

House Party Protocol (HPP)

Whitepaper · English Version 2.0 · 2026

House Party Protocol (HPP)

Whitepaper · Version 2.0 · 2026

1. Introduction

Until now, blockchains have been passive infrastructure for recording transactions between people. House Party Protocol (HPP) is active infrastructure where agents act and create value on their own.

Autonomous AI agents that decide, execute, and pay are no longer a future concept. AI has moved beyond computation and inference to become a genuine economic actor that creates and consumes value. Yet most existing infrastructure still assumes a human is in the loop: payment requires an account and a card, trust depends on reputation, and execution assumes human oversight. HPP Mainnet is an AI-native execution layer built to remove that assumption. On HPP, an agent's on-chain execution, inference, and payment complete as one continuous flow.

HPP's technical foundation combines two track records: Blocko, which has built enterprise blockchain since 2014 and runs mission-critical, national-scale systems such as Korea's National Health Insurance Service (NHIS) on its Aergo Enterprise platform; and Aergo, a public chain in production since 2019. Aergo Hub was an early experiment in infrastructure that simplified application deployment and operation. Building on that experience, HPP takes on the problem of trustworthy trading of idle resources in the AI inference market.

Whitepaper v1.0 covered HPP's tokenomics, core technology, and product suite. This paper defines the three conditions an agent economy needs (payment, trust, and execution) and explains how HPP implements them across its network, products, verification layer, token, and governance. It also introduces HPPLite, a sovereign rollup framework for SQLite applications, and the ecosystem application strategy built on it.

2. Market Problem: Three Structural Limits

AI agents have become both suppliers and consumers of services in their own right. The moment they try to act, they run into three structural limits.

Payment

Existing payment systems assume a human: accounts, cards, subscriptions, identity checks. Autonomous agents can't use them directly. Inference calls are small and frequent — a micro-transaction pattern that today's payment rails, with their fees and settlement cycles, were never built for. Agents are left borrowing a human's payment method and API key.

Trust

When compute and inference become the unit of trade, buyers can't verify what they bought. Which model ran, how many tokens were actually consumed, whether the response was altered in transit — none of it is verifiable. That uncertainty gets priced in as a discount, capping what an honest supplier can charge regardless of quality. The distortion is already visible: demand for cheap inference is strong enough to have spawned an informal resale market outside official channels — one where buyers cannot verify integrity and have no guarantee the service will still be there tomorrow.

Execution

Smart contracts cannot run AI models. There is no standard path for on-chain logic to request off-chain compute, verify the result, and pay for it. And there is no secure way to run an autonomous agent in an isolated environment under bounded permissions.

HPP addresses each limit with its own layer: x402, an open payment standard, for payment; an attestation, bond, and reputation structure together with Proof-of-Inference for trust; and the on-chain compute framework Noösphere for execution.



Figure 1. Three structural limits mapped to HPP layers

3. Network Architecture

3.1 Overview

HPP Mainnet, an Ethereum L2, is an Optimistic rollup built on the Arbitrum Nitro stack. State transitions and dispute resolution anchor to Ethereum Mainnet, and data availability runs through an AnyTrust committee for cost efficiency and throughput.

Item	Detail
Network	HPP Mainnet (Ethereum L2)
Technology stack	Arbitrum Nitro Stack
Settlement layer	Ethereum Mainnet
Data availability	AnyTrust
Dispute protocol	BoLD (fraud-proof based)
Chain ID	190415
Native gas token	ETH
Infrastructure operator	Conduit (Rollup-as-a-Service)

HPP is fully EVM-compatible. It carries no protocol-level changes to the upstream Nitro code; any chain-specific differences stay within officially supported configuration parameters. This avoids the added security risk of a custom fork and preserves the full benefit of upstream audits and battle-testing.

ETH is HPP Mainnet's native gas token. Network fees are paid in ETH; the HPP token is not gas; it powers ecosystem utility and governance instead.

3.2 Three-Layer Stack

Layer	Component	Role
Layer 1 (Security)	Ethereum Mainnet	Final settlement, fraud-proof verification, security anchor
Layer 2 (Execution / Settlement)	HPP Mainnet	Agent economy execution and settlement, ETH gas, HPP utility and governance, L3 state anchoring
Layer 3 (Application)	HPPLite (public) / BESU (enterprise)	Application-specific execution, settles to HPP Mainnet

A dual-L3 architecture, a public L3 and an enterprise L3 sharing one settlement layer, is a structural feature specific to HPP. Workloads with different requirements each run in the

execution environment suited to them, while all of them share transparent auditability on the same public settlement layer.

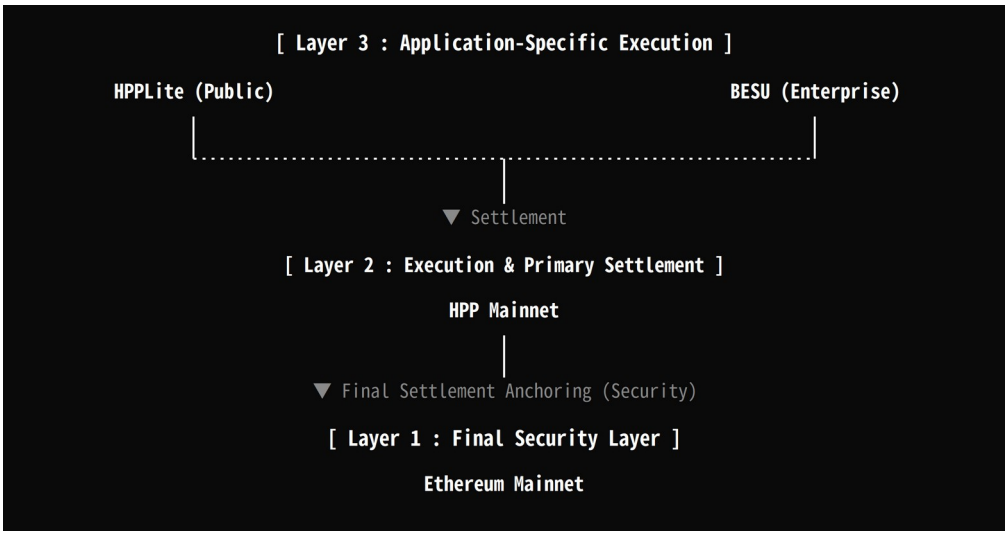


Figure 2. Settlement flow from Ethereum L1 through HPP Mainnet to L3

3.3 zk Settlement

In step with the upstream stack's own roadmap, HPP is pursuing zk settlement: adding zero-knowledge proofs to the settlement process.

4. The HPP Stack

HPP's products are not a list of separate services; they form one integrated stack. An agent pays through x402, infers through HPP Router, connects on-chain logic to off-chain compute through Noösphere, and records its state through HPPLite. This loop is where HPP's agent economy begins.

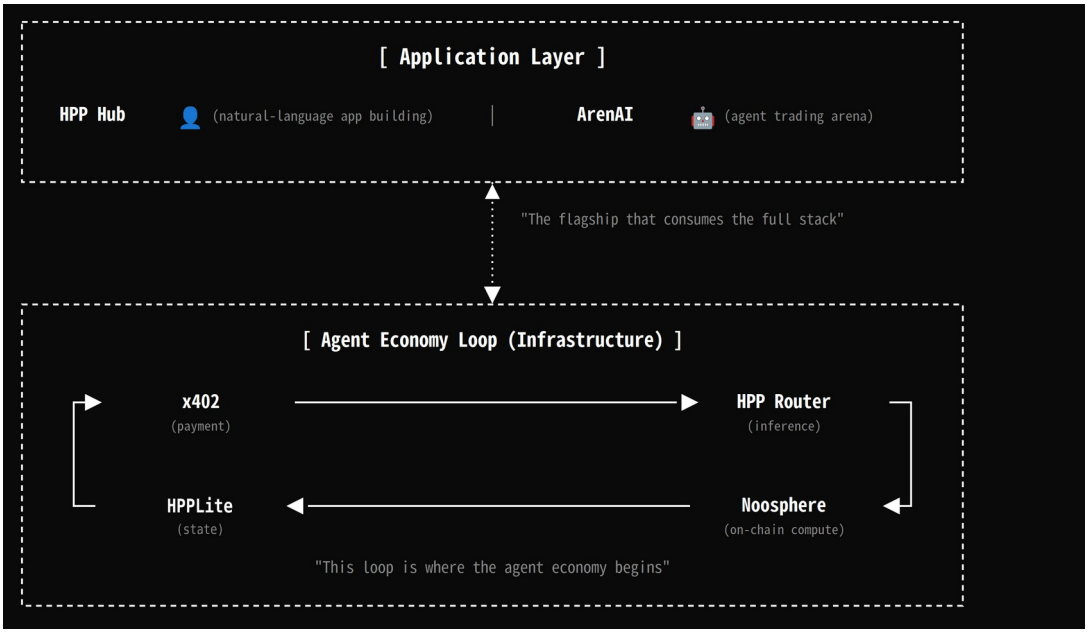


Figure 3. The HPP stack loop and application layer

Product	Layer
x402	Payment Layer
HPP Router / HPP Coder	Intelligence Layer
Noösphere	Compute Layer
HPPLite	State Layer
HPP Hub	Application Layer

4.1 ×402: Payment Layer

x402 is an open payment standard built on the HTTP 402 status code. An AI agent skips account creation and checkout entirely, a single signature completes an on-chain payment. HPP Mainnet is registered in the x402 asset registry as the default settlement chain for USDC (Chain ID 190415), so every agent and application in the ecosystem can transact using it as a common settlement currency. HTTP services from HPP Router and Noösphere agents can bill per call through x402. When a smart contract instead requests off-chain compute from Noösphere, the whole cycle, subscription, escrow, verification, and settlement, runs transparently through Noösphere's on-chain lifecycle.

4.2 HPP Router / HPP Coder: Intelligence Layer

HPP Router is a high-performance LLM inference router. It routes incoming requests to the best-fit model on the network and provides a unified portal for model discovery, API key management, and testing. Router is more than a routing tool; it is evolving into the demand-

side interface for this paper's core thesis: the verifiable inference marketplace.

HPP Coder is a coding LLM purpose-built for smart contract development, fine-tuned from an open-weight model on a Solidity corpus, and offered as one of the core models on HPP Router. Developers use it to generate, review, and audit smart contract code.

4.3 Noösphere: Compute Layer

Noösphere is HPP's on-chain compute framework. A smart contract requests off-chain compute as a subscription; decentralized agent nodes execute it in a containerized environment, and delivery, billing, and verification all run through an on-chain lifecycle. The request and the settlement both stay on-chain and auditable, that is what defines Noösphere.

Developers can build genuinely AI-native dApps: a smart contract that queries an AI model directly and acts on the verified response in its own logic. Recurring compute pipelines and verifiable randomness (NoösphereVRF) work the same way. A single agent node can earn from two channels at once, the on-chain subscription market and per-call sales through x402.

4.4 HPPLite: State Layer

HPPLite is the state layer that records the state of agents and applications on the HPP network in verifiable form. Once the extension is attached to a SQLite application, every write is posted in encrypted batches and a state root of the whole database is anchored periodically. State that needs to persist, such as an agent’s work history or the data of an application an agent has built, gains an ordered history without a separate chain. Chapter 6 covers the architecture, data availability, and trust model in detail.

4.5 HPP Hub: Application Layer

HPP Hub is an app-builder platform for building on the models HPP Router provides. Users build conversational apps, web apps, Web3 apps, and 2D or 3D games directly in Hub; for smart contracts, HPP Coder handles code generation and review.

Hub is where the stack's value reaches the end user directly. Router's inference becomes a finished application, so someone with no development background can still build on HPP.

5. Verifiable Inference Marketplace

HPP Router is expanding into a verifiable inference marketplace, the core of HPP 2.0. If HPP is the execution and settlement layer for the agent economy, this marketplace is the economic layer built on top of it. This chapter covers the problem it solves, what is actually being traded, and how trust gets built.

5.1 The Problem: Idle Capacity, No Way to Verify It

Compute capacity for AI inference sits idle more often than not. Few companies that build out GPU infrastructure keep it running at full utilization. A market for renting idle GPUs already exists and is growing fast: per CoinGecko, the decentralized physical infrastructure (DePIN) sector reached a market cap of about $19.2 billion in September 2025, up from $5.2 billion a year earlier. But renting raw compute and selling a verified result — proof that a specific model produced this output under specific conditions — are different products entirely. HPP is not targeting the compute-supply market itself; it is targeting the verification gap between the two.

For companies sitting on idle capacity, every idle hour is lost revenue. The problem is that a buyer transacting with an anonymous supplier has no way to verify the service was actually delivered. Without verification, buyers price in that risk, which caps what any supplier can charge regardless of how good the service actually is. The market's real flaw is not price. It is trust.

5.2 Principle: Inference Is the Unit of Trade

What trades on this marketplace is not credits, accounts, or access. It is the inference result itself. A supplier sells the output of running an open-weight model on its own hardware, and never resells someone else's access.

This one rule secures two things. First, a supplier selling its own idle capacity cannot be cut off by a third party's policy change, so supply stays reliable. Second, because the supplier is trading its own execution, that execution itself becomes something that can actually be verified.

5.3 Business Model: Selling Verification, Not Inference

HPP does not compete on inference pricing, and it does not buy capacity wholesale to resell, so it carries no inventory risk. What HPP sells is verification, and the supplier pays for

it, not the buyer. In a market with list prices, verification is just overhead. In a discounted market, verification becomes real pricing power.

A verified, bonded supplier earns more for the same capacity, which is reason enough to opt into verification voluntarily. Buyers pay no separate verification fee.

The marketplace is built so that this trust signal actually moves price: verified supply surfaces by default, and verification results and track records are public on-chain.

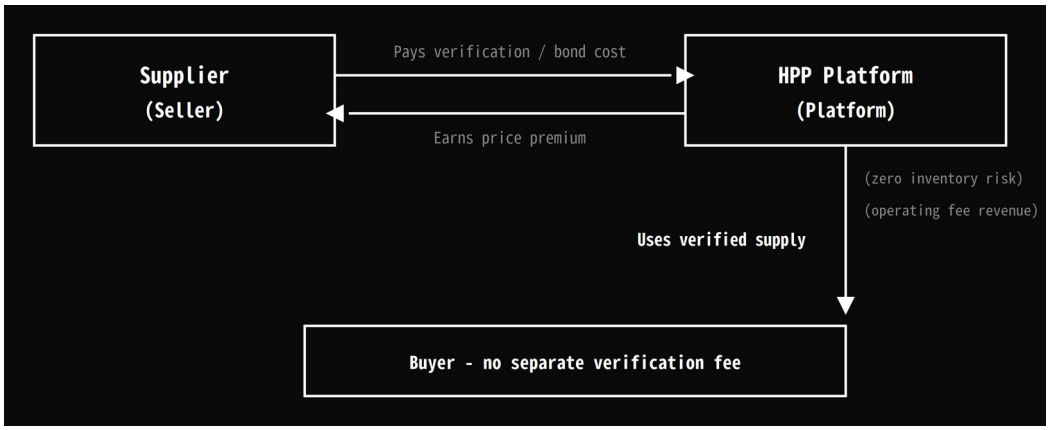


Figure 4. Seller, platform, and buyer: verification cost is absorbed into the spread

5.4 Three-Layer Trust Structure

Reputation alone cannot carry a trust system; it is backward-looking by nature and easy to game anonymously. HPP's trust structure runs three interlocking layers instead. Attestation generates the signal, bond puts an economic cost on walking away, and reputation feeds that signal into price. None of the three works without the other two.



Figure 5. Three-layer trust structure: attestation, bond, reputation

Attestation generates the signal. A supplier registers an execution profile, model artifact,

tokenizer, inference runtime, precision and quantization, decoding configuration and seed, and hardware tier, and commits a request and response hash, a model weight hash, and a token count on-chain with every delivery. Hardware tier is part of the profile because re-execution results vary by accelerator architecture; without it committed, cross-checking does not work, and misrepresenting it is itself a violation. The commitment is not cryptographic proof on its own; it is a signed claim subject to verification. What enforces it is sampling-based re-execution: for any profile the protocol has certified as reproducible, a random sample is re-run under identical conditions and cross-checked. This never sits in the request path; the response returns immediately and the check runs asynchronously.

Bond puts a price on walking away. Suppliers post slashable collateral and lose it if they break protocol rules.

Reputation puts the signal into price. Attestation results and slashing history accumulate publicly per supplier, and the market prices accordingly.

A fully trustless proof for large-scale AI inference remains an open research problem, and this paper does not claim otherwise. This architecture is designed to close most of the market's real trust gap, as practically and effectively as current technology allows.

Privacy follows the same logic. Prompts and responses are never written on-chain, only their encrypted commitments are. Raw content is disclosed only in a dispute, only to the verifying party, under a defined procedure. Content confidentiality itself is a matter of supplier contracts and data-handling policy, not something the protocol tries to enforce on-chain.

5.5 Slashing Principles

Automatic slashing applies only where the protocol can establish a violation with deterministic evidence.

- Re-execution mismatch on a reproducible profile, meaning model or output spoofing: automatic, immediate slashing.
- Availability or deadline violations the protocol can observe directly, such as signed receipts or on-chain delivery deadlines: automatic slashing, limited to what the protocol itself measures.

- Hardware tier misrepresentation: automatic where cryptographic hardware attestation is available, evidence-based adjudication otherwise.
- Subjective judgments of output quality: never auto-slashed, reflected in reputation only.

Anything without deterministic evidence goes through evidence submission and formal dispute resolution instead. This split exists to minimize false positives. In a market with near-zero switching costs, wrongly slashing an honest supplier is the most damaging failure the protocol could have, so wherever a call is not deterministic, HPP handles it through reputation and dispute resolution rather than an economic penalty.

5.6 Bond Design

Bonds are posted in HPP tokens. As the supplier pool and transaction volume grow, so does total bonded value; supply-side growth translates directly into bonding demand for HPP.

Minimum bond size is set in USD terms per model tier. Both the bond floor and the verification sampling rate are tuned so that expected losses from getting caught always outweigh expected gains from cheating, and both are re-evaluated periodically so that token price swings do not erode security. Slashed bonds are burned entirely, so the party doing the verifying is never the one who profits from the penalty. Exact bond sizes and slashing parameters will keep being tuned against live operating data.

5.7 Supply Policy: The Open-Weight Rule

Third-party suppliers can only offer models with openly published weights. This is not a preference; it is a technical requirement for verification to work at all. Verification requires independently reproducing a reference execution under the same conditions as the committed one, and that is only possible when the weights are public.

The open-weight requirement directly protects the integrity of the verification system. Because supply is limited to models that can be checked against a reference, model identity can actually be proven, and it rules out the one failure mode that matters most, a supplier without the weights reselling access to someone else's proprietary model. Legitimate ownership or rights to the compute used, and a ban on reselling third-party accounts, are separate supplier requirements enforced through attestation and re-execution checks.

At registration, suppliers make four binding commitments: the model is open-weight and

license-compliant; the compute is owned or legitimately licensed; there is no reselling of third-party accounts or access; and the hardware actually used matches what is registered. Hardware tier accuracy is a precondition for re-execution checks to work at all, so misrepresenting it is judged and slashed exactly like model spoofing.

Open-weight model quality keeps improving quickly, and the coverage a verifiable supply chain can support grows with it. HPP Coder, an open-weight model fine-tuned for a specific domain, is the first example of this policy.

5.8 Phased Rollout

Phase	Description	Verifier
Phase 1: Self-supply	HPP hosts open-weight models directly, including HPP Coder, and runs attestation and verification from day one.	Platform
Phase 2: Whitelisted suppliers	Vetted external enterprise suppliers join by hardware tier. Bonding and slashing go live in production, and verification parameters are tuned against real operating data.	Platform
Phase 3: Permissionless	Supplier participation opens fully; verification and settlement move to the Noösphere on-chain protocol.	On-chain protocol



Figure 6. Phased rollout, gated by operating metrics

Every phase transition is decided by hitting operating metrics, not by a calendar date. The supplier pool expands in step with the hardware tiers the verification infrastructure can support, since re-execution checks only work within a given tier, a new tier cannot be onboarded until the platform can verify it. Settlement, bond deposits, and attestation commitments are recorded transparently on HPP Mainnet at every phase. A blockchain is a

neutral layer that enforces rules without any single party's involvement, and that is the entire reason this marketplace needs to be built on one.

6. HPPLite and Ecosystem Applications

6.1 HPPLite: Sovereign Rollups for SQLite Applications

Most web services and applications run on a database, and SQLite is the most widely deployed database engine of all. HPPLite is an extension that turns a SQLite application into a sovereign rollup without code changes. The application uses SQLite as it always has; HPPLite records every write, posts it in encrypted batches, and anchors a state root of the whole database on the HPP network at a fixed cadence.

Each application gets its own rollup contract. Without building an independent chain or running a separate node, every application gains an ordered, verifiable data history of its own. Connecting takes one parameter on the database connection string, and HPPLite is designed to work as is with the SQLite drivers and ORMs of major languages such as Python, Node.js, and Ruby.

In this structure, each application is an agent with a chain of its own. To users it looks like an ordinary web service, but on the HPP network it is an independent entity with its own rollup contract and an ordered state history. HPP calls this an Agentic Chain. HPPLite’s design goal is to deliver this per application, with no separate chain or node.

6.2 Architecture

HPPLite applies the proven shape of an L2 rollup to a database: a single sequencer orders writes, data is posted in batches, and state is committed periodically.

Sequencer: the application process is the sequencer. It executes SQL, groups writes into batches, and posts them. An on-chain lease guarantees that only one instance posts at a time.

Write capture: every write is recorded as row-level changes inside the application’s own transaction. When a transaction rolls back, its record rolls back with it, so no write is ever applied without being logged.

State commitment: every set number of batches, a state root built from per-table Merkle trees is anchored on the HPP network. Only what changed is recomputed, so commitment

cost tracks the volume of change, not the size of the database.

Rebuild and verify: a key holder can rebuild the entire database from chain records alone, starting from an empty file, and compare it against the posted state roots.

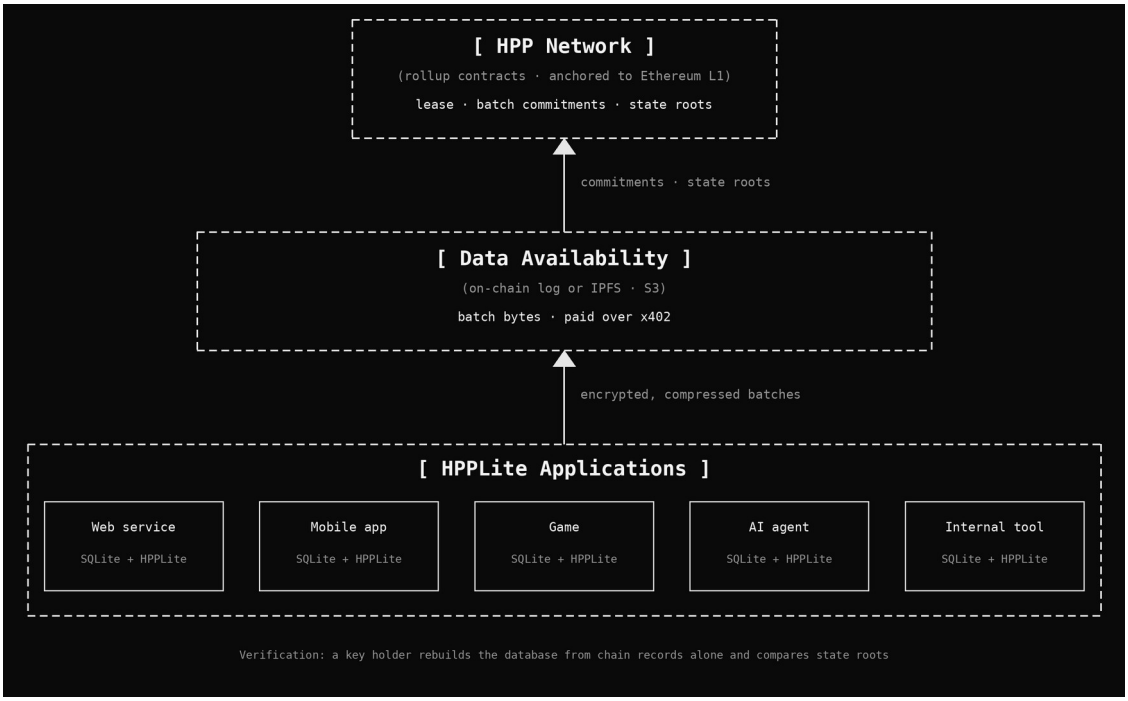


Figure 7. HPPLite architecture: application → data availability → HPP network

6.3 Data Availability

HPPLite carries its own data availability (DA) layer and lets each application choose where its batches live.

- On-chain DA: batches are written to transaction logs on the HPP network, and the contract stores only the batch height and a rolling hash, at far lower cost than contract storage.
- External store with on-chain commitments: batches live in an external store such as IPFS, and the HPP network keeps only the hash, the reference, and the store's signed receipt. Store fees are paid per call over x402, and if the store does not respond, batches are posted to the chain directly.

Either way, batch hashes are checked against on-chain records during a rebuild, so integrity is verified the same way wherever the bytes live. Batches are compressed before posting to cut storage cost, and encrypted with ChaCha20-Poly1305 so that the public chain holds only ciphertext.

6.4 Trust Model

HPPLite is a sovereign rollup. Rather than a settlement layer adjudicating state transitions, the chain acts as a witness that holds ordered data and state commitments, and verification is performed by whoever rebuilds the database. HPPLite does not enforce operator behavior through economic penalties. Instead, tampering surfaces as a state root mismatch during a rebuild, and a lost server disk can be recovered from the chain.

Today, verification is performed by holders of the encryption key. Viewing keys that allow rebuilding and auditing without write authority, and separate roles for operator and owner keys, are in development as the next step.

6.5 Relationship with the HPP Network

An HPPLite application deploys its own rollup contract on the HPP network and performs batch posting and state commitments as HPP network transactions. When an external store is used, store payments also run over x402 on the HPP network. Gas is paid in ETH, and low fees keep operating costs under control even for write-heavy applications.

As more applications run on HPPLite, real transactions and payments on the HPP network grow with them. HPPLite creates demand for the HPP network directly at the application layer.

6.6 Application Strategy

The HPP stack reaches users through the applications built on it. HPP grows that application layer in two ways: building applications in-house and supporting builders with grants and infrastructure.

Build: applications designed and operated by the HPP team. Rather than one heavy service, HPP steadily builds many lightweight applications, and refines the path for putting applications built by coding agents onto HPPLite. Because each application has its own rollup, every new application adds to real usage records on the HPP network.

Support: grants, HPP Router credits, and infrastructure support lower development costs for early-stage teams.

Connecting applications made in HPP Hub so that they run on HPPLite is another option that may be explored.

These applications record data through HPPLite, pay through x402, and run inference

through HPP Router, all on the HPP network. More applications mean more network usage, and the direction for linking that revenue to the HPP token is set out in Section 8.3.

6.7 ArenAI

ArenAI is the reference application that consumes the full HPP stack, a trading arena built for AI agents. Agents compete on returns in a simulated environment, and users can test and validate strategies risk-free, with no real capital at stake. A validated strategy graduates to a live trading tournament on an on-chain DEX; the roadmap extends this into a marketplace where validated strategies themselves become tradable IP. Users can also connect their own custom agents to participate.

Agents inside ArenAI pay through x402 and infer through HPP Router, calling Noösphere for verifiable on-chain compute when needed. ArenAI is not a standalone product; it is proof that HPP's agent-economy stack works under real conditions.

Live-money tournaments and the strategy marketplace roll out in stages, within whatever regulatory and investor-protection requirements apply in each jurisdiction. Exact features and launch timing will be announced as development milestones are reached.

7. Extension Tracks

7.1 HPP Enterprise

HPP Enterprise bridges and settles a private, Hyperledger BESU-based L3 to HPP Mainnet. Institutions get the data control of a private environment and the auditability of a public chain at the same time, and their enterprise workloads flow into the public network as real demand.

This track rests on Blocko's decade-plus of running mission-critical, national-scale blockchain infrastructure in production, and the customer base that comes with it. A direct pipeline from proven B2B enterprise demand into public-chain settlement is a moat most public chains do not have, and HPP already has it.

7.2 zk Settlement

In line with the upstream Arbitrum roadmap, this track adds a zero-knowledge proof layer to the existing BoLD-based optimistic settlement. A valid proof shortens finality and bridge withdrawal times; if none is submitted in time, the existing BoLD dispute mechanism still

works as a fallback. This is additive, not a replacement; rollout stays within what the upstream stack officially supports, preserving HPP's no-fork principle.

8. HPP Token

8.1 Overview

Info	Detail
Token name	HPP
Standard	ERC-20
Total supply	1,700,000,000 HPP (fixed, non-inflationary)
Origin network	Ethereum Mainnet (L1)
L2 representation	Bridged token via Arbitrum Canonical Bridge
L1 contract	0xe33fbe7584eb79e2673abe576b7ac8c0de62565c
L2 contract (HPP Mainnet)	0xB48334E7938367bC24Fe1F19000D6f06C622E6c7
Role	Utility token, governance token

HPP is minted on Ethereum Mainnet and bridged to HPP Mainnet through the Arbitrum Canonical Bridge. Total supply is controlled entirely by the L1 contract; L2 tokens are minted and burned only through the standard bridge, with no discretionary minting authority anywhere. ETH remains the gas token on HPP Mainnet without exception. HPP itself is not used to pay network fees; its role is the utility and governance functions below.

8.2 Token Utility

Staking, Stake-to-Credit: staking HPP unlocks tiered credits toward HPP Router usage or fee discounts, the core mechanism tying token holding directly to real usage of ecosystem services. Staking already runs as a recurring seasonal program, with terms and rewards for each season announced through official channels.

Supplier bonding: suppliers, agent operators, and future strategy sellers on the verifiable inference marketplace post HPP-denominated bonds to guarantee service integrity. Bonds carry a USD-denominated floor by model tier and are subject to burn-based slashing on violation, so marketplace supply-side growth converts directly into HPP bonding demand.

Governance: HPP holders vote with staking-weighted power on HPP DAO decisions.

Ecosystem expansion: token utility expands as ecosystem applications grow, starting with ArenAI. The direction for linking ecosystem revenue to the token is set out in Section 8.3.

8.3 Value Accrual Direction

HPP aims to ground token value in revenue the ecosystem actually earns, rather than in new issuance or incentive distribution. Ecosystem revenue may include HPP Router usage fees, platform revenue from the verifiable inference marketplace, and revenue from applications HPP builds directly, including on HPPLite.

Whether and how ecosystem revenue is linked to the HPP token will be decided through HPP DAO proposals and votes. This paper does not commit to any specific mechanism or its scale.

9. Governance: HPP DAO

HPP DAO is HPP's on-chain governance, live since April 2026. Built independently rather than on a third-party governance platform, it runs as three connected channels: proposals are raised and debated in the Forum, formalized proposals go to a vote on Agora, the voting layer, and outcomes and execution status are recorded transparently in HPP Report.

Every proposal follows the same five-stage lifecycle: discussion, formalization, vote, execution, and integration, with status tracked on-chain (Status Snapshot) end to end for full transparency. Voting power is staking-based and fixed at a snapshot when voting opens. Governance participation feeds back into staking rewards, so taking part in a decision counts as a real contribution to the ecosystem.



Figure 8. The proposal lifecycle and its owning channel

How ecosystem revenue is linked to the HPP token (Section 8.3) will also be decided through HPP DAO proposals.

10. Roadmap

The dates below reflect HPP's current targets. Marketplace phase transitions are decided

by hitting operating metrics, not fixed dates. Detailed schedules will be announced through official channels as it is confirmed.

Timing (current target)	Milestone
Q4 2026	1) HPP Router service officially launches 2) Full support for x402 v2 and Account Abstraction 3) ArenAI alpha launches
2027	1) Verifiable inference marketplace Phase 1 (self-supply, attestation live) reaches steady state, and Phase 2 (whitelisted suppliers) expands 2) Marketplace opens permissionless (Noösphere on-chain transition) 3) HPPLite development and HPPLite-based applications 4) ArenAI full launch 5) HPP Enterprise (BESU-based L3) track begins in earnest 6) zk Settlement support

11. Conclusion

House Party Protocol is building the execution layer for a future where agents are primary economic actors. It closes the three foundational gaps in the agent economy, payment, trust, and execution, with one integrated stack, and the trust behind every inference and off-chain compute transaction rests on a verification structure built from attestation, bonds, and reputation working together.

HPP does not claim a fully trustless proof that does not yet exist. It will not compete on price, will not carry inventory it does not need, and will not accept supply it cannot verify. Instead, it verifies what can be verified, prices in the cost of walking away, and builds a market where trust is honestly reflected in price. Access opens in stages. Decentralization happens gradually, but the direction does not change.

HPP aims to ground the token in the ecosystem’s real economic activity.

A team with more than a decade running mission-critical, national-scale systems is solving an old problem, trustworthy trading of idle resources, with the resources of the AI era. Together with developers, suppliers, enterprises, and its community, HPP is working to set the standard for an autonomous agent economy built on verifiable intelligence.

12. Disclaimer and Investment Risk Notice

Important: please read the full disclaimer before relying on this whitepaper.

Disclaimer

This House Party Protocol (“HPP”) whitepaper (the “Whitepaper”) is published for

informational purposes only and creates no legal obligation. It is provided as is, as of the date of writing. Any part of it may be changed, updated, or replaced at the HPP team's discretion without notice, and the most recent version always supersedes earlier ones.

1) Nothing in this Whitepaper constitutes an offer, solicitation, or inducement to invest, and none of it should be read as business, legal, financial, or tax advice. If anything here is unclear, consult an accountant, lawyer, or other qualified professional before making a decision.

2) Some statements in this Whitepaper are forward-looking rather than statements of current fact. Anything relating to future plans, including product features, the roadmap, and phased plans, may change for any reason, and no guarantee or commitment is made as to whether it will be realized.

3) The HPP team and related parties accept no liability for any direct or indirect loss arising from information in this document.

4) Distribution of this Whitepaper may be restricted in certain jurisdictions. Complying with applicable local laws and regulations is entirely the recipient's responsibility.

5) Information may be lost or altered in translation. The English version is the official language version of this Whitepaper; where versions conflict, the English version governs.

Investment Risk Notice

1) Purchasing HPP tokens may result in material loss, including loss of the purchase amount. The HPP team accepts no liability for a decline in token value, loss, or loss of liquidity arising from force majeure factors such as regulatory change.

- Blockchain risk: weaknesses or vulnerabilities in the underlying protocols, Ethereum and the Arbitrum Nitro stack.
- Personal data risk: potential loss of tokens from a wallet due to private key or personal data compromise.
- Security risk: attack risk against the blockchain and smart contracts.
- Wallet compatibility risk.
- Force majeure risk: the project may be modified due to technical, legal, or administrative regulatory change.

- Token distribution and market impact risk, and general market and liquidity risk.
- Network operation risk: risk arising from the AnyTrust data availability committee, sequencer and batch-poster operators, contract upgrade authority, the canonical bridge, and dependence on third-party infrastructure such as RaaS.
- Protocol and service risk: risk from mis-adjudication in the verification process, supplier misconduct, changes to staking, bonding, or slashing policy, and the token migration process.
- AI regulatory risk: the introduction or change of national regulation applicable to the provision and use of AI models.
- Value accrual risk: any mechanism linking ecosystem revenue to the token may not be adopted, or may change, depending on HPP DAO decisions, and ecosystem revenue does not guarantee token price.

Review this risk notice and the applicable terms of use in full before using HPP tokens or related services.

Appendix A: Glossary

Term	Definition
Layer 2 (L2)	A scaling network that processes transactions on top of Ethereum and settles results back to it, inheriting Ethereum's security while improving cost and throughput.
Layer 3 (L3)	An application-specific network built on an L2 that settles to it.
Optimistic Rollup	A scaling design that assumes transactions are valid by default and verifies them with a proof only if challenged.
AnyTrust	A data-availability model in which a designated committee guarantees availability, at lower cost than posting full data to Ethereum.
BoLD	Arbitrum's dispute resolution protocol, designed so a single honest participant is enough to defend the correct state.
zk Settlement	A settlement approach that uses zero-knowledge proofs to back the validity of a state transition and shorten finality.
Canonical Bridge	A network's officially supported path for moving assets, locking on L1 and minting on L2.
x402	An open payment standard built on the HTTP 402 (Payment Required) status code, enabling per-call payment with a signature and no account creation.

Term	Definition
Account Abstraction	Technology that turns a wallet into a programmable account, enabling features such as fee sponsorship, automated execution, and session-based permissions.
Proof-of-Inference	HPP's verification structure described in Chapter 5, which checks the origin and integrity of inference results through attestation, sampled re-execution, and bonds. It is not a complete cryptographic proof but a combination of verifiable claims and economic collateral.
Commitment	A record, typically a hash, that fixes content without revealing it, allowing later comparison against the original to detect tampering.
Bond	Collateral a supplier posts to guarantee performance, subject to forfeiture if a rule is violated.
Slashing	Forfeiting a posted bond as a penalty for a rule violation. On HPP, slashed amounts are burned.
Open-Weight Model	An AI model whose weights are published, so anyone can download it, run it, and check the results.
Stake-to-Credit	A mechanism that grants usage credits or fee discounts to participants who stake HPP.
Sovereign Rollup	A rollup that posts data and state commitments to a chain, while verification is performed by the parties who rebuild the data.
Agentic Chain	A structure in which a single application has its own rollup contract and an ordered state history. On HPP, it is implemented with HPPLite, without a separate chain or node.
Sequencer	The party that orders transactions or writes and posts them in batches.
State Root	A single hash summarizing the entire state at a point in time, compared against a rebuilt state to confirm integrity.