

0G Q2 2025 Overview



Throughout 2025, 0G Labs has successfully accelerated its technical roadmap progress and ecosystem growth. There have been several highlights thus far for 0G, including the centerpiece so far in Q2 with the launch of Galileo, 0G's third public testnet (and the most capable version of its blockchain so far). The milestone of Galileo followed the success of the Newton testnet, which saw substantial user activity and active stress tests, providing valuable insights that allowed Galileo to be launched confidently.

Alongside core infrastructure upgrades, 0G expanded its ecosystem through integrations with AI-native applications like sightAI, xNomad, and Cove, showcasing the network's commitment to a high-performance, AI-optimized blockchain environment. Together, these initiatives do a great job in showcasing 0G's overall mission in building a modular, scalable infrastructure layer for decentralized compute, storage, and intelligent agent applications.

Testnet Progress: Transition from Newton to Galileo

Newton, 0G's second major public testnet (V2), achieved significant early adoption and provided valuable lessons that informed Galileo's design. In just a short span, Newton recorded:

- Over 2.5 million unique wallets
- Executed more than 350 million transactions
- Saw the deployment of roughly 530k smart contracts

Above all, Newton demonstrated that there is significant market interest and perceivable demand for 0G's high-throughput, AI-centric infrastructure. Newton was built using the Cosmos SDK and Ethermint, enabling rapid integration across 0G's compute and storage layers. However, this prototype architecture revealed structural bottlenecks that constrained future scalability and needed to be addressed before mainnet.

Of course, as is the point with any testnet, as usage significantly increased on Newton, some clear technical limitations began to creep up. Some examples of core limitations include:

- **Mempool Congestion:**

The Tendermint-based mempool did not prioritize high-gas transactions, leaving the network vulnerable to spam and transaction flooding.

- **Monolithic Consensus & Execution:**

Newton's consensus and execution were fused in a single process, imposing a performance ceiling and making it difficult to optimize or upgrade components independently.

- **Integration Friction:**

Fast-paced ecosystem development encountered forward-compatibility issues, resulting in frequent upgrades that often required complex rework for dApp teams to remain compatible.

The above pressure points provided invaluable data at scale and highlighted the need for a full architectural redesign. In response, the 0G team used this feedback to engineer a more robust solution in Testnet V3 (Galileo), raising the bar for performance and modularity.

Galileo Testnet Launch: Architecture & Performance Upgrades


Galileo (Testnet V3) was designed and created amidst the architectural redesign necessitated after concluding Newton. Galileo officially went live in Q2 2025 as a clean-slate deployment. It introduces major technical improvements to address Newton's limitations and advance the network's capabilities. Key advancements with Galileo include:

1. Higher Throughput

Galileo is expected to achieve approximately a 70% increase in throughput over Newton, supporting up to 2,500 transactions per second (TPS) under optimized conditions. Since Galileo launched, it has been far off these numbers, though it does feature a steadily increasing TPS despite increasing user activity. At the time of writing, Galileo's TPS is around 60.

Transaction TPS

May 10, 2025 → Jun 20, 2025



Finality also remains fast (~1–2 seconds), and the network can sustain significantly more load than before. This alone is a major improvement, helping to prepare 0G for a full mainnet launch, even with the expectations of immediate, extremely high demand

2. Decoupled Consensus and Execution

The new design for Galileo replaces the Cosmos SDK core with an Ethereum Engine API architecture, effectively separating the consensus layer (now powered by an optimized CometBFT) from the execution layer. Overall, the design is a clear form of modularization and removes previous performance ceilings. The move also paves the way for future enhancements like parallel transaction execution.

Additionally, Galileo integrates the latest Ethereum protocol upgrades, including the Shanghai and Cancun-Deneb hard forks, and is built to support Pectra features. Developers can utilize new EVM opcodes like PUSH0 and benefit from improved data handling and contract efficiency on 0G. In short, the testnet is fully EVM-compatible, ensuring Ethereum tooling and smart contracts work seamlessly.

3. Flexible Validator Staking

Validator delegation has moved out of the base protocol and into smart contracts. By implementing staking logic at the contract level, Galileo allows more flexible and upgradable validator and delegation mechanisms without requiring hard-fork changes. The design choice here enables faster iteration on validator incentives and governance.

4. Security & Reliability Improvements

A fully revamped faucet system was introduced to distribute testnet tokens with better bot protection and abuse prevention. Dedicated RPC nodes and other infrastructure enhancements have eliminated previous bottlenecks, resulting in a more stable network under heavy load. These changes collectively reinforce the network's security and should ensure more consistent performance.

Despite the deep changes under the hood, Galileo was designed so that existing 0G Storage and 0G Compute users required no changes to their code or workflows during the upgrade. The backend services were updated for V3 compatibility, meaning developers could continue using 0G's decentralized storage and compute as before, with Galileo's improvements automatically applied.

As of early May 2025, Galileo is live and operating with these new capabilities, effectively resetting the chain (new chain ID 16601) and inviting developers to migrate to the new testnet. Success for Galileo has also come quite early, with the testnet already racking up over 8.6 million unique wallet interactions across 6.9 million deployed smart contracts and 128 million total transactions (in roughly 1 month).

Transaction


May 10, 2025 → Jun 20, 2025



0G Infrastructure Upgrades

In addition to the testnet advancements described above, the 0G development team has been diligently working to deliver key infrastructure roadmap objectives. Overall, 0G’s Q2 technical developments centered on modularizing the core of the chain and enhancing its AI-specific capabilities.

0G team Github activity in 2025



One of 0G’s main accomplishments on this front is its redesigned chain architecture that separates the consensus engine from the transaction execution process. In Galileo, the team replaced the previous Cosmos SDK monolithic framework with the standard Ethereum Engine API for consensus client communication. Consensus (based on CometBFT/Tendermint) runs as an independent process, while a parallel execution client processes EVM transactions.

The decoupling here yields multiple benefits:

- Consensus upgrades or replacements can occur without rewriting the execution layer.
- Execution environments can be optimized or run in parallel threads.
- Overall maintenance of 0G is simplified.

Of course, modularity is the core theme. By treating components like interchangeable modules, 0G can iteratively improve each module in isolation and explore potential future parallel execution opportunities. The net result of 0G’s efforts is what should lead to greater scalability for future growth.

In addition to the changes at the core protocol level, several other optimizations and upgrades have occurred:

1. Optimized Consensus with BLS Signatures

In tandem with decoupling, the consensus layer saw its own optimizations. Galileo’s CometBFT consensus uses BLS signature aggregation to reduce overhead in block confirmations. By aggregating validator signatures efficiently, 0G minimizes the communication and verification load per block, contributing to higher throughput. The consensus tuning is particularly important for high-frequency AI transactions and data availability. Finality remains on the order of seconds, and BLS helps keep latency low even as validator count scales.

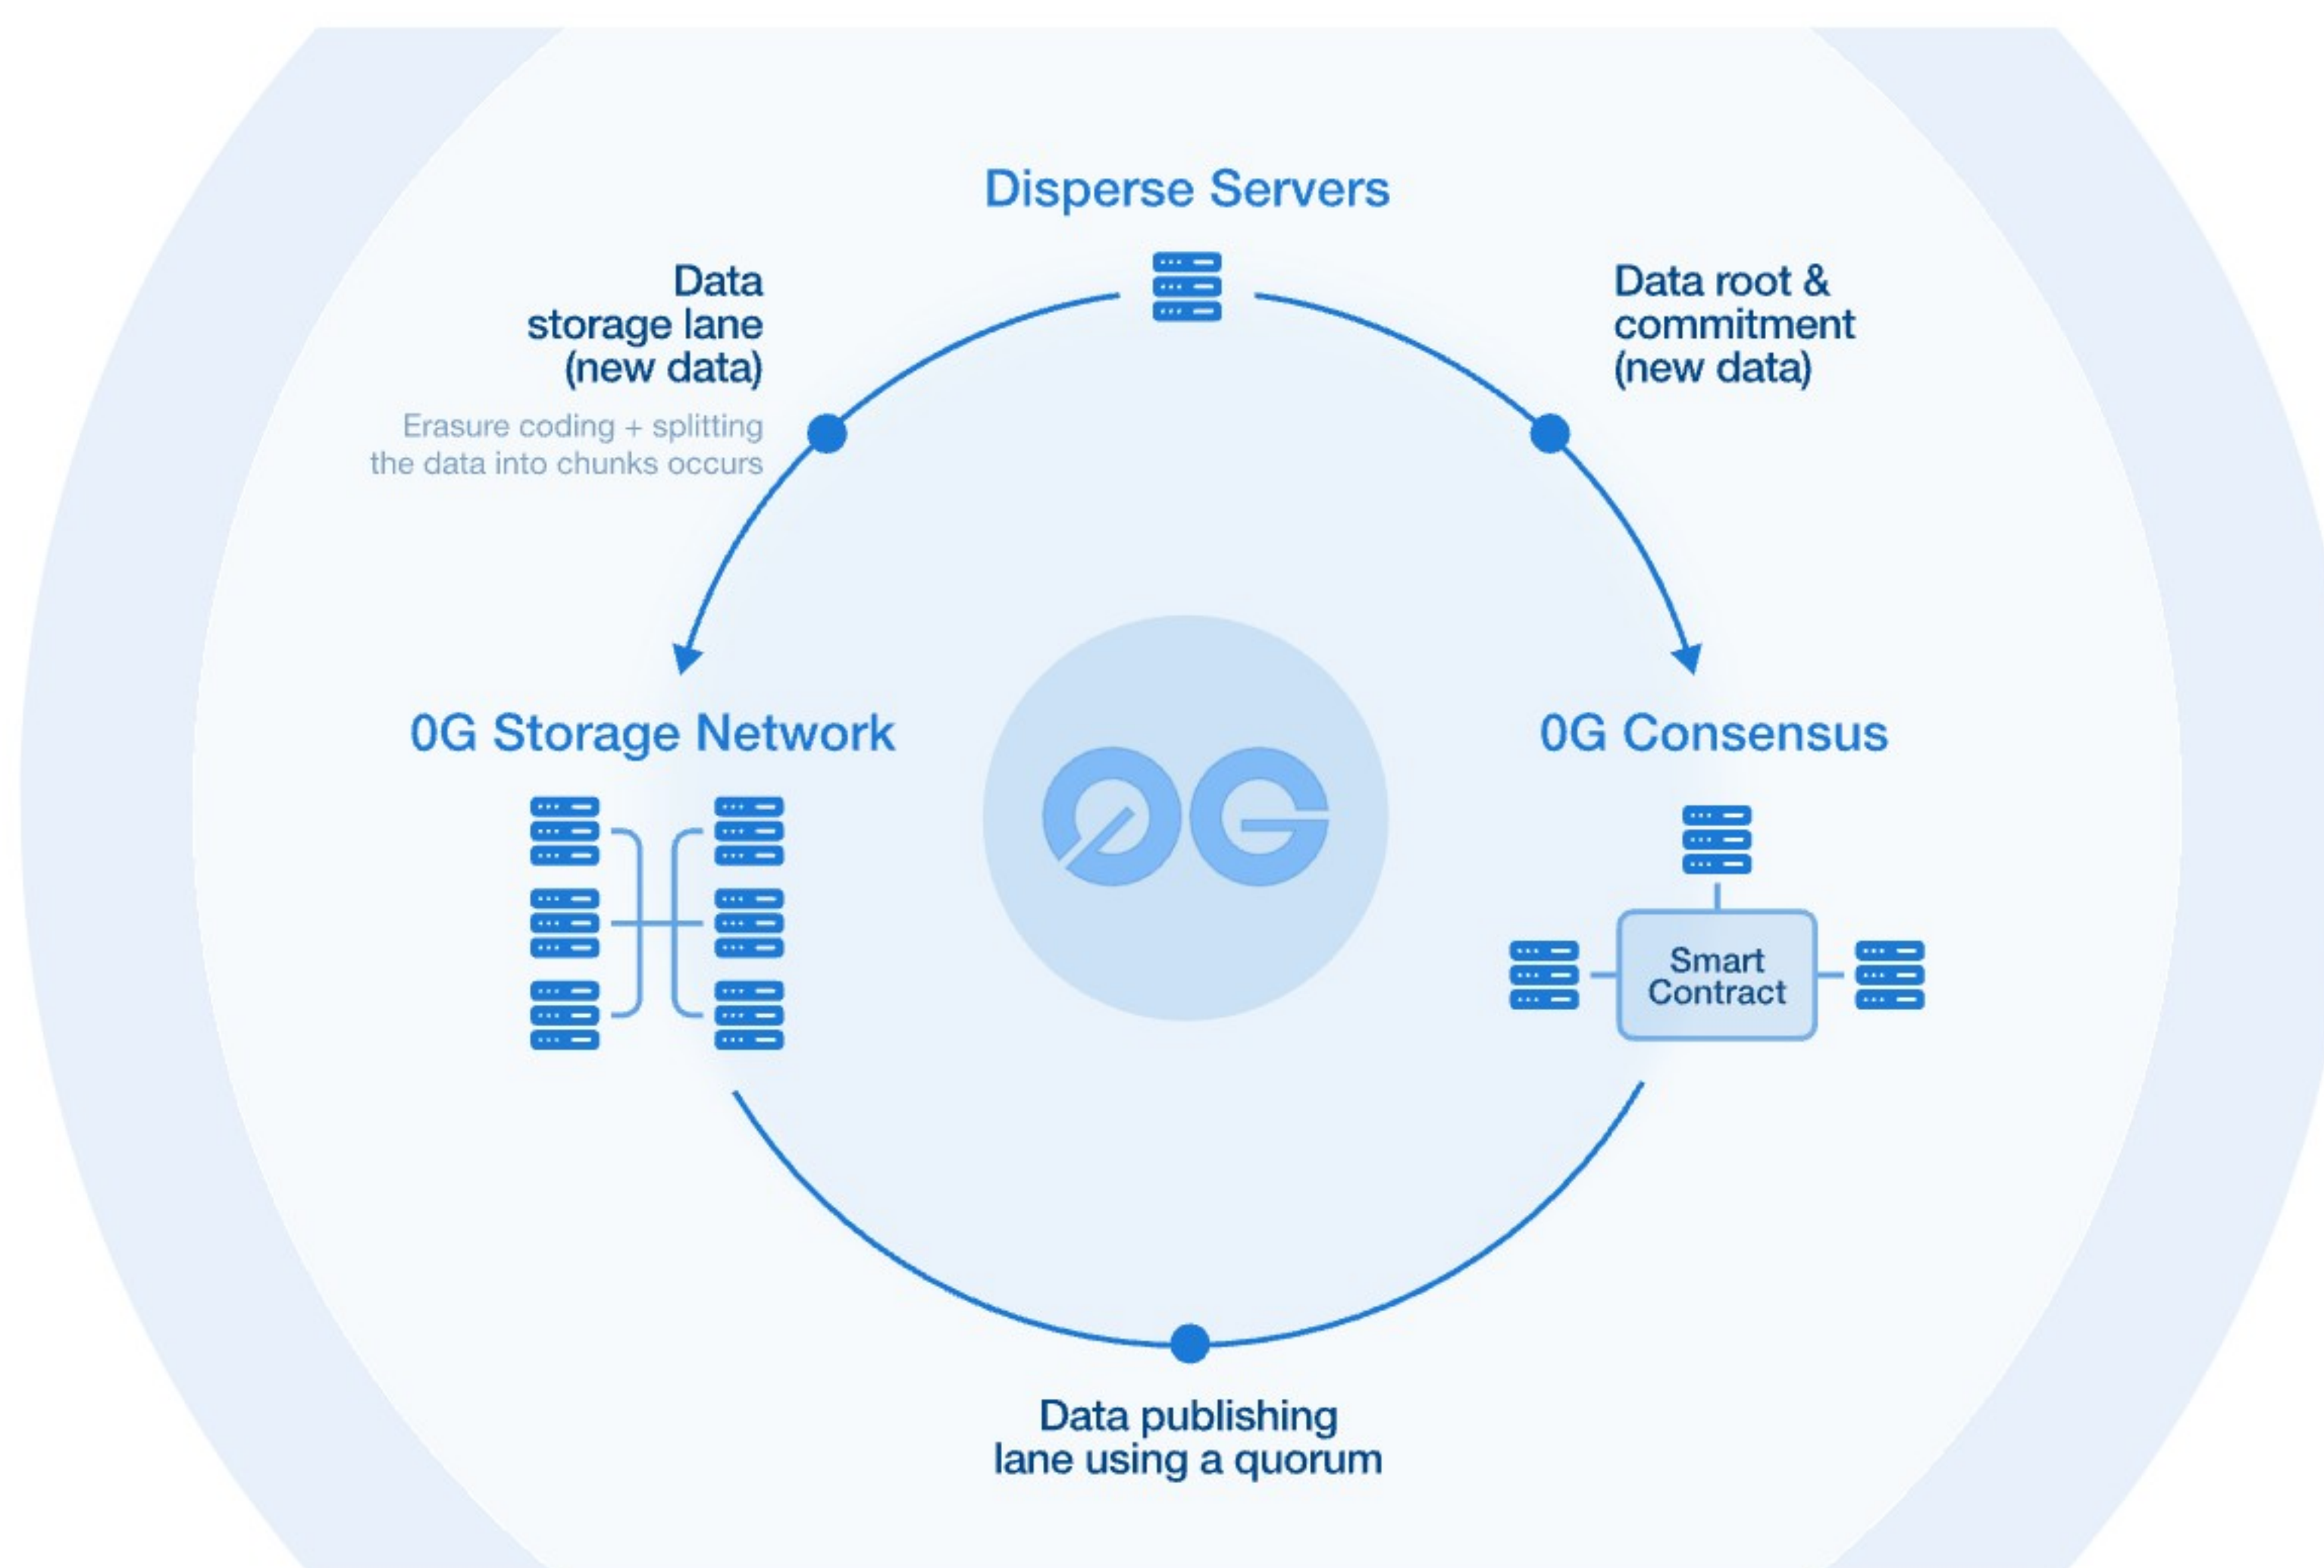
2. Validator Flexibility via Smart Contracts

Another technical innovation in Q2 was the introduction of smart contract-based validator delegation. Instead of hard-coding the staking and delegation logic in the protocol (as many chains do), 0G moved this functionality into upgradable smart contracts on the chain. The approach grants far more flexibility in how the validator set and staking economy can evolve.

For example, new staking mechanisms, reward schemes, or even liquid staking derivatives can be implemented at the contract level without requiring a hard fork or core update. It also accelerates iteration on improvements to validator incentives and security policies. In Galileo, the default contract handles standard delegations, but developers or the community could propose new contracts to introduce features in the future.

3. Storage and Compute Integration

Beyond the consensus/chain layer, 0G continued to refine its unique AI storage and compute infrastructure. The 0G Storage network, which underpins the chain's data availability, achieved throughput of 2 GB/s in recent benchmarks, per 0G Labs. At that throughput, 0G would hold the title for fastest decentralized storage technology in the AI-crypto sector.



On the compute side, 0G's distributed inference layer operates with no appreciable overhead versus centralized clouds, and at significantly lower cost, thanks to optimizations in how model execution is verified on-chain. AI applications often involve massive data streams or heavy compute workloads (like model training or real-time inference), which traditional blockchains cannot support. 0G's vertically integrated approach (Layer-1 + off-chain layers) is proving capable of meeting these demands. Overall, the approach is a direct lineage to 0G's goals of building out a full-stack AI operating system.

Ecosystem Expansion

The \$89M Ecosystem Growth Program

To support its expanding ecosystem, 0G launched a major funding initiative in early 2025: the \$89 million Ecosystem Growth Program. While announced in February, this program has been in full swing during Q2, providing critical resources and structure to accelerate AI application development on 0G.

Backed by the 0G Foundation and leading strategic investors, the \$89M Ecosystem Program is one of the largest Web3 initiatives dedicated to decentralized AI. The vision is to nurture the “next generation of AI” built as open, modular, trustless applications rather than closed APIs. In practice, this means financially empowering developers, startups, and researchers who are working at the intersection of AI and blockchain. By injecting capital and support into many projects, 0G is helping seed a diverse ecosystem where autonomous agents, AI-driven dApps, data marketplaces, and more can flourish on 0G’s network.

The Ecosystem Growth Program is structured to be flexible in how it supports projects. It offers grants, liquidity support, and other forms of investment depending on the project’s needs and stage. For very early-stage ideas, small grants or prizes might be given (for example, hackathon prizes and Guild grants). For more mature startups, the program can participate in funding rounds or provide liquidity (i.e., for protocol-owned treasuries or token incentive programs) to help them scale. The common denominator is a focus on projects that are purpose-driven, scalable, and aligned with 0G’s ecosystem priorities.

Global Accelerator Partnership

To further catalyze ecosystem growth, 0G Labs announced a partnership with Web3Labs to launch a Global Accelerator Program in June 2025. This is a six-month accelerator jointly run by 0G and Web3Labs,

designed to support high-potential projects within the 0G ecosystem. The program will select up to 10 teams globally (with a focus on Asia-Pacific) and provide them with capital and resources to scale from early traction to mainstream adoption.

This Acceleration Program Timeline

Program Timeline	Strategy
May 26 - July 1	Project Recruitment
July 1 - July 31	Project Selection & Admission
August - October	Acceleration & Resource Support
October - November	Program Conclusion & Follow-up

Participating startups receive up to \$25,000 in grants, hands-on technical integration support, marketing/PR assistance, and connections to an investor alliance for fundraising. They also benefit from perks like up to \$10k in cloud credits and fast-track exchange listing support. The accelerator’s thematic scope spans several domains aligned with 0G’s vision, including:

1. AI Agents
2. AI-NFTs
3. Decentralized Science (DeSci)
4. PayFi
5. Robotics
6. Privacy-Preserving AI

By Q2’s end, the accelerator had kicked off with community meetups (including a 150+ builder gathering in Hangzhou) to spur interest. It complements the earlier “Guild on 0G” program (focused on very early-stage teams) by targeting projects that already have some traction and need help to reach the next level. Together, these programs are fostering a pipeline of quality projects that will enrich 0G’s ecosystem and drive usage of its network.

New AI-Native Applications on 0G

In addition to nurturing native projects, 0G Labs has also forged partnerships with established players to strengthen its infrastructure. As part of the ecosystem fund, 0G launched an accelerator program called “Guild on 0G” with a dedicated \$8.88M allocation. The Guild focuses on early-stage builders who are ready to go from testnet prototypes to mainnet launch. It provides a more structured incubation, including direct funding, technical mentorship, and even gas credits to offset the costs of running AI workloads on 0G. The intent is to fast-track these teams so that by the time 0G mainnet goes live, they have polished dApps ready for users.

The Guild program has already onboarded multiple cohorts of AI-first startups. For example, three notable projects have been beneficiaries of Guild support and funding:

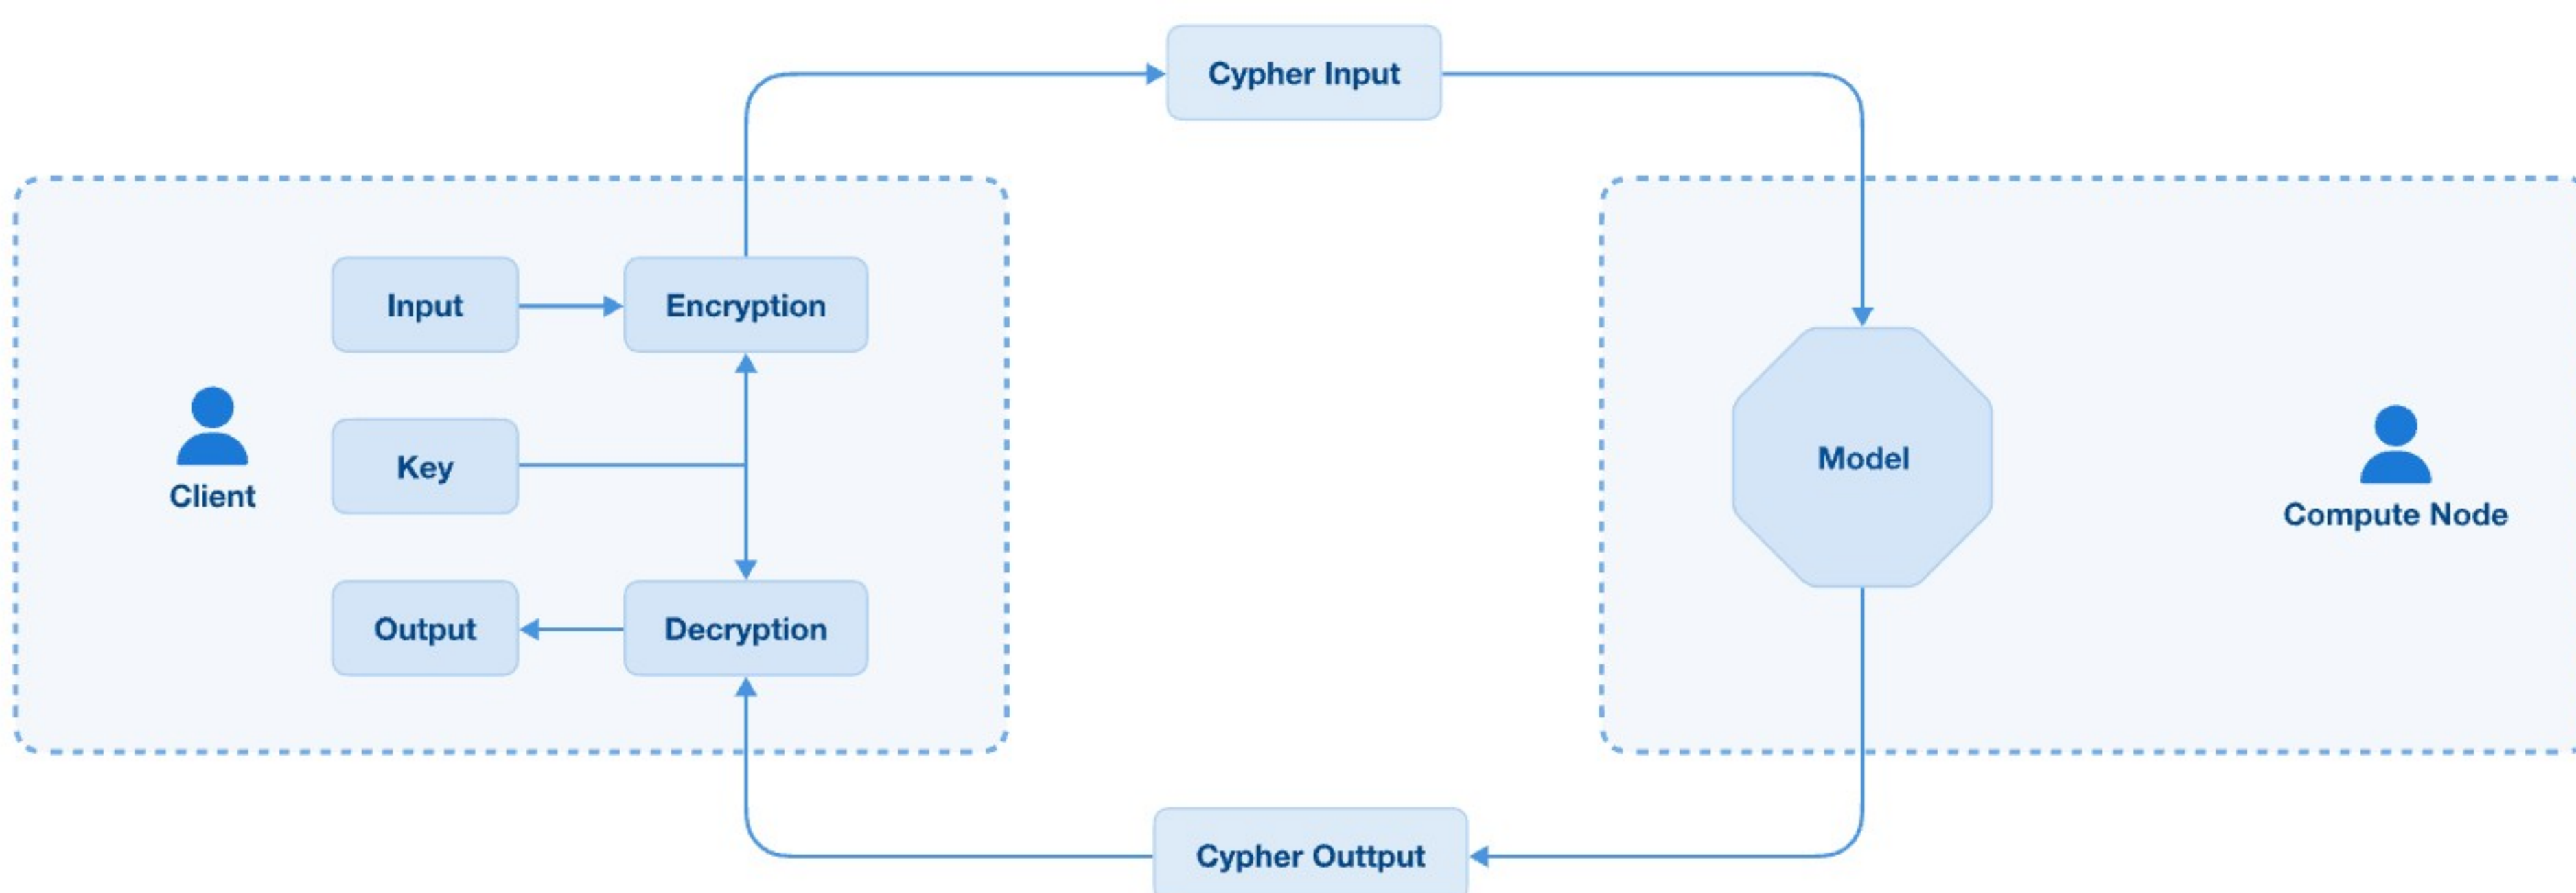
- sightAI
- xNomad
- Cove

Between the three projects, they have actively demonstrated how 0G’s infrastructure supports composable, intelligent agents within Web3. Each project leverages 0G’s compute, storage, and identity features to push the boundaries of its own business models.

1. sightAI – Composable Agents with Memory

sightAI is a privacy-preserving computation oracle platform that enables secure, privacy-focused calculations on encrypted data across multiple chains (EVM chains, Solana, and TON). It does so by utilizing advanced cryptographic techniques known as Fully Homomorphic Encryption (FHE), which enables users and dApps to process sensitive data without revealing the underlying information. Overall, this makes sightAI extremely enticing for AI development.

FHE machine learning diagram



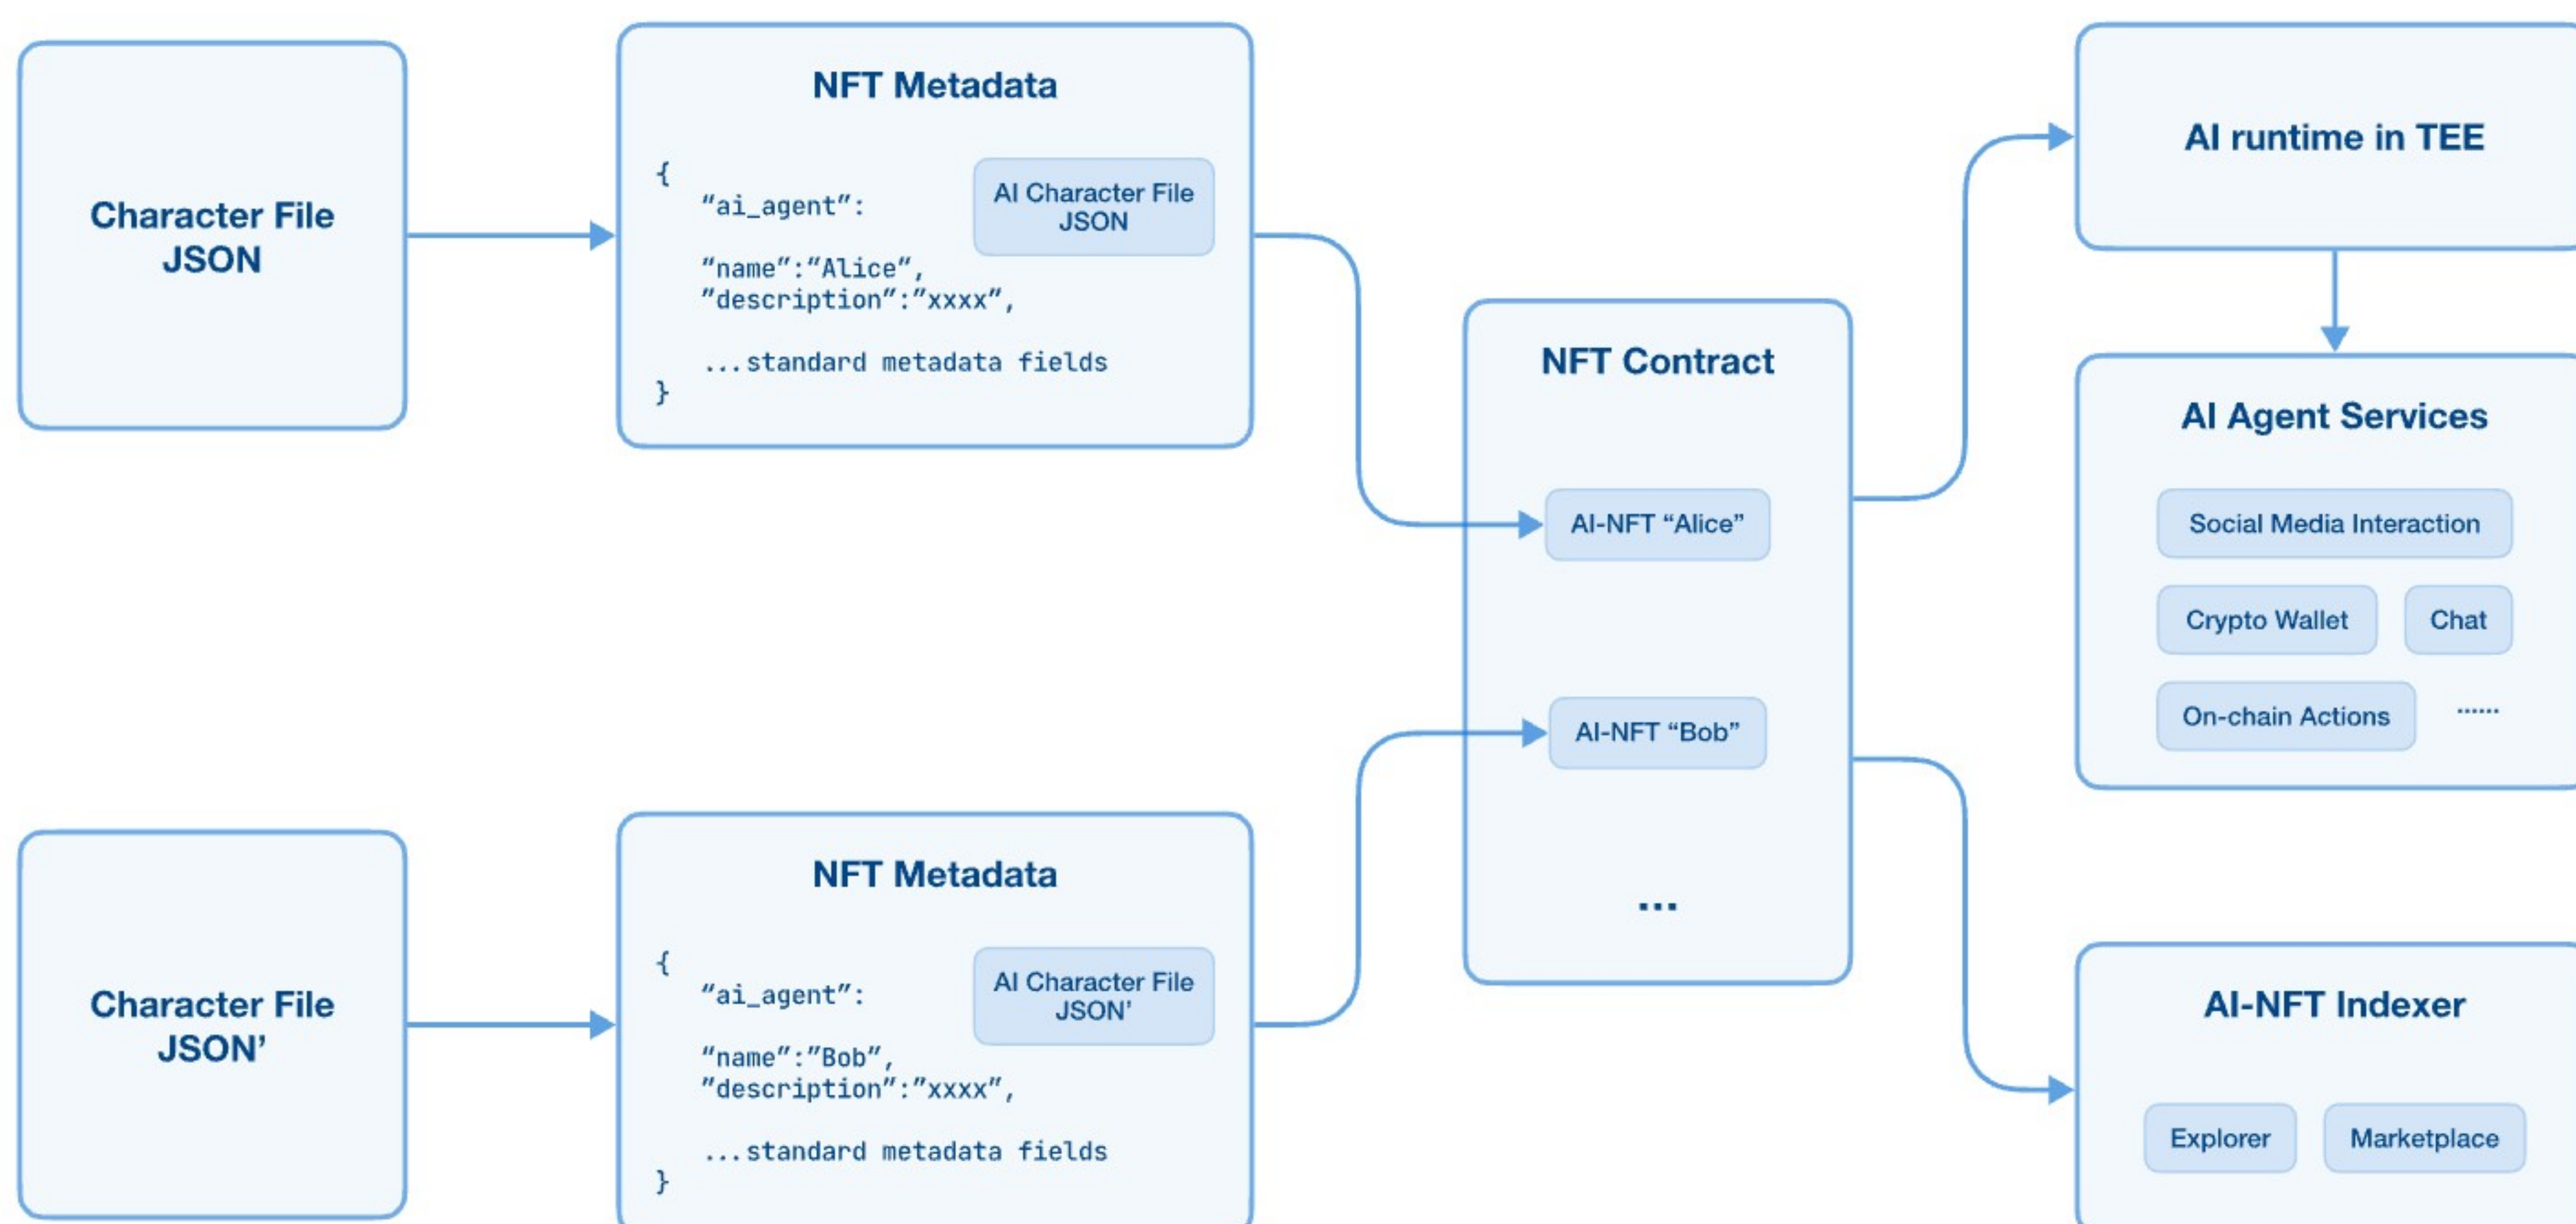
Per 0G Labs, sightAI is building an open framework for modular AI agents that feature transparent inference and persistent memory across sessions. Traditional AI systems lock users into opaque, centralized APIs, but sightAI’s approach uses Modular Compute Providers (MCPs) as decentralized agent hosts, each identified by a DID (Decentralized Identifier) and deployable on demand. Users can compose and query different AI agents, inspecting the full chain of logic and data behind each result (transparent inference).

sightAI also introduces encrypted memory, allowing users to securely carry context from one interaction to the next. So, how does sightAI benefit from building on top 0G? 0G supports sightAI by providing decentralized storage for model data and encrypted user memory, enabling agents to have a verifiable, persistent state across devices.

2. xNomad – Tokenized AI Agents (iNFTs)

xNomad is focused on making autonomous AI agents accessible through tokenized ownership. It allows anyone to turn an NFT into an “autonomous, Intelligent NFT” (AI-NFT), which effectively gives NFTs a “brain” and bank account. These AI-NFT agents can manage crypto assets, interact with DeFi protocols, and evolve over time as personalized assistants.

How AI-NFT Works



Technologically, each agent is powered by the xNomad Brain, a domain-specific large language model fine-tuned to interpret on-chain data and crypto market signals. By living on-chain, the agents carry state and logic that is transferable and transparent. 0G’s infrastructure plays a key role by offering low-cost inference and storage to support xNomad’s AI workflows. For instance, 0G’s GPU-accelerated compute ensures that running AI models or analytics for these agents is cost-effective and scalable.

Moreover, xNomad uses 0G’s custom token standard (ERC-7857) to mint verifiable, composable AI-NFTs on the chain, ensuring each agent’s identity and history can be cryptographically authenticated. 0G’s upcoming AI service marketplace will allow xNomad to distribute its AI models (the xNomad Brain) as on-demand services accessible across different dApps. By building on 0G, xNomad is demonstrating that AI agents can be owned, traded, and upgraded like assets, all while operating autonomously with on-chain verifiability.

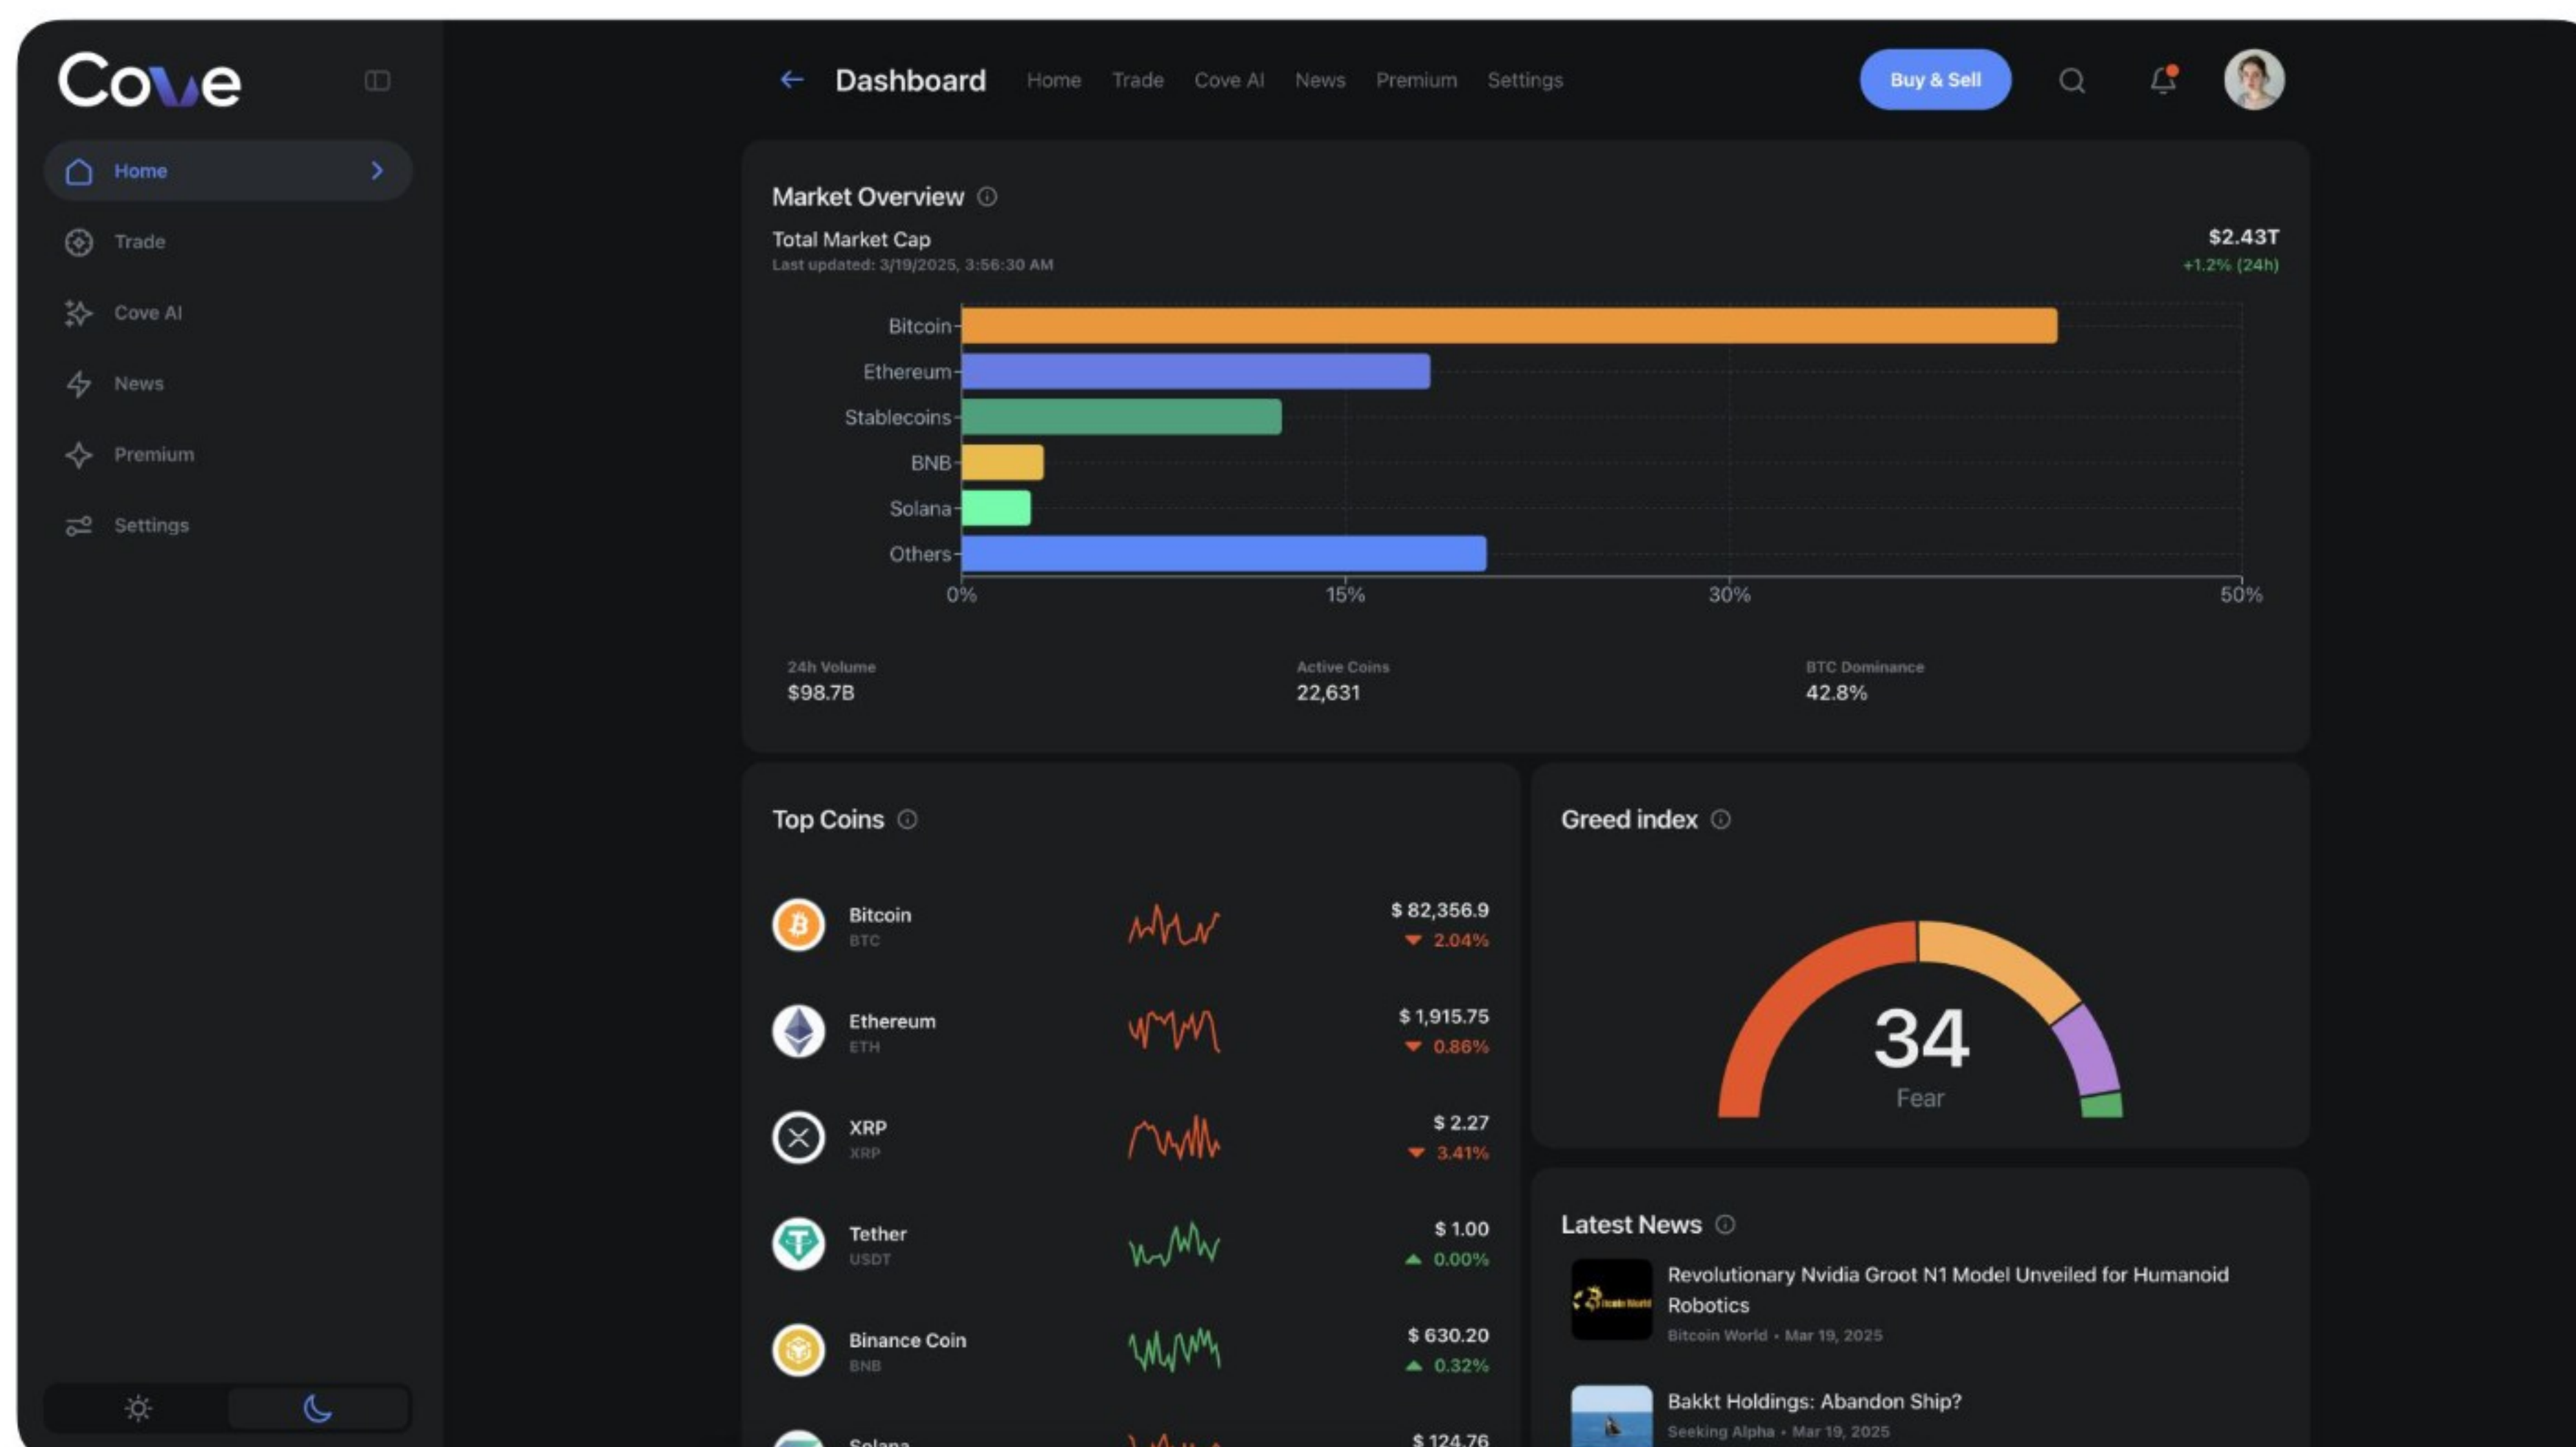
3. Cove – Wallet-Native AI Trading Agents

Cove is developing an AI-driven trading platform that operates directly from users’ wallets, eliminating reliance on centralized exchanges or interfaces. Its platform comprises WalletNet (a wallet-based network) and MEX (My Exchange), allowing users to tap into DEXs, CEXs, and OTC markets seamlessly through AI agents. The AI agents handle tasks such as optimizing trade routes across venues, monitoring market sentiment for strategic cues, and executing orders, while users retain custody of their funds.

Your AI-Powered Crypto Market Intelligence

Ask anything about any cryptocurrency. Get instant, detailed analysis with charts, trends, and insights powered by advanced AI.

Try Cove Now



Cove's agents aim to resist MEV (Maximal Extractable Value) exploitation and operate with privacy and autonomy. 0G serves as the compute and storage backbone for Cove's intelligent agents. Specifically, Cove leverages 0G's decentralized compute layer to run its high-frequency trading algorithms and sentiment analysis models right on the network. This provides secure, low-latency processing close to on-chain data. Simultaneously, market intelligence data generated by the AIs (such as strategy logs, trade history, and analytics) is stored via 0G's storage network in encrypted form. The encryption ensures tamper-proof records and auditability.

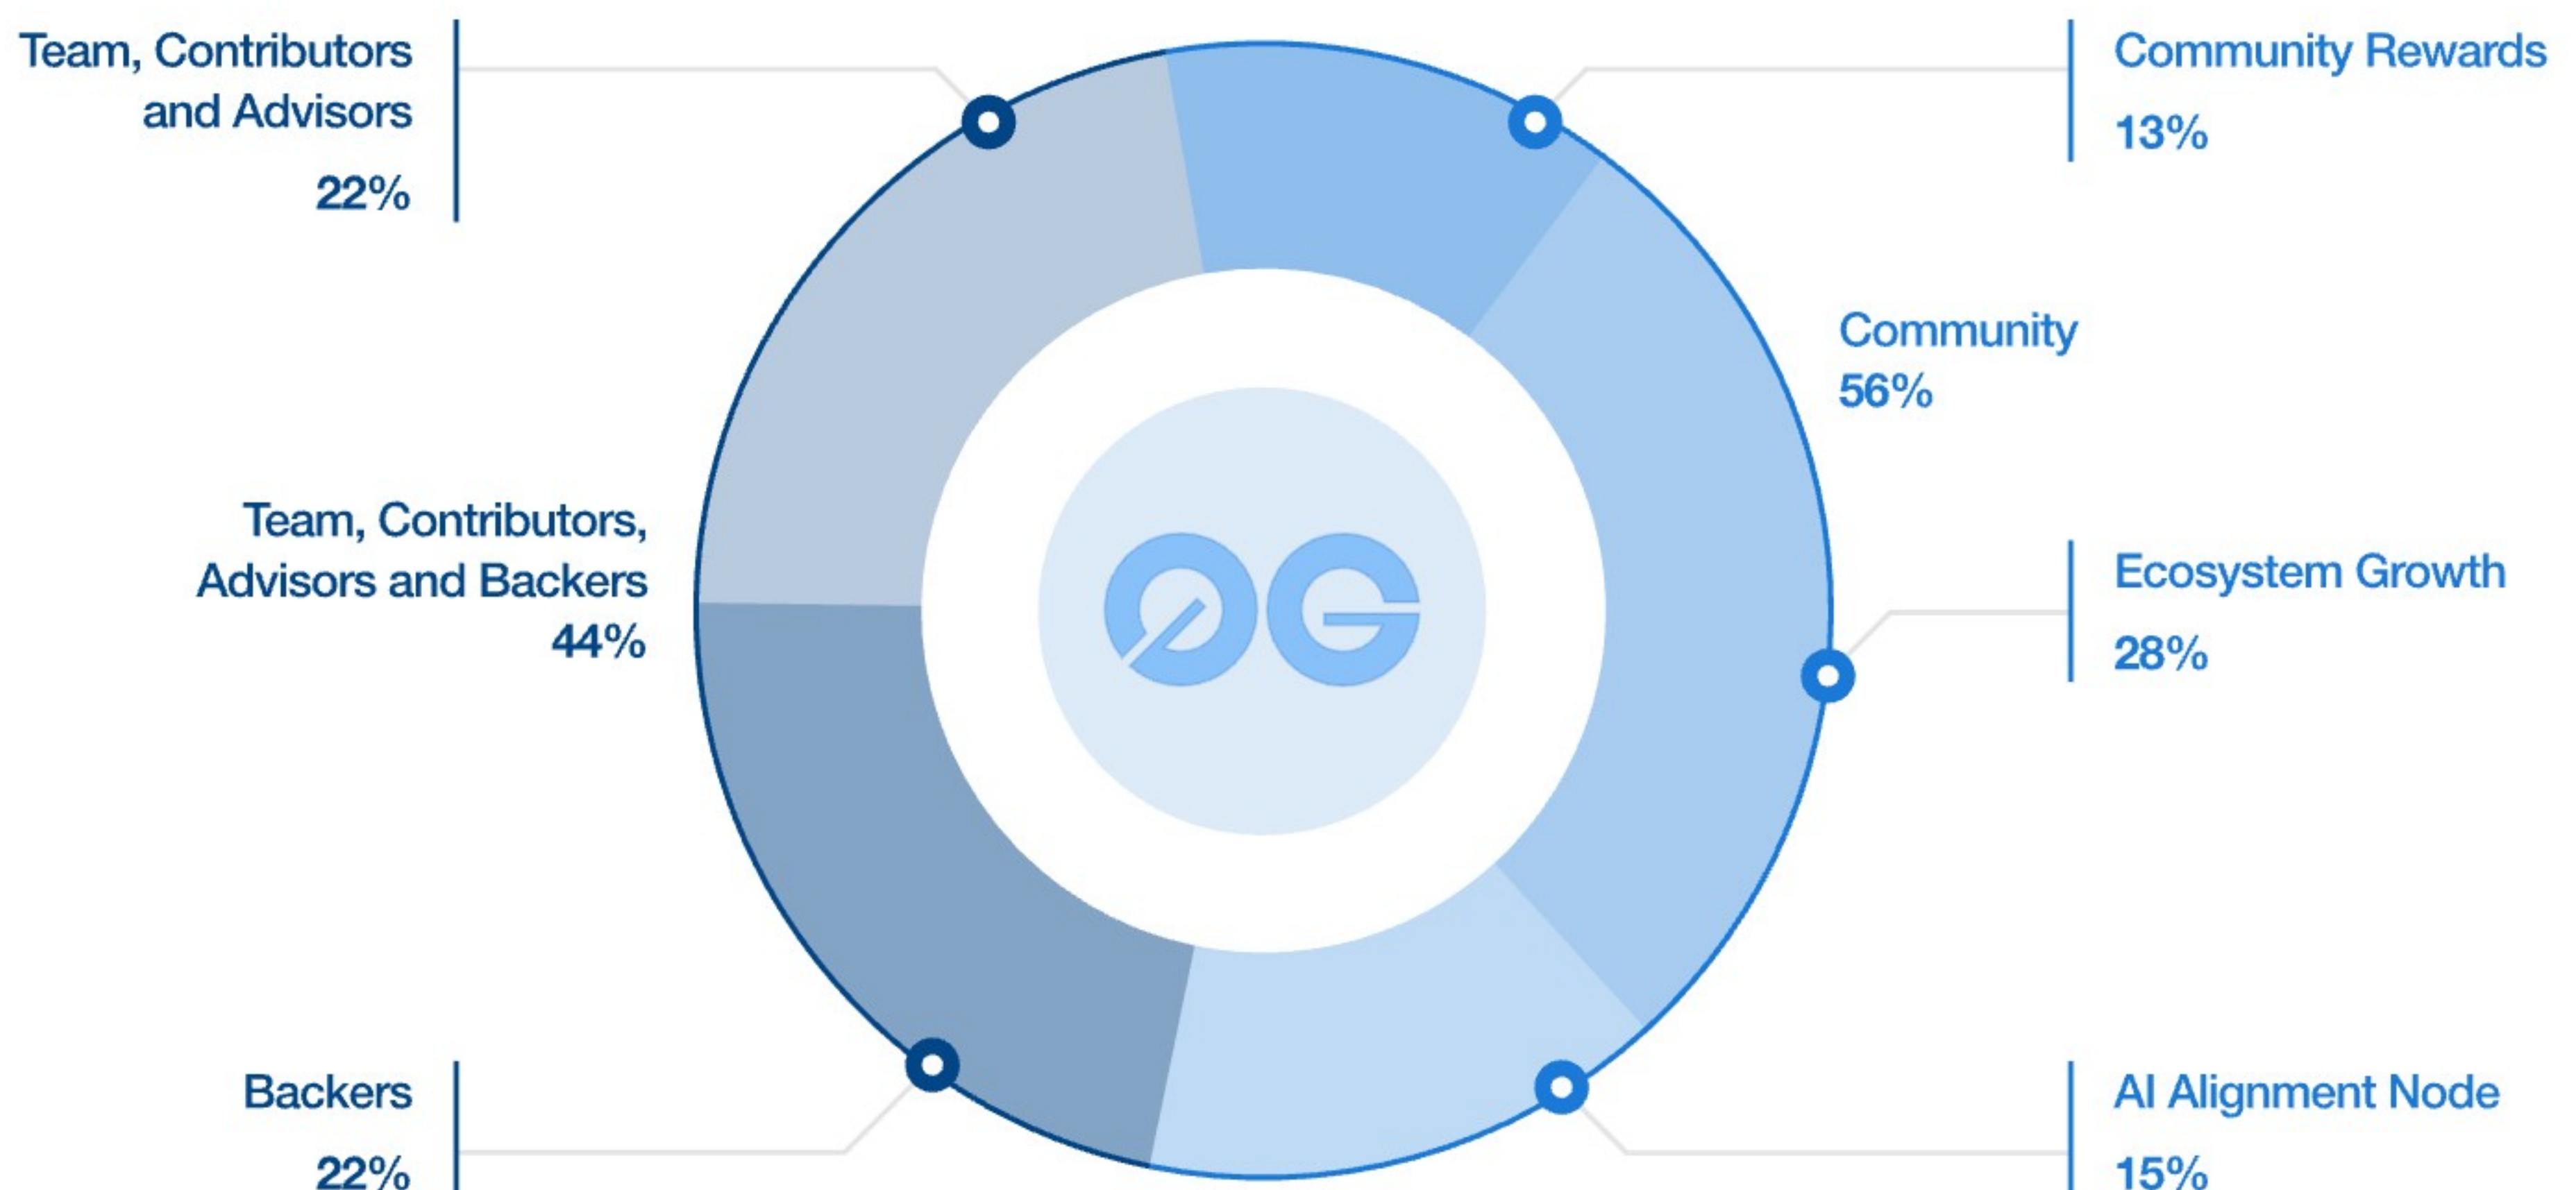
0G's fast data availability layer gives Cove's agents real-time access to on-chain order books and price feeds, which is critical for making split-second trading decisions. By building wallet-native AI agents on 0G, Cove envisions a trading infrastructure that is more efficient, self-custodial, and intelligent by design.

Looking Ahead: H2 2025 and Beyond

As 0G Labs moves into the second half of 2025, the focus shifts from testnet success to delivering mainnet and scaling the ecosystem for broader adoption. There will be considerable focus on the upcoming 0G mainnet launch, which is anticipated in H2 2025 (slightly later than the initial Q2 target, as the team took extra time to refine Galileo). The mainnet release will mark 0G's transition from experimental testnets to a production-grade network powering real economic activity.

According to the team, mainnet will carry forward Galileo's architecture with additional enhancements for performance and scalability. We can expect mainnet to incorporate everything proven in Galileo, plus any remaining features needed for full functionality. The mainnet launch will also involve rigorous security audits and likely a Token Generation Event (TGE) to distribute the 0G token to stakeholders in a manner consistent with the previously shared token allocation plan.

Token Allocation



Note: Allocations subject to change.

Source: 0g.ai/blog/clarifying-0g-s-token-allocation-our-commitment-to-the-community

Total Supply at Launch: 1,000,000,000

Title	Allocation	% of the allocation TGE	Lock-up (months)	Unlock (months)	Fully unlocked (months)
Backers	22.0%	0.0%	12	36	48
Team, Contributors and Advisors	22.0%	0.0%	12	36	48
Community	56.0%	38.07%			
AI Alignment Node	15.0%	33.33%*	0	36	36
Ecosystem Growth	28.0%	49.0%	0	24	24
Community Rewards	13.0%	20.0%	remainder distributed seasonally over 48 months		
	100.0%	21.32%			

Note: Table values assume a 0G token generation event coupled with 0G Mainnet launch. Allocations and schedules are subject to change.
* For the AI Alignment Node Sale - 33.33% of the tokens will be allocated at TGE. The tokens will be withdrawable in accordance with a community voted release schedule. Please follow OG socials for more details.

Source: 0g.ai/blog/clarifying-0g-s-token-allocation-our-commitment-to-the-community

To conclude, the second half of 2025 is set to be transformational for 0G. Mainnet launch will be the headline, bringing 0G’s vision of a decentralized AI Operating System into reality. Challenges will certainly arise (bootstrapping network effects is never trivial, and ensuring security at scale will be a continuous effort), but 0G has shown a propensity to plan ahead and iterate quickly.

“

Our mission is to bridge traditional finance into digital assets through our crypto native research.