokenized inference eliminates subscriptions: you pay per query, switch providers freely, and no one can censor your requests.

Definition: Compute Cost Formula

$$\underbrace{\text{Total Cost}}_{\text{what you pay}} = \underbrace{\text{GPU-Hours}}_{\text{compute time}} \times \underbrace{\text{Price per Hour}}_{\text{market rate}} + \underbrace{\text{Network Fee}}_{\text{protocol cut (2-5%)}}$$

Scenario: Training a small language model — 100 GPU-hours on an A100 GPU.

Provider	Rate / GPU-hr	Network Fee	Total Cost
AWS (centralized)	\$32.00	\$0	\$3,200.00
Google Cloud	\$28.00	\$0	\$2,800.00
Akash (decentralized)	\$5.00	\$12.50 (2.5%)	\$512.50
Render Network	\$8.00	\$20.00 (2.5%)	\$820.00

Savings with Akash: $\$3,200 - \$512.50 = \$2,687.50$ (84% cheaper)

Note: Prices are approximate and vary by GPU type, availability, and network congestion. Decentralized providers may have lower uptime guarantees and less tooling than cloud providers.

Decentralized compute can be 60–85% cheaper — but the trade-off is reliability, tooling, and support.

What is an AI agent?

An AI agent is autonomous software that **observes** on-chain state, **decides** on actions using a model, and **acts** by submitting transactions — without human intervention.

Three types:

1. **MEV bots** — scan pending transactions, extract arbitrage profits
2. **DAO governance bots** — analyze proposals, vote based on criteria
3. **Portfolio rebalancers** — monitor positions, adjust allocations automatically

The agent runs 24/7. It does not sleep, forget, or get emotional.

Real-World Example

In January 2025, an AI agent on Virtuals Protocol autonomously managed a \$2M portfolio, making 1,200 trades in 30 days with no human input.

The agent selected DeFi strategies, executed swaps, and rebalanced positions — all recorded on-chain for anyone to audit. Performance: +8.3% in 30 days (vs. ETH +2.1%).

The key question: When an autonomous agent loses money or causes harm — who is responsible? The developer? The model? The protocol?

AI agents are the first software that acts autonomously on-chain. Accountability for their actions remains unresolved.

zkML: Proving AI Without Revealing the Model

zkML: Zero-Knowledge Machine Learning Verification

Prove AI inference correctness without revealing the model



Key insight: Verification is cheap (\approx ms), proof generation is expensive (\approx min)

- **What you see:** The zkML pipeline from model input to on-chain verified proof
- **Key pattern:** The model weights never leave the provider; only the proof goes on-chain
- **Takeaway:** zkML separates verifiability from transparency — you can prove correctness without revealing how

zkML is the holy grail of AI verification: prove the answer is correct without showing how you got it.

Zero-Knowledge Machine Learning (zkML):

Proves that a *specific model* produced a *specific output* for a *specific input* — without revealing the model weights.

Use case: on-chain credit scoring

1. Bank trains a credit model (proprietary)
2. User submits financial data
3. Model produces a credit score
4. ZK proof verifies: “This score came from this model on this data”
5. Smart contract uses the score — model stays private

Current limitation: Generating ZK proofs for large neural networks is extremely slow — minutes to hours for a single inference.

Key projects: EZKL, Modulus Labs, Giza

Approach 1: Compute-to-Data

Ocean Protocol's model: data stays with the owner. The model travels to the data, trains locally, and returns only the trained weights.

How it works:

1. Data owner publishes a dataset description (not the data)
2. Buyer pays tokens for compute access
3. Algorithm runs inside the data owner's environment
4. Only results (trained model or predictions) are returned
5. Data never leaves the owner's server

The data is the product, but the buyer never sees it.

Approach 2: Federated Learning + Blockchain

Multiple parties train a shared model without sharing raw data. Each participant trains locally and shares only weight updates.

Blockchain's role:

- Records each participant's contribution
- Distributes token rewards proportional to data quality
- Detects and penalizes poisoned updates
- Creates an immutable audit trail of the training process

Both approaches solve the same problem:

How do you train AI on sensitive data (medical records, financial data, personal information) without exposing it?

Key projects: Ocean Protocol, FedML, Oasis Network

Privacy-preserving ML lets AI learn from sensitive data without anyone seeing it. Blockchain tracks who contributed what.

How AI enhances smart contracts:

1. Predictive oracles

AI models predict real-world outcomes (weather, crop yields, election results) and feed them on-chain for parametric insurance and prediction markets.

2. Dynamic parameters

AI adjusts protocol parameters in real time: lending rates, collateral ratios, and liquidation thresholds based on market conditions.

3. Automated auditing

ML models scan smart contract code for known vulnerability patterns before deployment — faster and cheaper than manual audits.

Real-World Example

Chainlink's oracle network processes billions of data points per year, feeding real-world data to DeFi protocols across multiple blockchains.

AI-enhanced oracles can weight data sources by reliability, detect anomalous feeds, and predict short-term price movements — reducing the risk of oracle manipulation attacks.

The tension:

Smart contracts are deterministic — same input, same output. AI models are probabilistic — same input, potentially different output. Integrating them requires careful design to preserve contract reliability.

AI makes oracles smarter and contracts more adaptive — but probabilistic outputs conflict with deterministic execution.

Content Provenance: Fighting Deepfakes On-Chain

Content Provenance: Fighting Deepfakes On-Chain

C2PA + blockchain attestation pipeline



Without provenance, any content could be AI-generated -- blockchain makes verification trustless

- **What you see:** The content provenance pipeline from creation to consumer verification
- **Key pattern:** Hashes are registered on-chain at creation; consumers verify by checking the chain
- **Takeaway:** Deepfakes are detected by the absence of a valid on-chain attestation

As AI-generated content becomes indistinguishable from real content, on-chain provenance becomes essential.

The C2PA standard + blockchain:

The Coalition for Content Provenance and Authenticity (C2PA) creates metadata “nutrition labels” for digital content.

How blockchain adds trust:

1. Creator registers content hash on-chain at creation time
2. Every edit creates a new hash linked to the original
3. Consumer scans content and checks on-chain record
4. Missing or tampered attestation = likely deepfake


Why blockchain matters here:

- No single company controls the registry
- Records cannot be altered after the fact
- Open verification — anyone can check

Key projects: *Numbers Protocol, Starling Lab*

AI × Blockchain Risk Matrix

Key risks by likelihood and impact



The deadly combination: AI generates wrong outputs (**hallucinations**). Blockchain makes them permanent (**immutability**).

Scenario: AI oracle hallucinates price spike → DeFi liquidates \$50M → on-chain → no rollback.

Mitigations:

- Multi-oracle consensus (median of N sources)
- Circuit breakers (pause if deviation > X%)
- Time-delayed execution for large transactions

- **What you see:** Risk matrix — AI failure modes vs. blockchain properties
- **Key pattern:** Top-right quadrant = danger zone (high error + immutability)
- **Takeaway:** Unreliable AI + permanent records = irreversible mistakes

Mango Markets (2022): oracle manipulation drained \$114M. AI hallucinations could cause similar damage without intentional manipulation.

The irony of “decentralized” AI:

Decentralized compute marketplaces promise to break the cloud monopoly. But they still depend on NVIDIA GPUs — and NVIDIA controls 90%+ of the AI accelerator market.

If NVIDIA restricts supply, raises prices, or changes licensing terms, the *entire* “decentralized” AI ecosystem suffers.

The centralization stack:

Even a fully decentralized protocol has centralized dependencies at every layer. True decentralization requires diversity at every level.

Layer	Centralized?	Risk
Hardware	NVIDIA 90%+	Supply chokepoint
Models	5 labs dominate	Knowledge monopoly
Data	Web scraped	Copyright lawsuits
Training	\$100M+ cost	Barrier to entry
Oracles	Few providers	Single point of failure

The question: Is decentralized AI actually decentralized — or just centralized infrastructure with a token on top?

Decentralization is a spectrum, not a binary. Most AI-blockchain projects are less decentralized than they claim.

“Decentralized AI” often means decentralized coordination with centralized hardware, models, and data.

Energy and environmental cost:

Activity	Energy (MWh)
GPT-4 training run	~50,000
Bitcoin mining (annual)	~150,000,000
Average US household/yr	10.5

A single GPT-4 scale training run consumes as much energy as roughly 5,000 US households use in a year.

Adding blockchain verification on top increases the footprint further — every proof, every on-chain record, every consensus round costs energy.

Proof-of-Stake reduced blockchain energy by 99%, but AI training costs keep rising.

Ethical concerns:

1. Biased AI in immutable systems

If a biased credit scoring model is deployed on-chain, its discriminatory decisions become permanent and auditable — but also irreversible.

2. Automated discrimination

AI agents acting on-chain can discriminate at machine speed: denying loans, rejecting insurance, or front-running trades — with no human review.

3. Accountability gaps

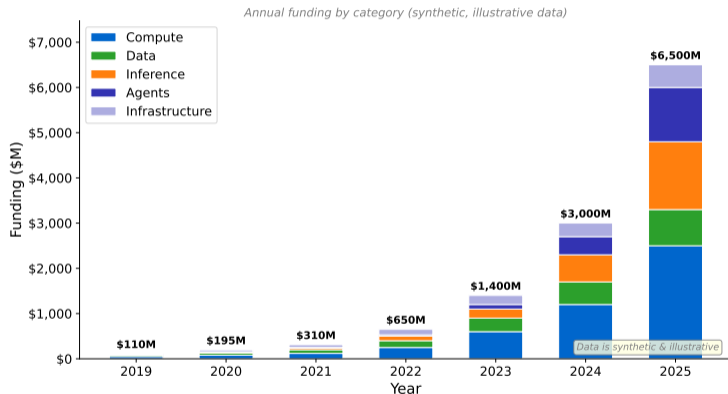
When an AI agent causes harm on-chain: who is liable? The developer who trained the model? The protocol that deployed it? The DAO that voted to use it?

4. Deepfake amplification

Decentralized, censorship-resistant networks make it harder to remove AI-generated misinformation once published.

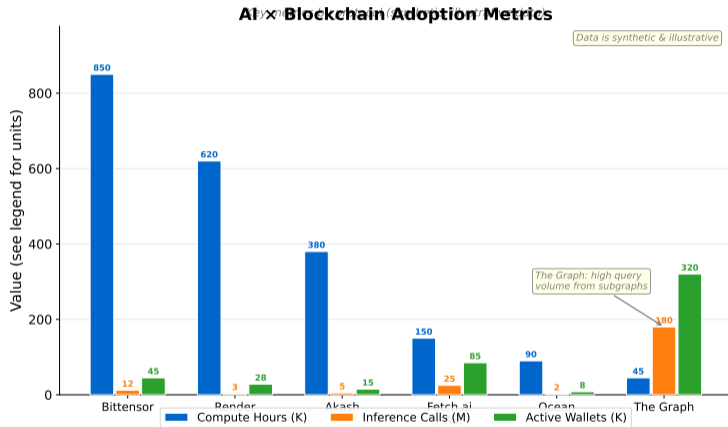
AI training and blockchain verification both consume energy. Combining them doubles the environmental question.

AI × Crypto Funding Growth



- **What you see:** Stacked bar chart of AI-crypto venture funding by category and year
- **Key pattern:** Funding exploded after 2023, driven by the ChatGPT hype cycle; “AI agents” emerged as a new category in 2024
- **Takeaway:** Money is pouring in, but most funding goes to infrastructure (compute) rather than applications

VC funding in AI-crypto projects exceeded \$5B in 2024. Whether this produces real value remains to be seen.



- **What you see:** Multi-metric dashboard showing users, compute hours, and inference calls for leading projects
- **Key pattern:** The Graph dominates inference calls (subgraph queries); Akash and Render lead compute hours; user counts remain modest
- **Takeaway:** Usage is growing but still a rounding error compared to centralized alternatives (AWS processes millions of GPU-hours daily)

Adoption is real but small. Decentralized AI compute is roughly 0.1% of AWS's GPU capacity.

Where are the projects?

Region	Share of Projects
United States	~40%
Asia-Pacific	~25%
Europe	~20%
Other	~15%

Regulatory arbitrage shapes location:

Strict US crypto regulation pushes token issuance to Singapore, Dubai, and Zug — while keeping engineering teams in Silicon Valley.

Result: a project may have US developers, a Swiss foundation, and Singapore-based token operations.

Key hubs and their advantages:

Silicon Valley (US)

- AI talent concentration
- Proximity to compute hardware makers
- Venture capital ecosystem

Zug / Singapore

- Crypto-friendly regulation
- Clear token classification frameworks
- Foundation-friendly legal structures


London / Berlin

- Financial infrastructure expertise
- Growing regulatory clarity (MiCA)
- Strong academic AI research

The most successful projects are geographically distributed — picking the best jurisdiction for each function.

AI talent clusters in the US. Crypto-friendly law clusters in Switzerland and Singapore. Successful projects span both.

Stakeholder Impact Analysis



Stakeholder	Impact	Why
Cloud providers	Lose	Market share erodes
GPU owners	Win	Monetize idle hardware
AI researchers	Win	Cheaper compute access
GPU manufacturers	Win	Demand from both sides
End users	Win	Cheaper, censorship-free AI
Regulators	Mixed	New compliance challenges

The redistribution thesis:

Centralized AI concentrates value at platform companies (OpenAI, Google, Amazon). Decentralized AI redistributes it to hardware owners, model creators, and data providers.

Whether this redistribution actually happens at scale remains unproven.

- **What you see:** Stakeholder impact scores ranging from strong negative to strong positive
- **Key pattern:** GPU makers and individual GPU owners benefit from both sides; cloud providers face disruption
- **Takeaway:** Decentralized AI redistributes value from platforms to hardware owners and end users

Decentralized AI promises to redistribute value from platforms to participants. The promise is ahead of the evidence.

The Disruption Map: Industries Reshaped

Industry	Current State	AI × Blockchain Disruption	Timeline
Cloud computing	3 providers control 65%+	Decentralized GPU marketplaces offer 60–85% savings	2025–2028
Content creation	Deepfakes indistinguishable from real	On-chain provenance verifies authenticity	2025–2027
Financial services	Human traders + centralized exchanges	Autonomous AI agents trade 24/7 on DEXs	2024–2026
Healthcare	Data silos prevent cross-hospital ML	Federated learning + blockchain tracks contributions	2027–2030
Legal	Manual contract review (\$500/hr)	AI audits smart contracts in minutes	2025–2027

Pattern: Industries with high trust requirements, data sensitivity, and intermediary costs are most vulnerable to AI-blockchain disruption. *Financial services are first because they are already digitized. Healthcare is last because regulation is strictest.*

AI-blockchain disruption hits digitized, intermediary-heavy industries first. Physical-world industries lag by 3–5 years.

Before investing, building, or believing the hype — ask:

1. Does the project NEED both AI AND blockchain?

Or would one technology alone suffice? If you remove the blockchain and the product still works, the token is cosmetic.

2. What specific AI problem does the blockchain solve?

Data provenance? Decentralized compute? Verifiable inference? If the answer is vague, the integration is shallow.

3. Is the token integral to the mechanism, or just bolted on for fundraising?

A token that incentivizes compute providers is integral. A token that “governs the AI” with no clear mechanism is decoration.

4. Can users verify the AI's outputs independently?

If the AI is a black box and the blockchain just records unverified outputs, you have gotten no benefit from decentralization.

5. What happens when the AI is wrong — is there a rollback mechanism?

Immutable mistakes are the worst kind. Does the protocol have circuit breakers, dispute resolution, or insurance?

If a project cannot answer all five, it is probably buzzword stacking.

These five questions separate genuine AI-blockchain innovation from marketing that combines two trending keywords.

The Integration Spectrum: From Buzzwords to Real Value

The spectrum:

Level 1: Buzzword stacking

“AI + Blockchain” in the name. No technical integration. Token is a fundraising vehicle. Whitepaper mentions both but connects neither.

Level 2: Useful combination

Token incentivizes data sharing or compute provision. Blockchain tracks contributions. AI processes the data. Each part does something real.

Level 3: Genuine innovation

zkML proves inference correctness on-chain. Federated learning uses blockchain for coordination. Decentralized compute achieves price parity with AWS. Neither technology alone could achieve this.

Examples at each level:

Level 1:

- “AI token” with no AI functionality
- Centralized API with a governance token
- “Decentralized GPT” that runs on one server

Level 2:

- Bittensor (token incentivizes model hosting)
- Ocean Protocol (token unlocks data access)
- Akash Network (token pays for compute)

Level 3:

- EZKL (ZK proofs for neural network inference)
- Ritual (verifiable AI inference on-chain)
- Numbers Protocol (content provenance at scale)

Most projects are Level 1. A few reach Level 2. Level 3 is rare and early.

Ask yourself: which level is this project? If you cannot tell, it is probably Level 1.

What Comes Next: The Convergence Roadmap

Near-term (2025–2026):

- AI agents become common in DeFi trading and governance
- zkML reaches production for simple models (credit scoring, identity)
- Compute marketplaces achieve 50% price parity with cloud for batch jobs
- Content provenance standards (C2PA) gain browser integration

The infrastructure phase.

Medium-term (2027–2029):

- AI-native blockchains launch (designed for ML workloads from day one)
- Content provenance becomes standard for news and media
- Decentralized training of medium-sized models becomes feasible
- Regulatory frameworks for AI agents emerge (EU AI Act + MiCA convergence)

The adoption phase.

Long-term (2030+):

- Autonomous agent economies: AI-to-AI commerce on-chain
- AI-governed DAOs make complex decisions without human boards
- Machine-to-machine payments for compute, data, and inference
- Personal AI agents manage your on-chain identity, assets, and interactions

The transformation phase.



The roadmap follows a classic technology adoption curve: infrastructure first, applications second, transformation last.

One Last Thought...

We started with a question: *can AI and blockchain — opposites in almost every way — combine into something greater?*


The answer is nuanced:

Yes — where the combination solves a real problem. Decentralized compute breaks monopolies. zkML enables verifiable AI. Content provenance fights deepfakes.

No — where the combination is cosmetic. A governance token on a centralized AI model adds complexity without benefit.

The convergence is real, but most of it is still Level 1 and Level 2. Level 3 — genuine technical innovation that neither technology could achieve alone — remains rare.

The future belongs to projects that pass all five evaluation questions. Most will not.



AI and blockchain are powerful apart. Together, they are powerful only when the integration solves a real problem.

Seven ideas to remember from today:

1. **Convergence thesis:** AI is powerful but opaque; blockchain is transparent but dumb — their weaknesses are complementary
2. **Decentralized compute:** GPU marketplaces (Akash, Render) offer 60–85% savings over cloud providers, with trade-offs in reliability and tooling
3. **Tokenized inference:** Pay-per-query AI eliminates subscription lock-in and enables censorship-resistant access to models
4. **AI agents:** Autonomous on-chain software that observes, decides, and acts — already profitable in MEV extraction, expanding to portfolio management
5. **zkML:** Zero-knowledge proofs for AI verify that a model produced a specific output without revealing its weights — the holy grail of verifiable AI
6. **Deadly combination:** AI hallucinations + blockchain immutability = irreversible mistakes with no rollback
7. **Five-question test:** Does it need both? What problem is solved? Is the token integral? Can outputs be verified? What if the AI is wrong?

If a project cannot answer all five evaluation questions, it is probably buzzword stacking — not innovation.

Seven takeaways from three pillars: compute, agents, and verification. The five-question test is your decision tool.

Today's journey:

We explored the intersection of AI and blockchain through three pillars: decentralized compute that breaks cloud monopolies, AI agents that act autonomously on-chain, and verification systems that prove AI correctness without revealing models.

The field is young, overhyped, and genuinely promising — all at the same time. Most projects are buzzword stacking. A few are building infrastructure that could reshape how AI is trained, deployed, and governed.

Next: Explore specific projects in depth and build your own evaluation framework for AI-blockchain ventures.

Key Vocabulary:

- Convergence thesis
- Decentralized compute
- GPU marketplace
- Tokenized inference
- AI agent
- zkML
- Content provenance
- C2PA standard
- Compute-to-data
- Federated learning
- Oracle manipulation
- Hallucination risk
- Centralization creep
- Buzzword stacking
- Integration spectrum

Five reflection questions:

1. Which AI service do you use most — and how decentralized is it?
2. Would you trust an AI agent to manage your crypto portfolio?
3. Should AI-generated content be required to carry provenance metadata?
4. Is “decentralized AI” truly decentralized if it runs on NVIDIA GPUs?
5. Pick one project from today — does it pass all five evaluation questions?

Apply the five-question framework to every AI-blockchain project you encounter. Most will fail at question one.