The moat: we patented the hard part.

Patented scheduling

Intelligently routes each task to the fastest, most reliable nodes in real time. US patent granted, EU pending.

Fault tolerant

Detects node failures and reroutes work automatically, so every job finishes.

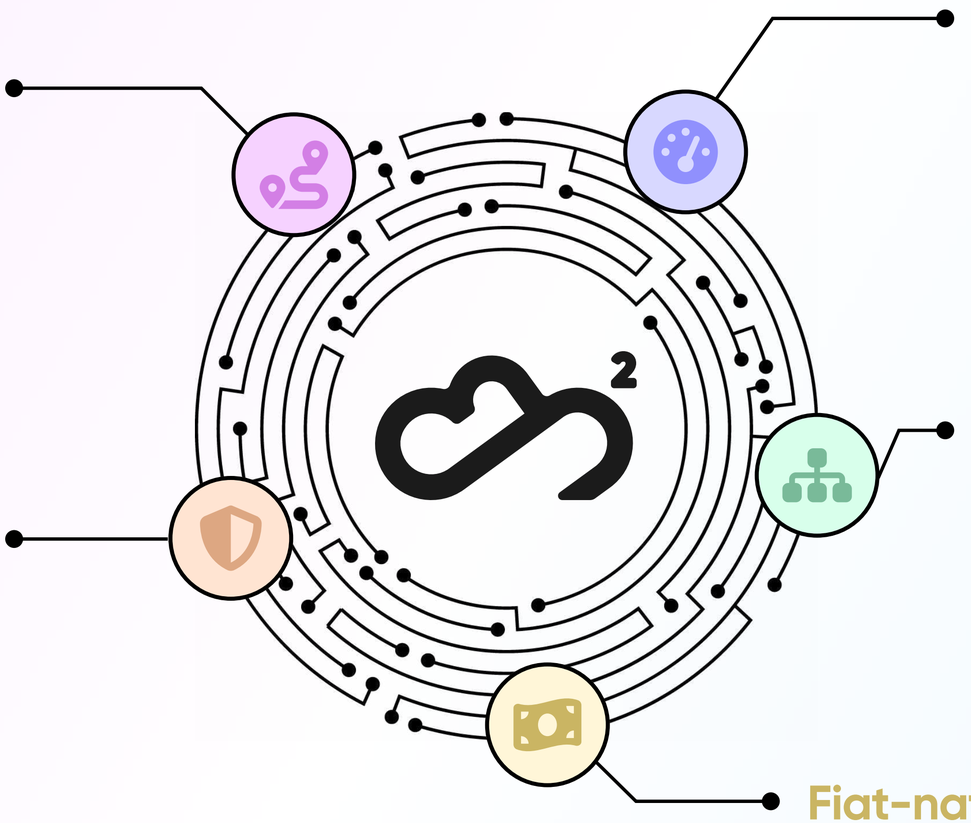

Reliable by design

Turns an unreliable, heterogeneous fleet into SLA-backed, predictable compute.

Automatic parallelization

Splits any parallel job across the fleet and reassembles the results, no orchestration code.

Fiat-native

Real-money billing and payouts. Enterprises procure with confidence; workers cash out.

The market

CAGR: 44%

SOM: $8.7B

Global GPU-as-a-Service (GPU rental) market

Source: Fortune Business Insights, 2026

SAM: $20.6B

Public cloud end-user spending on AI inference and batch processing services.

Source: Deloitte, 2026

TAM: $142.8B

Global spending on AI infrastructure

Source: Evolvance, 2026

Reliable like a hyperscaler.
Cheap like a marketplace.

ByteNite	DePIN GPU (Aethir)	Hyperscaler	Consumer GPU
Low cost	Low cost	High cost	Low cost
Full managed stack	Marketplace only	Full managed stack	Marketplace only
Fiat billing	Crypto-native	Fiat billing	Fiat billing
QoS / SLA-backed	No SLA	QoS / SLA-backed	No SLA

Customers pay per second for the CPU, GPU, RAM, and network they actually use.

Usage is tracked and itemized in real time, so customers can monitor spend and plan their budget.

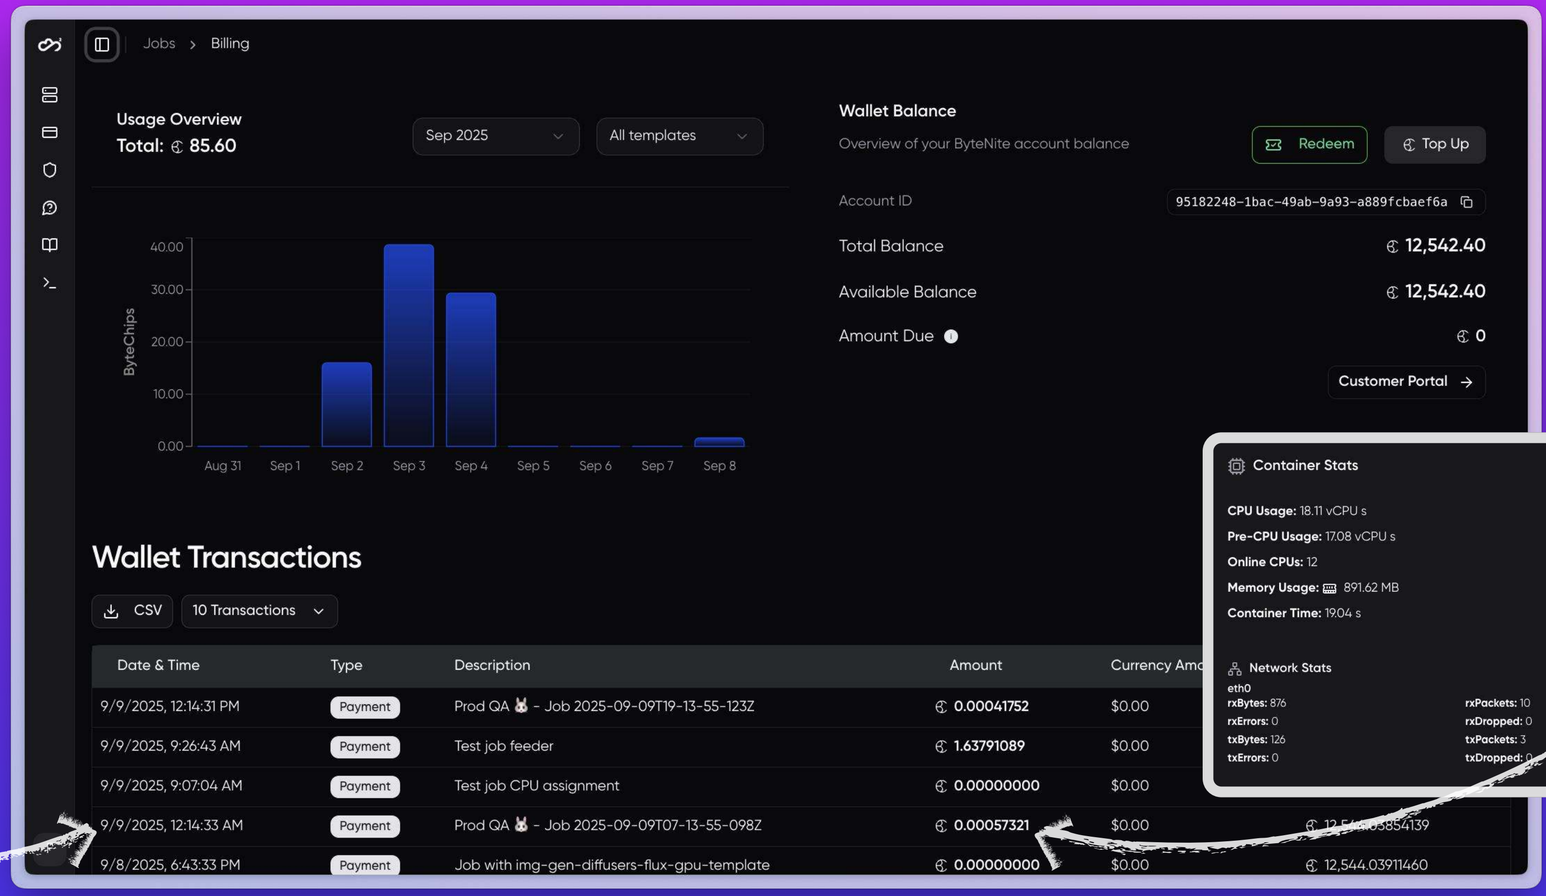

Every job is a transaction: suppliers earn a real-money payout, and ByteNite keeps a margin.

ByteNite

Thank You

ByteNite

The distributed cloud for AI

LinkedIn

linkedin.com/company/bytenite

Website

www.bytenite.com

Documentation

docs.bytenite.com

Fabio Caironi, CEO & Founder, fabio@bytenite.com