SHARDY

DECENTRALIZED
MESH NETWORK

VISION

Shardy: The Decentralized Mesh for AI

Rescuing the AI computation bottleneck

Transforming every browser tab into a high-performance AI supercomputer

PROBLEM:

The AI Compute Bottleneck
AI is Hungry. Hardware is Scarce

Abstract:

- The monopoly of cloud giants (AWS, Google) is expensive and centralized;
- The chip shortage (H100/A100) is hindering innovation;
- Billions of consumer GPUs (MacBooks, Gaming PCs) are idle 90% of the time

Our point:

Enormous power is at our fingertips, but it's not being used

SOLUTION:

Browser-Native Supercomputing
Zero Install. Infinite Scale

Abstract:

- Bio-Mechanical Runtime: No installation required. Any browser becomes a computing node via WebGPU;
- WebGPU + WASM: Direct access to graphics card hardware from a browser tab;
- Permissionless: Becoming part of the network is easier than opening YouTube

Our point:

We're removing the barrier to entry for computing power providers and node validators

PRODUCT:

How Shardy Works
From Task to Settlement in Seconds

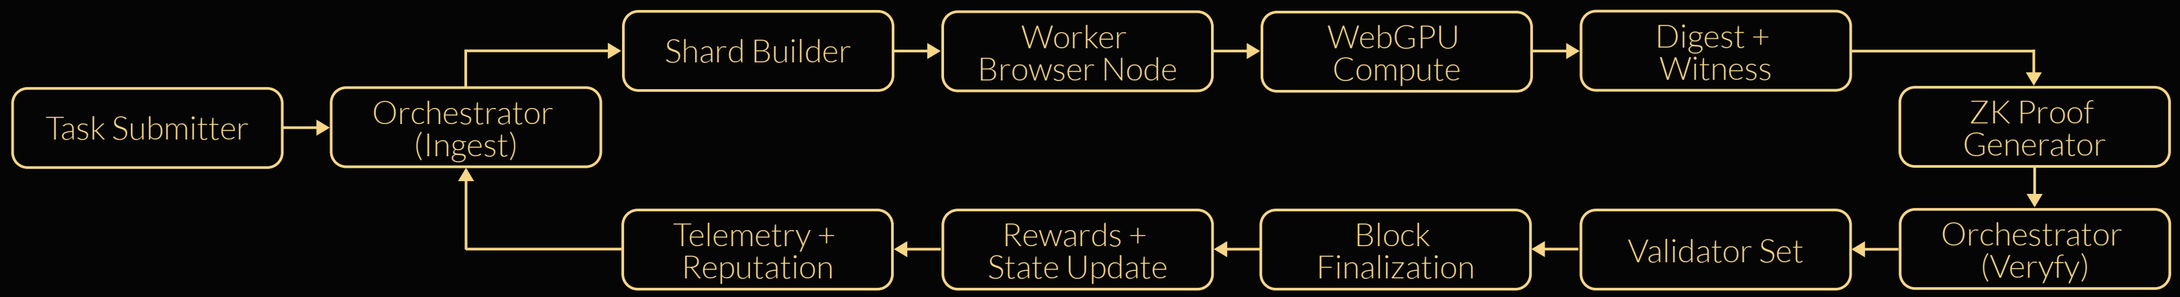

Key features:

- Intelligent Routing: the system automatically understands whether your iPhone can handle a given task or whether a powerful PC is needed;
- Redundancy: Each task is duplicated to protect against errors

TECHNOLOGY:

Mathematical Certainty (The ZK-Factor)
Trustless Compute via ZK-SNARKs

Abstract:

- How to trust an anonymous browser? Mathematics;
- Groth16 (Circom): Generating a short ZK proof that the computation was performed correctly;
- Verification in less than 10 ms: Verification of a powerful computing device takes milliseconds on the network side

Our point:

The results are mathematically impossible to forge

CONTROL PLANE:

The Orchestrator Mesh
A High-Performance Control Plane

Tech stack:

- Libp2p / GossipSub: Decentralized task broadcasting that eliminates bandwidth bottlenecks;
- Intelligent Matchmaking: Mathematical distribution of tasks across WorkerTiers based on GFLOPS and device VRAM;
- Byzantine Fault Tolerance: Using a redundancy factor (REDUNDANCY_FACTOR = 2) to instantly detect manipulation;
- Dead-Letter Queues (DLQ): Automatic isolation of problematic tasks in SQLite to maintain network stability.

Key takeaway:

Shardy is a self-organizing network where management occurs on the fly, with no central point of failure

for more info

EXECUTION PIPELINE:

The Bio-Mechanical Runtime
From Raw Data to ZK-Settlement

01

Zero-Copy Execution:

- Direct tensor mapping to the GPU via WebGPU (TypeGPU);
- Use of SharedArrayBuffer for instant data transfer without unnecessary memory copying

02

Deterministic Integrity:

- WASM EMA Smoothing: A Rust/WASM algorithm that guarantees identical computation results on Nvidia, AMD, and Apple Silicon;
- OPFS Mass Storage: Caching gigabyte models directly to the user's disk via the browser file system.

03

ZK-SNARK Verification (The Trust Factor):

- Groth16 (Circom 2.0): Generates a short cryptographic proof of task execution;
- Sub-10ms Verification: The network verifies the authenticity of the result in milliseconds without recalculating terabytes of data.

for more info

MARKET & STRATEGY:

Global Cloud Computing Market:

1. Global cloud computing market: Total size of the industry controlled by monopolies (AWS, Google, Azure);
2. Target potential: Total capacity of idle consumer GPUs (Apple Silicon, Nvidia);
3. Driver: Global shortage of AI infrastructure and hardware

Serviceable AI Inference & DePIN segment:

1. AI Inference & DePIN Segment: Market size for neural network inference and decentralized infrastructure;
2. Available capacity: 110M+ gaming PCs and MacBooks capable of WebGPU computing;
3. Advantage: Cost reduction of 40-60% due to the absence of capital investment in data centers

Obtainable market share for decentralized mesh-compute:

1. Lattice Era Share (3-5 years): Achievable volume by capturing the market for burst workloads and AI startups;
2. Product: Scalable LLM inference and image generation via a browser-based network;
3. Economics: Target cost per task is $1.45 per 25 TFLOPs, ensuring high competitiveness

for more info

MARKET & STRATEGY:

Go-to-Market Strategy (GTM)

1. Zero-installation deployment: leverage millions of idle GPUs through a browser interface (zero barrier to entry);

2. Developer SDK: Tools for easily deploying LLM and neural networks on the Shardy network;

3. Enterprise Pilot Projects: Pilot projects for companies requiring elastic AI inference

ECOSYSTEM:

Multi-Chain Synergy
Empowering Any Ecosystem

Our point:

- Off-chain Compute, On-chain Settlement: Heavy AI tasks are solved in Shardy, and the final ZK proof and payment are sent to your chain;
- AI Oracle: Shardy can serve as a decentralized oracle for AI data;
- Low Barrier: We bring millions of users to Web3 through their browsers, without requiring them to install wallets at once

THE $SHRD ECONOMY:

Aligning Incentives through Math

TOKENOMICS: THE FUEL OF DECENTRALIZED COMPUTE

01

Key Roles and Utility:

- Compute Providers (Workers): Stake $SHRD to protect against Sybil attacks and receive token rewards for completing AI tasks;
- Validators: Stake significant capital to validate ZK proofs and receive network fees;
- Submitters (AI Clients): Pay for computations in $SHRD tokens

Key takeaway:

Payout = Base_Compute * Quality_Score * Reputation_Multiplier

02

Economic Mechanisms:

- Dynamic Pricing: Computation costs adapt to supply and demand in real time;
- Reputation Multiplier: Honest nodes with a high SSR (Shard Success Rate) earn more;
- Slashing Mechanics: Attempts to manipulate computations result in immediate loss of stake (mathematical penalty)

for more info

Mobile ZK-Proofs

$HARDY

ROADMAP:

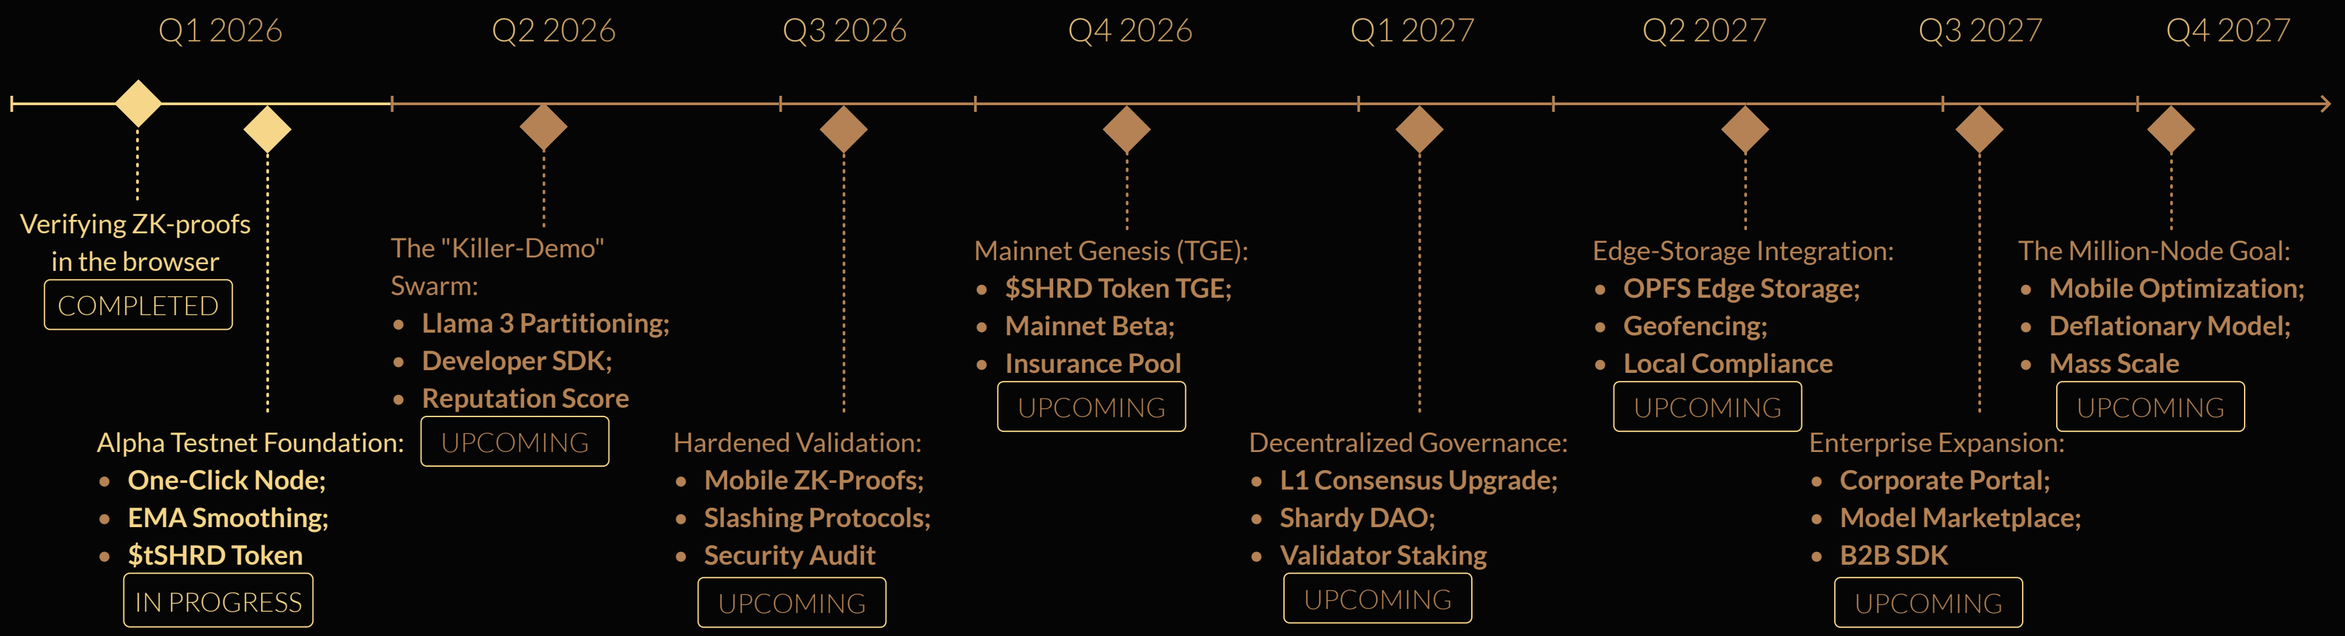

for more info

CORE TEAM:

Ihor Sokolov

CEO & Web3 Developer
foundermafstat@gmail.com
https://x.com/iampublion

Irina Semichasova

Co-founder & Full Stack
isemichasova@gmail.com
https://x.com/irine_es

Join the Mesh

Shardy isn't just a network. It is the distributed brain of the open internet

CONNECT NODE

Website

Documentation