

A DECENTRALIZED NETWORK FOR HIGH- EFFICIENCY AI COMPUTE

Anastasia Matveeva, Co-creator, Gonka
PhD, UPC Barcelona

THE AI COMPUTE BOTTLENECK IS NOW PHYSICAL, NOT ALGORITHMIC

2026 US DATA CENTER PIPELINE:
announced VS under construction



PRIMARY DRIVERS:

- power transformers: 2-4 year lead times
- grid interconnection: up to 5-year utility queues
- local permitting and community opposition

FEDERATION IS THE ARCHITECTURE THAT SCALES

25 YEARS OF FEDERATED COMPUTE:



FEDERATED AI COMPUTE IS IN PRODUCTION TODAY

4K GPUS OVER 26 COUNTRIES*



*as of April 2026

Anastasia Matveeva, Co-creator of Gonka protocol

GONKA

PERMISSIONLESS, VERIFIABLE, DECENTRALIZED AI COMPUTE

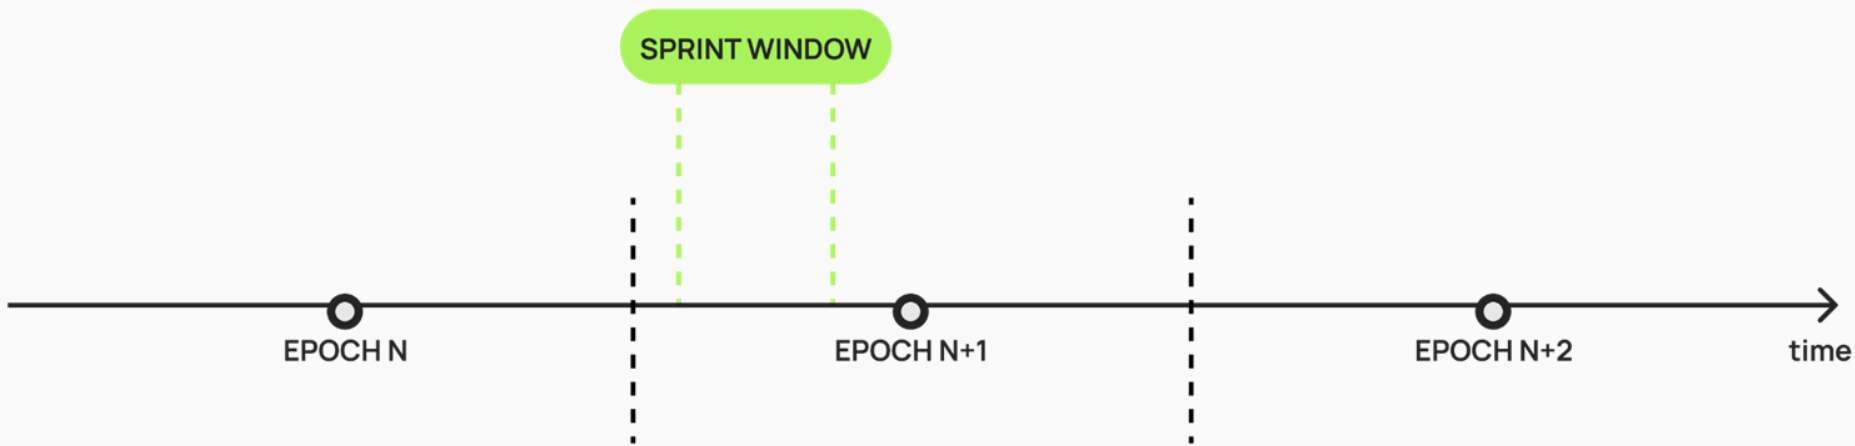
Gonka is a permissionless, verifiable, decentralized compute network.

Permissionless and decentralized means geo-distributed, trustless, and free of any single point of control; blockchain provides those governance primitives.

The cloud is reshaped by AI. Blockchain is shaped by finance. Gonka was designed from scratch with AI alignment as the foundation.

ALIGNMENT WITH AI WORKLOADS IS EFFICIENCY

THE SAME COMPUTE THAT SECURES THE CHAIN PRODUCES AI OUTPUT.



VERIFICATION IS BUILT IN

The chain secures the exact model parameters running on the network: no silent quantization and no silent model swaps.

TEE is one optional layer for workloads requiring hardware-isolated confidentiality. ZKP and FHE are on the research frontier.

The system is verification-agnostic by design: whatever primitive matures, Gonka can adopt it without redesign.

THE CHAIN IS THE RECORD LAYER, NOT THE COMPUTE LAYER

The chain handles coordination and recordkeeping for inference and training to run off-chain.

The chain itself shards by independent model-specific workloads.

Result: horizontal scaling of chain coordination without sacrificing integrity.

FOR HARDWARE PROVIDERS

| MINIMUM HARDWARE TODAY:

Hardware-agnostic by design: NVIDIA, AMD or ASICs, anything that produces verifiable work and performance for supported models.

| EARNINGS HAVE TWO STREAMS:

Protocol-emission rewards and per-work revenue from inference users.

| NO LOCK-IN:

No exclusivity, no long-term commitment, join when your hardware sits idle

**OPERATORS ARE SHAPING THE FUTURE OF GLOBAL AI INFRASTRUCTURE.
GOVERNANCE FOLLOWS COMPUTE.**

FOR WORKLOAD OWNERS

| COST

As the network grows, order of magnitude lower than centralized providers

| EASE OF USE

Easy integration. OpenAI-compatible API. Top-tier open-source models (Kimi K2.6 and others).

| INTEGRITY

Chain commits to exact model parameters. No silent quantization, model or configuration tweaks. Reproducibility and sovereignty over your inference pipeline.

THE LONG ARC

| CO-LOCATION IS A CONSTRAINT

Single-site density runs into power, permitting, and grid limits. Lift that requirement and the GPUs already deployed worldwide are plenty.

| INFERENCE WORKS

Operational reality. The architecture works at scale today.

| TRAINING IS THE OPEN FRONTIER

Active research building on DiLoCo (DeepMind), Nous Psyche, Prime Intellect, and others to pretrain a top-tier model in a fully decentralized setup.

CONTACTS



Technical documentation



anastasia.matveeva@productscience.ai



LEARN MORE

