

BARANOS AI

A Deterministic, Verifiable LLM Deployed on Fogo

Neutral intelligence as blockchain infrastructure

• DETERMINISTIC

Same committed job → same canonical result

• CHALLENGEABLE

Incorrect execution can be disputed and resolved

• COMPOSABLE

Results can become state consumed by smart contracts

THESIS

Blockchains made computation neutral. Baranos AI extends that property to inference: a committed model, committed inputs, deterministic execution rules, and a result that can be independently reproduced, challenged, and resolved.

Version 2.0 / September 2026

Current deployment: Fogo

Status note: deployment, model-performance, cost, challenge-resolution, and cross-chain capability claims should be supported by reproducible public benchmarks before being treated as independently verified production guarantees.

Abstract

Baranos AI is a verifiable onchain inference system deployed on Fogo that enables large language model outputs to be reproduced, challenged, settled, and consumed as blockchain state. Each inference job is defined by commitments to a specific model and weights, tokenizer and runtime, inputs, prompt construction, decoding parameters, output constraints, and execution policy, producing a canonical computational result under a deterministic specification. In the normal Confirmation mode, computationally intensive inference is executed offchain and the result, together with the required execution commitments or verification artifacts, is posted onchain for verification and settlement. If a result is challenged, the disputed computation can be replayed and adjudicated onchain. Replay mode provides fully onchain execution when that execution path is required.

This architecture combines efficient offchain inference with blockchain-native verification, challenge resolution, and composability. Settled results can become durable onchain state available to smart contracts and applications, allowing software to rely on a committed model computation without depending solely on the operator that performed the inference. Baranos is designed around approximately minute-scale result and challenge resolution for the current deployment, subject to workload and benchmark conditions.

The initial users envisioned for Baranos are prediction-market applications that operate under their own terms of service and policies and require deterministic resolution of natural-language questions from committed evidence. The same verifiable-inference architecture can extend to other applications that require reproducible, challengeable AI outputs as part of onchain workflows.

THE CORE PRODUCT

Baranos AI is not best understood as “a chatbot on a blockchain.” It is a trust-minimized execution layer for AI inference. Its normal Confirmation mode executes inference offchain and verifies/settles the result onchain using an optimistic challenge mechanism; disputed computation can be replayed onchain, while Replay mode provides fully onchain execution.

1. From Smart Contracts to Verifiable Intelligence

Smart contracts are powerful because their execution is public, deterministic, and shared. The moment an application needs interpretation rather than arithmetic, however, that trust model usually breaks. A protocol calls a private AI API, relies on a committee, asks token holders to vote, or accepts an offchain answer from a privileged operator.

Baranos AI is designed to close that gap. Its objective is to make model inference behave more like a smart-contract state transition: the model is identified, the inputs are committed, the execution policy is fixed, and the output can be verified.

This is particularly important when AI output can move money, resolve a market, grant permission, classify collateral, trigger insurance settlement, or direct an autonomous agent. In those settings, “the API said so” is a weak foundation. Baranos replaces that with a procedure that can be inspected and challenged.

BARANOS THESIS #1

AI becomes materially more useful to blockchains when the execution of intelligence can inherit blockchain-style guarantees.

2. What Deterministic LLM Inference Means

A deterministic inference job has one canonical computational outcome. In practical terms:

same committed model + same weights/state + same input + same execution rules → same result

This is stronger than receiving a similar answer twice. For a blockchain-grade system, determinism means that independent parties can agree on the exact computation that should have occurred.

2.1 Why ordinary LLM inference can vary

- Sampling can deliberately introduce randomness through temperature, top-p, top-k, or other decoding choices.
- Different numerical formats, kernels, hardware paths, parallel schedules, or runtime implementations can produce small arithmetic differences that compound through generation.
- A provider can silently change weights, tokenizers, system prompts, quantization, context construction, or decoding defaults.

2.2 How Baranos defines a canonical job

A Baranos job binds the variables that determine inference: model commitment, tokenizer and runtime version, numerical/quantization profile, prompt construction, input/evidence commitment, decoding parameters, seed where relevant, context/output limits, and execution policy. Oracle-grade policies can require greedy or otherwise deterministic decoding.

The result is no longer an informal request to “ask an AI.” It is a defined computation over committed state.

3. Verification, Challenges, and Fast Resolution

The defining architectural insight is to separate execution from settlement without giving up enforceability. Baranos uses an optimistic design rather than a zero-knowledge proof system: in the normal Confirmation path, expensive inference executes offchain and its result is posted and verified onchain. The protocol does not require every validator to repeat every multiply on the happy path. Instead, any disputed portion can be replayed onchain under the same committed deterministic specification, allowing the chain to adjudicate the dispute.

3.1 The basic flow

- 1 An application creates an inference job under a defined policy that commits to the model, inputs, prompt template, evidence rules, decoding settings, payer, deadline, permitted output structure, and execution mode: Confirmation or Replay. The initial product is designed around structured, policy-governed application workflows rather than an unrestricted general-purpose prompt interface.
- 2 In Confirmation mode, an eligible executor performs the inference offchain and posts the result plus the required execution receipt, commitment, or verification artifact onchain. In Replay mode, the computation is executed fully onchain.
- 3 For Confirmation-mode jobs, the posted result enters an optimistic settlement path: absent a valid challenge it can be confirmed according to the job policy; if challenged, the disputed computation is escalated to onchain replay.
- 4 The chain adjudicates a challenge by replaying the disputed portion under the committed execution specification and determining which result conforms to that specification.
- 5 Once confirmed or resolved, the canonical result becomes onchain state that downstream programs can consume.

The current Baranos deployment is designed around approximately minute-scale result/challenge resolution. That latency matters: verification becomes useful for interactive markets and applications rather than only for slow, post-hoc auditing.

WHAT VERIFICATION PROVES

Verification proves execution integrity: “this is the result produced by model M on input X under policy P.” It does not prove that the model’s statement is factually true, unbiased, safe, or wise. A perfectly verified model can still be wrong.

4. Current Deployment: Fogo

Baranos AI is deployed on Fogo. The project selected Fogo because the workload demands a combination of high throughput, low latency, predictable execution, and low transaction cost that is difficult to achieve on conventional blockchain infrastructure.

Fogo is an SVM Layer 1 built around a Firedancer-based validator implementation and performance-oriented consensus. Baranos uses Fogo as its underlying consensus and settlement layer for verifiable inference.

Baranos’ current engineering position is stronger: for the present implementation and workload, the team considers Fogo the only blockchain environment tested that simultaneously provides the throughput headroom, economics, and latency needed to return and settle a verifiable LLM result in approximately one minute. That “only blockchain” statement should be read as a project-supplied deployment conclusion, not as an independently established universal benchmark across every chain.

4.1 Why performance matters for AI

- Inference jobs generate more data movement and computational coordination than ordinary token transfers.
- Optimistic challenge systems add commitments, receipts, verification steps, and—only when a result is disputed—onchain replay and adjudication.
- Interactive applications cannot tolerate settlement that takes hours.
- If verification costs are too high, applications fall back to trusted offchain servers and lose the core benefit.

5. Architecture

Baranos separates inference into independently committed objects so that model selection, evidence, policy, execution, verification, and settlement cannot be silently bundled behind an opaque endpoint. A key design property is that every component can execute onchain, even when the normal path runs computationally intensive work offchain. Jobs reference explicit inference policies rather than relying on hidden or mutable execution settings. Each accepted job can also be attributable to its originating application through an operator or application identifier, creating durable provenance without requiring end-user identity to become part of the inference protocol.

OBJECT	ROLE
Model Registry	Model identifier, weight commitment, tokenizer/runtime metadata, quantization profile, activation state.
Evidence / Input Registry	Documents, structured inputs, canonical hashes, timestamps, provenance and availability policy.
Inference Policy	Prompt template, decoding rules, context/output bounds, deterministic settings, challenge and settlement rules.
Inference Job	Binds originating operator/application identifier, market identifier, model, input root, policy, query or query commitment, payer, deadline and execution mode (Confirmation or Replay).
Result Account	Output commitment/plaintext, execution receipt or verification artifact, confirmation/challenge state, replay status and settlement metadata.

5.1 Execution modes: Confirmation and Replay

Baranos exposes two execution modes. Confirmation is the normal, optimistic verification mode: inference executes offchain, the result and required proof/receipt artifacts are posted onchain, and the chain verifies and settles the result. If the result is challenged, the disputed portion is replayed and adjudicated onchain. Replay is the heavyweight mode: the relevant computation is executed fully onchain from the outset. This dual-mode design is central to the trust model. Offchain execution improves cost and latency on the happy path, while universal onchain executability provides a credible fallback that keeps executors honest. Conceptually, this resembles an offchain matching engine whose results are settled by a chain, except that Baranos can deterministically replay disputed AI computation onchain. The design is optimistic rather than ZK-based: correctness is enforced through commitments, challenge rights, and replay, not by requiring a zero-knowledge proof of every inference.

6. The Current Model: An OpenWeight LLM

The current Baranos deployment uses an Open Weight LLM. The nationality or origin of the model is not the core protocol property; the important point is that Baranos can bind a specific model version and execution environment into a verifiable job.

This makes the deployment a proof of a broader abstraction: the infrastructure is not inherently tied to one vendor, one model family, one language, or one geography. A future Baranos registry can support multiple models concurrently, provided each model has an explicit commitment, reproducible runtime semantics, and a policy that defines how deterministic execution is achieved.

Historical jobs remain pinned to the exact model and policy under which they settled. Upgrading the preferred model therefore does not rewrite the meaning of previous results.

7. Baranos as an AI Oracle

The initial application category we envision is prediction markets whose resolution questions cannot be reduced to a price feed. We expect early users in this category to operate under their own terms of service and policies governing their markets and users. A market can precommit to a model, evidence policy, prompt template, deterministic decoding rules, permitted output schema, and deterministic resolution procedure. End users interact with the market; Baranos functions as infrastructure for the market's structured inference workflow rather than as the market's consumer-facing chatbot.

- **Prediction markets:** determine whether a natural-language event condition was satisfied from a frozen evidence set.
- **Insurance:** evaluate submitted evidence against precommitted policy language.
- **Governance:** classify proposals, check policy compliance, or generate structured determinations from public records.
- **Corporate-action and disclosure monitoring:** classify filings, announcements, or event evidence.

The protocol does not eliminate interpretation. It makes the interpreter, evidence, and rules of interpretation explicit before settlement.

8. Initial Application Scope and Future Integrations

Initial prediction-market users

We envision the initial users of Baranos to be prediction-market applications that maintain terms of service and policies governing their markets and users. The protocol itself focuses on deterministic inference, verifiability, evidence commitments, and resolution mechanics rather than prescribing how those applications govern access to their own products.

Registered inference policies

A registered inference policy can bind the model and runtime, prompt template, evidence-source rules, input and output bounds, deterministic decoding settings, permitted output schema, challenge rules, settlement rules, and execution mode. These policies define the computation that Baranos is expected to perform and make the relevant resolution procedure explicit and reproducible.

Public evidence and committed inputs

Prediction-market resolution should prefer public, attributable evidence and committed input roots. Raw prompts, documents, or sensitive evidence need not be published as arbitrary plaintext merely because the resulting computation is verifiable; the protocol should commit only the data required by the selected policy and verification path.

Application policies and resolution criteria

Prediction-market applications can define the market and resolution criteria under which they use Baranos, including how resolution questions are framed, what evidence is admissible, and what terms and policies apply to their users and markets. Baranos inference policies can then encode the deterministic computational procedure used for a particular job. This separation keeps application-level governance distinct from the protocol's role in executing, verifying, and settling the committed computation.

Operator identity and provenance

An integration may be associated with an operator or application identifier. An inference job can record or cryptographically bind that originating identifier, a market identifier, and the applicable inference or resolution policy. This creates a persistent provenance trail linking a settled result to the application that requested it and the policy under which the computation was performed. The exact registry and attestation mechanics can evolve as the production framework develops.

Application-level terms, market rules, and user policies remain the responsibility of the integrating application. Baranos remains responsible for the deterministic execution, verification, challenge, and settlement mechanics defined by the applicable inference policy. This separation preserves provenance and accountability without making Baranos the operator of the underlying prediction market.

Future integrations

Additional application categories such as insurance, governance, risk engines, or autonomous systems may use the same architecture as their security, data-availability, privacy, and challenge requirements are defined.

Open verification, structured job creation

Baranos can preserve transparent verification and challengeability while supporting structured, policy-defined inference jobs. Verification of a result is conceptually separate from the application-level terms and policies governing the product in which that result is used.

Composable settled results

Applications and downstream contracts may consume settled Baranos results as blockchain state, subject to the interfaces and inference policies used by the deployment.

No unrestricted public prompt endpoint

The initial Baranos product is infrastructure for bounded, policy-governed inference, not a consumer-facing public chatbot or an unrestricted askAI(prompt) primitive. Broader interfaces can develop alongside explicit safety, data-handling, and application-policy frameworks.

9. The Trust Model

Baranos changes what an application must trust. A conventional AI API asks the caller to trust the provider's infrastructure and configuration. Baranos instead uses an optimistic verification model built around a public execution specification, cryptographic commitments, economic incentives, onchain challenge/settlement, explicit inference policies, and persistent job provenance. The normal Confirmation path executes inference offchain, but the chain can replay and adjudicate any disputed portion. Replay mode can execute the computation fully onchain when that stronger, heavier-weight path is desired. Application-level terms and policies remain separate from deterministic verification of the jobs submitted to the protocol.

SYSTEM	WHAT THE APP TRUSTS	WHAT CAN BE AUDITED
Private AI API	Provider + hidden configuration	Usually limited
Human / token oracle	Committee or economic voting incentives	Votes and published evidence
Baranos AI	Committed model/policy + optimistic verification + onchain replay fallback	Model, inputs, policy, result, challenge and replay path

9.1 Fogo Consensus, Validator Governance, and the FOGO Token

Baranos does not ask FOGO holders or Fogo validators to decide what an AI model should answer. The model commitment, inputs, runtime, and deterministic inference policy define the computation. In Confirmation mode, executors compute offchain and post results onchain; challengers can dispute a result, causing the relevant computation to be replayed and adjudicated onchain. Replay mode performs the computation fully onchain. Fogo provides the underlying consensus and settlement layer that determines which protocol state is canonical.

This distinction matters for AI governance. Fogo validators are not intended to vote on individual model outputs or substitute social judgment for deterministic execution. Their role is lower in the stack: they secure the chain on which Baranos jobs, results, challenges, slashing events, model-registry changes, and protocol upgrades are recorded.

Fogo uses proof-of-stake consensus in which validator voting weight is based on FOGO stake, including delegated stake. Fogo also operates with a limited validator set subject to minimum stake and operational requirements. Its published architecture describes validator governance as a mechanism for enforcing network health, including the ability to remove persistently underperforming or harmful validators. Protocol-level governance can require a two-thirds supermajority of staked tokens.

For Baranos, this creates a second security layer around verifiable inference. The Baranos protocol determines whether an inference was executed correctly; Fogo consensus determines the canonical history in which that determination settles. If Baranos becomes infrastructure for autonomous agents, markets, treasuries, or other economically consequential systems, the economic security of Fogo therefore becomes part of the security budget of onchain intelligence.

FOGO holders can participate in that security model by staking or delegating FOGO to validators. This should not be described as a direct vote over what Baranos is allowed to think or say. A more precise formulation is that FOGO provides economic participation in, and security for, the validator infrastructure that ultimately settles Baranos state. In that sense, Fogo supplies consensus over what the AI computation actually was—not editorial control over what the AI ought to conclude.

10. Security and Failure Domains

Verifiable execution does not remove model, data, or publication risk. Baranos therefore treats policy definition, job bounds, evidence handling, and publication safeguards as part of the security model, alongside the following security domains:

Model integrity — weights and runtime artifacts must match their registered commitments.

Prompt and policy integrity — instructions, decoding settings, and challenge rules must be versioned and immutable for settled jobs.

Input integrity — every document or data item used by the model must be committed under the job's input root.

Execution integrity — the accepted output must conform to the canonical computation.

Data availability — applications must define how long raw evidence remains retrievable.

Prompt injection and adversarial evidence — untrusted content must be treated as data rather than privileged instructions.

Model risk — hallucination, bias, weak reasoning, adversarial examples, and domain error remain possible even when execution is perfect.

Resource and publication safety — job bounds, registered policies, constrained prompt templates, and evidence rules can reduce adversarial workloads and the risk of placing inappropriate or sensitive arbitrary content into permanent blockchain state.

11. Economics

A verifiable AI network must make the honest path economically preferable to both trusted centralization and adversarial execution. Jobs therefore need explicit pricing for execution, storage, verification, challenge capacity, and optional data availability.

The developer abstraction should remain simple: specify the model job, latency/security policy, and maximum fee; the protocol handles the eligible execution and settlement path. Fogo's low-latency, high-throughput design is central to the current Baranos implementation because excessive chain fees or congestion would erase the advantage of trust-minimized inference.

11.1 FOGO as the Settlement and Security Asset

Baranos inference creates native demand for Fogo blockspace. A typical Confirmation-mode optimistic flow can require an onchain job request, a posted completion and verification artifact, and—only when necessary—a challenge followed by onchain replay of the disputed computation and dispute resolution. Replay-mode jobs are heavier because execution is fully onchain. Downstream applications may then execute additional transactions based on the settled result.

Those transactions consume FOGO-denominated gas. Under Fogo's published fee design, base transaction fees are low, while users can attach priority fees when faster inclusion has greater economic value. Priority fees accrue to the block producer. This is especially relevant to AI workloads because the value of latency is heterogeneous: a low-value classification job may tolerate delay, while a liquidation, arbitrage, market-resolution, or autonomous-agent decision may rationally pay more for urgent inclusion.

Baranos can therefore create two distinct economic markets: an inference market that compensates executors and challengers for correct computation, and a blockspace market in which FOGO pays for settlement, security, and latency. These layers should remain conceptually separate. A future Baranos-specific asset, stablecoin bounty, or other payment mechanism could price AI computation without replacing FOGO's role as the native gas and staking asset of the underlying chain.

If Baranos usage scales, the relevant FOGO demand is not simply the number of integrations. It is the aggregate chain activity generated by job requests, completions, challenges, registry operations, and AI-directed downstream actions, multiplied by the mix of base and priority fees those applications are willing to pay. This relationship should be measured rather than assumed and included in future public benchmarks.

12. Performance and Benchmarking

The central Baranos claims are measurable and should be published as reproducible benchmarks. In particular, the project should disclose enough information for third parties to reproduce Confirmation-mode happy-path jobs, challenged jobs with onchain replay, and full Replay-mode execution.

BENCHMARK	METRIC	REQUIRED DISCLOSURE
Inference latency	Time from accepted job to posted result	model, precision, context, output tokens, executor hardware
Challenge resolution	Time from challenge to canonical settlement	challenge mechanism, disputed unit, verifier hardware, network conditions
Cost	Total chain + execution + verification cost	fee assumptions, compute provider, storage/data availability
Determinism	Bit-/token-level agreement across independent runs	runtime, kernels, numerical format, decoding policy
Concurrency	Jobs/sec and p95/p99 latency under load	network load, model mix, scheduler policy

CURRENT TARGET

Baranos is deployed on Fogo and is designed to return a Confirmation-mode result and, where challenged, resolve it through onchain replay in approximately one minute. Public benchmark artifacts should separately define Confirmation happy-path latency/cost, challenged replay conditions, and full Replay-mode execution, including model size, context, output length, hardware, fees, and finality.

13. Why This Is More Than “Onchain AI”

Running every neural operation onchain on every job is not the breakthrough. The important combination is deterministic specification, optimistic offchain execution on the happy path, onchain verification and settlement, challengeability, onchain replay of disputed computation, an optional fully onchain Replay mode, fast settlement, and composability.

If those properties hold, the output of intelligence becomes something another program can safely depend on. A smart contract can bind itself in advance to a model's canonical answer in much the same way it binds itself to ordinary program logic.

That turns Baranos from a model demo into infrastructure. The long-term opportunity is a policy-governed intelligence layer in which models are registered, applications submit bounded jobs under explicit inference policies, inference can be executed competitively offchain under Confirmation mode, disputed computation can be replayed onchain, Replay mode remains available for fully onchain execution, and settled results become shared state. The architecture separates verifiable computation from the application-level terms and policies governing how that intelligence is used.

14. Roadmap

Phase 0 – Fogo deployment and initial application focus

Operate the current Open Weight LLM on Fogo; validate Confirmation-mode offchain execution and onchain verification, challenged onchain replay, full Replay mode, settlement latency, determinism, cost, and contract-level integration. Initial production work is focused on prediction-market use cases.

Phase 1 – Public benchmarks and policy specification

Publish reproducible end-to-end measurements, model/runtime commitments, challenge traces, cost disclosures, and the initial inference-policy specification. Define the policy fields that constrain prompt templates, evidence sources, decoding, context/output bounds, output schemas, challenge rules, and settlement behavior. Document how prediction-market applications can map their own terms, market rules, and resolution criteria into deterministic inference jobs.

Phase 2 – Prediction-market integrations

Launch production integrations with prediction markets. End users interact with those markets rather than querying Baranos as a consumer-facing chatbot. Each market can submit structured jobs under defined inference policies, using frozen or otherwise policy-compliant evidence and deterministic resolution procedures. Each accepted job can be attributable to the operator/application and market identifiers associated with the originating integration.

Phase 3 – Hardened policy and data controls

Expand inference-policy enforcement, evidence commitments, data-availability rules, publication safeguards, and monitoring. Establish explicit rules for what information is committed onchain versus retained through an appropriate evidence-availability mechanism.

Phase 4 – Multi-model registry and additional integrations

Support additional open-weight models with explicit versioning and historical pinning. Extend the architecture to additional application categories with application-specific inference policies as their requirements become defined.

Phase 5 – Open verification and broader composability

Broaden executor and challenger participation where technically and economically appropriate, and allow supported applications to consume settled inference as a native dependency. Continue developing structured interfaces and policy frameworks as the range of applications expands.

15. Conclusion

Blockchains turned execution into a public fact. Baranos AI applies the same principle to machine inference while separating verifiability from unrestricted query access.

The model is known. The committed inputs are known to the verification procedure. The execution policy is known. The intended computation is deterministic. We envision the initial users to be prediction markets that maintain terms of service and policies governing their markets and users, while Baranos provides the verifiable inference infrastructure beneath those applications. In the normal Confirmation mode, inference executes offchain and the result is verified and settled onchain through an optimistic mechanism. If challenged, the disputed portion can be replayed and adjudicated onchain. Replay mode can instead execute the computation fully onchain. A settled output can then be consumed by another program.

Baranos AI is deployed on Fogo because the project requires unusually high throughput, low cost, and low latency. For the current workload, the team's engineering assessment is that Fogo is the only tested blockchain environment capable of satisfying all three requirements while targeting approximately minute-scale result and dispute resolution.

If the deployment continues to validate those properties under reproducible load, the result is not simply an onchain LLM. It is a new blockchain primitive: optimistic, verifiable intelligence that can use efficient offchain execution without asking software to trust the machine or company that produced the answer, because disputed computation remains reproducible onchain. The initial product is designed around bounded, policy-defined application workflows, reducing the privacy, abuse, and immutable-data risks associated with arbitrary public prompt submission.

The deployment also creates an important separation of responsibilities. Baranos governs the specification and verification of deterministic intelligence; Fogo supplies canonical settlement and proof-of-stake economic security. FOGO holders and delegators therefore do not choose individual AI answers, but they participate economically in the validator layer that secures the history in which those answers become authoritative onchain state. If autonomous systems increasingly depend on Baranos, that settlement layer becomes part of the governance and security architecture of verifiable intelligence. For prediction-market jobs, application and market identifiers can provide persistent provenance while the integrating market remains responsible for its own terms, policies, and market rules.

Fogo protocol references for Sections 9.1 and 11.1: Fogo Architecture documentation (docs.fogo.io/architecture.html), Fogo Protocol Litepaper v2 (api.fogo.io/litepaper.pdf), Fogo whitepaper (fogo.io/whitepaper.pdf), and FOGO Tokenomics (fogo.io/blog/introducing-fogo-tokenomics--performance-without-compromise).

END STATE

A smart contract should be able to call intelligence with the same confidence that it calls code.