



THE DECENTRALIZED GPU ECONOMY.

# FUELING THE FIRE OF OPEN COMPUTE

**A SOLANA-NATIVE GPU COMPUTE MARKETPLACE**

A permissionless, on-chain protocol that transforms idle GPU capacity into a globally accessible compute layer for artificial intelligence, rendering, and the next generation of decentralized applications.

**80%**

BELOW CLOUD PRICING

**PER-SECOND**

BILLING ENGINE

**ON-CHAIN**

PROOF-OF-COMPUTE

// SECTION 00

# TABLE OF CONTENTS

NAVIGATE THE PROTOCOL

01	<b>EXECUTIVE SUMMARY</b>	The mission, the market, the moment.
02	<b>THE COMPUTE CRISIS</b>	Why centralized cloud cannot scale AI.
03	<b>PROTOCOL OVERVIEW</b>	Pyronet as a decentralized GPU layer.
04	<b>ARCHITECTURE</b>	Five-layer protocol stack on Solana.
05	<b>PROOF-OF-COMPUTE</b>	Trustless verification of every workload.
06	<b>MARKETPLACE MECHANICS</b>	Renters, providers, and per-second settlement.
07	<b>GPU SUPPLY LAYER</b>	From consumer cards to enterprise fleets.
08	<b>USE CASES</b>	Six demand surfaces, one network.
09	<b>WHY SOLANA</b>	The only chain fast enough for compute.
10	<b>SECURITY MODEL</b>	Slashing, escrow, and economic guarantees.
11	<b>GOVERNANCE</b>	Network ownership and protocol evolution.
12	<b>ROADMAP</b>	Path to one million distributed GPUs.
13	<b>CONCLUSION</b>	The fire is lit.

---

THIS DOCUMENT OUTLINES THE TECHNICAL ARCHITECTURE, ECONOMIC DESIGN, AND STRATEGIC VISION OF THE PYRONET PROTOCOL — A SOLANA-NATIVE COMPUTE MARKETPLACE ENGINEERED FOR THE AI ERA.

# EXECUTIVE SUMMARY

THE MISSION · THE MARKET · THE MOMENT

Pyronet is a decentralized GPU compute marketplace built natively on the Solana blockchain. The protocol coordinates a global, permissionless network of graphics processing units — transforming consumer cards, mining rigs, and underutilized enterprise hardware into a single, on-chain compute layer accessible to anyone with a wallet.

The protocol delivers compute capacity at up to **eighty percent** below the cost of centralized hyperscalers such as AWS, Google Cloud, and Microsoft Azure. Workloads are billed by the second. Settlements occur on-chain. Verification of completed jobs is enforced by cryptographic proof-of-compute. There are no contracts, no lock-in periods, and no gatekeepers.

## // THE PYRONET THESIS

Artificial intelligence is the defining workload of this decade, and GPU capacity is the bottleneck. The hyperscalers cannot meet demand at any reasonable price. Pyronet routes demand to the dormant compute already in the world — verified on-chain, settled per second, and owned by the network.

The convergence of three forces — the artificial intelligence supercycle, the rise of Decentralized Physical Infrastructure (DePIN), and the emergence of Solana as a high-throughput settlement layer — has created the conditions for an open compute economy. Pyronet is engineered for that economy.

Three pillars define the protocol. First, a **permissionless marketplace** where any qualified GPU can join the network and earn for completed work. Second, a **cryptographic verification system** that ensures every workload is genuinely executed before payment is released. Third, a **per-second settlement engine** built on Solana's sub-cent fees and sub-second finality.

**80%**

BELOW CLOUD PRICING

**1s**

BILLING GRANULARITY

**\$0.00**

AVG TX FEE

**PoC**

VERIFICATION

# THE COMPUTE CRISIS

WHY CENTRALIZED CLOUD CANNOT SCALE AI

The world is running short of compute. Demand for GPU resources has outpaced the ability of centralized providers to deliver capacity at any price. Enterprise AI teams report waiting lists measured in quarters, not weeks. Academic researchers are priced out entirely. Independent studios watch render queues stretch for days. The bottleneck is structural, not temporary.

Three failures define the centralized cloud model:

// 01

### ARTIFICIAL SCARCITY

Hyperscalers throttle access to premium GPUs through allocation queues and minimum spend tiers, despite vast idle capacity worldwide.

// 02

### OPAQUE PRICING

Hourly rates, hidden egress fees, and complex commitment tiers make true unit economics impossible to calculate at scale.

// 03

### VENDOR LOCK-IN

Proprietary APIs, custom storage formats, and migration costs create deep moats that capture customers indefinitely.

Meanwhile, an enormous shadow fleet of capable GPUs sits idle. Gaming PCs run at a fraction of their potential outside peak hours. Cryptocurrency mining rigs went offline after the Ethereum merge and never returned. Enterprise data centers operate at thirty to forty percent utilization. Render farms cycle dark for entire calendar quarters. This is the largest unmonetized compute reserve in human history.

### // THE OPPORTUNITY

The global GPU-as-a-Service market is projected to surpass four hundred billion dollars by the end of the decade. The centralized cloud cannot meet this demand alone — and would not if it could. The natural response is a coordination layer that pools distributed hardware into a single liquid market. That coordination layer is Pyronet.

The historical analogue is precise. Telecommunications, ride-sharing, lodging, and energy production each underwent the same transition: centralized incumbents gave way to networked marketplaces that pooled distributed supply. Compute is the last category that has not yet been unbundled. Pyronet completes the pattern.


# PROTOCOL OVERVIEW

PYRONET AS A DECENTRALIZED GPU LAYER

At its core, Pyronet is a coordination protocol. It does not own GPUs. It does not operate data centers. It does not warehouse hardware. Instead, it provides the open, on-chain infrastructure through which independent compute supply and compute demand discover each other, transact, and settle — without intermediaries.

The protocol abstracts every participating GPU into a unit of measurable, verifiable compute. Providers register their hardware, declare their capabilities, and accept jobs. Renters submit workloads, specify their requirements, and pay only for the seconds their job actually runs. An automated matching engine, an on-chain escrow mechanism, and a cryptographic verification system guarantee both sides of the trade.

## // NETWORK TOPOLOGY



CORE: PROTOCOL ROUTER

MID: GPU PROVIDERS

EDGE: COMPUTE RENTERS

**Three principles guide protocol design.**

## // PERMISSIONLESSNESS

Anyone with a qualified GPU can join the supply side. Anyone with a Solana wallet can join the demand side. No application process, no credit check, no geographic restriction.

## // VERIFIABILITY

Every claimed unit of work must be cryptographically proven before settlement. Renters never pay for compute that did not occur. Providers always receive payment for compute that did occur.

## // COMPOSABILITY

Pyronet is a primitive, not a destination. Its compute layer is callable from any Solana program, integrable into any AI pipeline, and exposed through open APIs.

# ARCHITECTURE

FIVE-LAYER PROTOCOL STACK ON SOLANA

Pyronet is structured as a five-layer protocol stack. Each layer performs a discrete function and exposes a clean interface to the layer above. This modular design allows individual components to be upgraded, audited, or replaced without disturbing the rest of the system.



**L1 — Hardware Layer.** The physical foundation: thousands of independent GPU nodes operated by gamers, miners, studios, and data centers. The protocol supports a wide spectrum of devices, from consumer cards in the 4060 class through enterprise accelerators including the H100 and A100 families. Every node runs a lightweight client that reports availability, executes assigned workloads, and submits cryptographic proofs.

**L2 — Settlement Layer.** All economic activity is settled on Solana. Funds are escrowed at job initiation, streamed in per-second increments as work progresses, and released upon successful verification. Solana's sub-second finality and fractional-cent fees make per-second billing economically viable — a primitive impossible on any other public chain at scale.

**L3 — Verification Layer.** A proof-of-compute system ensures that providers cannot claim payment for work they did not perform. Workloads emit verifiable execution traces, which are sampled, challenged, and validated by an independent verifier set. Dishonest providers are economically slashed.

**L4 — Orchestration Layer.** The matching engine. It accepts incoming workloads, parses their requirements (GPU class, memory, geography, latency), and routes them to the optimal available provider. Service-level guarantees are enforced through automated failover and reassignment.

**L5 — Application Layer.** The surface that builders touch. AI training jobs, inference endpoints, render pipelines, scientific simulations, and on-chain AI agents — all expressed through standard interfaces, all running on the same underlying compute fabric.

# PROOF-OF-COMPUTE

TRUSTLESS VERIFICATION OF EVERY WORKLOAD

The defining technical challenge of decentralized compute is verification. In a centralized cloud, the provider's reputation guarantees the work; if a customer claims their job never ran, the hyperscaler can produce logs from infrastructure it owns end-to-end. In a permissionless network, no such trust anchor exists. Pyronet replaces it with cryptographic proof.

## // COMPUTE LIFECYCLE



The verification protocol operates in three phases.

**// PHASE ONE — INSTRUMENTATION.** Before execution begins, the workload is wrapped in a deterministic runtime that emits cryptographic checkpoints at fixed intervals during execution. These checkpoints capture the state of the computation in a form that cannot be forged after the fact.

**// PHASE TWO — ATTESTATION.** Upon completion, the provider submits a final attestation containing the checkpoint hashes, the result, and a signed claim of the resources consumed. This package is committed to Solana as a single transaction.

**// PHASE THREE — CHALLENGE.** A subset of independent verifiers is randomly selected to challenge the attestation. Challenged providers must reveal a deterministic re-execution of the contested segment. Failure to produce a valid response — or any divergence from the original output — triggers automatic slashing of the provider's stake.

## // ECONOMIC FINALITY

Proof-of-compute is not only cryptographic, it is economic. Providers post stake before accepting jobs. Successful work returns the stake plus payment. Detected fraud results in slashing. Over time, only honest providers survive in the network. The protocol does not ask providers to be trustworthy — it makes dishonesty unprofitable.

# MARKETPLACE MECHANICS

RENTERS · PROVIDERS · PER-SECOND SETTLEMENT

Pyronet operates as a continuous, open marketplace. Prices are not set by the protocol — they are discovered by the market, in real time, through the independent decisions of thousands of competing providers. The protocol's role is purely infrastructural: it provides the rails, the matching, the verification, and the settlement. Everything else emerges from open competition.

## // LIVE MARKETPLACE SNAPSHOT

Representative pricing across the GPU spectrum. Per-second billing applies; hourly rates shown for reference.

GPU MODEL	VRAM	PRICE / HOUR	USE CASE
NVIDIA RTX 3060	12 GB	\$0.048/h	Inference · Light Training
NVIDIA RTX 3090	24 GB	\$0.192/h	Stable Diffusion · Fine-Tuning
NVIDIA RTX 4090	24 GB	\$0.320/h	Mid-Scale Training · Rendering
NVIDIA RTX 5090	32 GB	\$0.400/h	Modern Inference · Generative AI
NVIDIA A6000	48 GB	\$0.400/h	Enterprise Workloads
NVIDIA A100 40GB	40 GB	\$0.610/h	Large Model Training
NVIDIA A100 80GB	80 GB	\$0.900/h	Foundation Model Workloads
NVIDIA H100	80 GB	\$1.500/h	Frontier-Scale Compute

**Pricing dynamics.** Providers set their own hourly rates and can adjust them in response to demand. Renters see a unified order book of available capacity and choose by price, location, and provider reputation. The result is a price discovery surface that consistently lands well below centralized cloud benchmarks — often by a factor of five to ten on equivalent hardware.

**Per-second granularity.** Inference workloads often run for seconds, not hours. Centralized clouds round up to minutes or hours, charging for capacity the customer never used. Pyronet bills exclusively for the seconds compute is actually consumed. Over a typical inference fleet, this granularity alone reduces costs by an additional twenty to forty percent beyond raw hardware savings.

**No lock-in.** There are no minimum commitments, no reserved capacity tiers, no exit fees, and no migration penalties. Workloads can be ported in or out in real time. The protocol's openness is its product.

# GPU SUPPLY LAYER

FROM CONSUMER CARDS TO ENTERPRISE FLEETS

The protocol welcomes any qualified GPU — consumer, professional, or data-center class. This deliberate breadth is the source of Pyronet's pricing advantage. Where centralized providers must amortize the capital cost of every card across customer billing, distributed providers operate on already-sunk hardware. Their marginal cost is electricity and bandwidth. The savings flow to renters.

// TIER 1

### CONSUMER NODES

Gaming PCs, prosumer workstations, and home labs. Operators contribute idle hours between gaming sessions and earn for completed jobs.

// TIER 2

### MINING FLEETS

Former cryptocurrency mining operations redeploy GPUs for AI inference. Existing power infrastructure becomes new productive capacity.

// TIER 3

### DATA CENTERS

Independent operators and tier-three facilities offer enterprise GPUs (A100, H100) at competitive rates without hyperscaler markups.

**Onboarding is friction-light.** A prospective provider installs the Pyronet node client, completes a one-time benchmark to establish the device's performance profile, posts the stake required for the tier they wish to serve, and begins accepting jobs. The benchmark prevents misrepresentation of hardware capabilities. The stake prevents bad behavior. Together they preserve marketplace integrity without requiring centralized review.

**Reputation accrues over time.** Every completed job, every uptime measurement, every verification challenge contributes to a provider's on-chain reputation score. Higher reputation unlocks access to higher-value workloads and preferred routing in the matching engine. The system rewards consistency and punishes neglect.

**Geographic diversity is a feature.** Pyronet nodes operate across every populated continent. This distribution provides natural latency advantages for edge inference, robust availability against regional outages, and access to workloads that specifically require data residency in particular jurisdictions.

# USE CASES

SIX DEMAND SURFACES · ONE NETWORK

Pyronet is designed for the workloads that define this computational era. The protocol's abstraction is intentionally general: any application that requires GPU acceleration can be expressed against the network. Six categories of demand have emerged as primary surfaces.

## // 01 — AI MODEL TRAINING

Early-stage AI teams, fine-tuning specialists, and independent labs train and adapt foundation models on Pyronet at a fraction of hyperscaler cost. Long-running training jobs can be checkpointed and resumed across nodes, allowing dynamic capacity reallocation as priorities shift.

## // 02 — INFERENCE AT SCALE

Production AI applications — chatbots, copilots, recommendation engines — serve inference traffic through Pyronet's distributed pool. Geographic distribution reduces latency. Per-second billing matches the bursty nature of real-world inference demand.

## // 03 — 3D & VFX RENDERING

Independent animation studios, VFX houses, and architectural visualization firms tap a global render farm without committing to fixed infrastructure. Jobs verified on-chain, payment streamed by the second.

## // 04 — ON-CHAIN AI AGENTS

Web3 builders constructing autonomous AI agents on Solana require infrastructure that the agents themselves can call directly, without centralized API keys or rate-limited endpoints. Pyronet provides this primitive.

## // 05 — SCIENTIFIC COMPUTE

Independent researchers run molecular dynamics, climate models, computational biology, and red-team simulations on cost-efficient hardware. Compute that previously required grant applications now requires a wallet connection.

## // 06 — GENERATIVE MEDIA

Image generation, video synthesis, voice cloning, and audio production pipelines run at consumer-tier costs while preserving enterprise-grade verifiability.

# WHY SOLANA

THE ONLY CHAIN FAST ENOUGH FOR COMPUTE

Pyronet's per-second billing engine cannot exist on most blockchains. A typical compute job generates dozens to hundreds of state updates per minute as usage accrues, escrow streams, and verification events occur. To settle these economically, the underlying chain must offer three properties simultaneously: high throughput, low fees, and sub-second finality. Solana is the only production network that delivers all three at the scale Pyronet requires.

<b>65,000+</b>	<b>400ms</b>	<b>\$0.00025</b>	<b>24/7</b>
PEAK TPS CAPACITY	BLOCK FINALITY	AVG FEE	GLOBAL UPTIME

**Throughput.** Solana's parallel transaction processing model handles tens of thousands of operations per second. Pyronet can settle high-frequency billing events for thousands of concurrent jobs without competing for blockspace.

**Fees.** A typical Pyronet settlement event costs a fraction of a cent in network fees. This makes per-second billing economically rational; on a chain with dollar-level fees, the overhead alone would exceed the cost of the compute itself.

**Finality.** Sub-second finality means that providers see payment confirmed almost immediately upon valid job completion. Renters see escrow movements in real time. Trust is established at human-perceptible speed.

**Composability.** The Solana ecosystem — including its stablecoin liquidity, decentralized exchanges, and growing AI-agent layer — is composable with Pyronet by default. Any Solana program can call Pyronet's compute marketplace as a primitive, and any Pyronet workload can transact with the broader on-chain economy.

## // WHY NOT ETHEREUM OR L2S

Ethereum mainnet's settlement costs are incompatible with per-second compute billing. Layer-two networks introduce additional bridging complexity, withdrawal delays, and fragmented liquidity. Solana's single, fast, cheap layer-one settlement is the cleanest fit for a compute marketplace. The protocol meets the workload where the workload lives.

# SECURITY MODEL

SLASHING · ESCROW · ECONOMIC GUARANTEES

Pyronet's security model rests on three reinforcing mechanisms: cryptographic verification, economic staking, and on-chain escrow. Together they create a system where honest behavior is the only profitable strategy.

## // CRYPTOGRAPHIC VERIFICATION

Every workload produces a deterministic execution trace. Verifier nodes sample, replay, and confirm these traces independently. Divergence between claimed and actual output is detectable with high probability.

## // ECONOMIC STAKING

Providers post stake proportional to the value of jobs they accept. Stake is held in on-chain escrow throughout job execution. Detected dishonesty results in slashing — the dishonest party loses their stake, which is partially burned and partially redistributed to honest verifiers.

## // ESCROWED PAYMENT FLOWS

Renter funds are locked in protocol escrow at job initiation. They are released to providers only as work is verified, in per-second increments. Renters cannot underpay; providers cannot overcharge. The contract executes itself.

## // REPUTATION DECAY

Provider reputation grows slowly through honest behavior and decays rapidly upon any verification failure. High-stakes workloads are routed exclusively to high-reputation nodes. The result is a continuous selection pressure favoring reliability.

## // WORKLOAD ISOLATION

Each job runs in a hardware-isolated execution environment. Multi-tenant nodes cannot leak data across workloads, and renters cannot extract information about other jobs running on the same hardware.

## // OPEN-SOURCE TRANSPARENCY

The protocol's core components are open-source. Independent security researchers, auditors, and the broader community can review every line of code that mediates trust on the network.

# GOVERNANCE

NETWORK OWNERSHIP · PROTOCOL EVOLUTION

Pyronet is designed to be owned by its users. The protocol's evolution — its parameters, its upgrades, its treasury — is directed through transparent on-chain governance. Decision rights accrue to participants who have demonstrated genuine engagement with the network: providers who serve compute, renters who consume it, and stakers who secure it.

## // PROPOSAL LIFECYCLE

Any qualifying participant may submit a proposal. Proposals progress through three stages: **discussion**, during which the community refines the proposal through public deliberation; **signaling**, during which informal votes gauge community alignment; and **execution**, during which a formal on-chain vote determines whether the proposal is adopted. Adopted proposals are executed automatically by the protocol — there is no privileged role with the discretion to override or delay community decisions.

## // PARAMETERS UNDER GOVERNANCE

- Protocol fee rate and distribution between treasury, burn, and verifier rewards
- Minimum staking requirements for each provider tier
- Verification challenge frequency and sampling parameters
- Slashing severity for various classes of detected misbehavior
- Hardware benchmark requirements for new GPU classes
- Treasury allocation toward ecosystem grants and protocol development
- Network-wide reputation curve calibration

### // DECENTRALIZATION AS PROCESS

Governance decentralization is a process, not a state. Early protocol parameters are set by the founding contributors to ensure stability during the network's bootstrap phase. As the protocol matures, control over each parameter is progressively transferred to community governance. The trajectory is one-way: every step releases more authority to the network itself, and no step returns it.

# ROADMAP

THE PATH TO ONE MILLION DISTRIBUTED GPUS

Pyronet's development is organized into four distinct phases, each defined by a clear technical milestone rather than a fixed calendar date. Progression between phases is gated by the network achieving the relevant capability — not by speculation about when it might.

---

## PHASE I IGNITION

Core protocol launch on Solana mainnet. Initial provider onboarding focused on consumer-tier GPUs. Marketplace operational with full per-second billing. Proof-of-compute live in baseline configuration. First production workloads from launch partners.

---

## PHASE II EXPANSION

Enterprise-tier GPU onboarding including A100 and H100 classes. SOC2 compliance pathway for nodes serving regulated workloads. Geographic redundancy across three continents. Native SDKs for major AI frameworks. Provider tooling for fleet operators.

---

## PHASE III ACCELERATION

On-chain AI agent integration layer. Native compute primitives callable from any Solana program. Cross-protocol composability with leading on-chain stablecoin, DEX, and lending venues. Advanced verifier set with formal challenge games.

---

## PHASE IV AUTONOMY

Full progressive decentralization complete. Governance authority fully resident with network participants. Protocol-funded ecosystem grants operating at scale. Network capacity targeting one million distributed GPU nodes. Pyronet as the default compute layer of the open AI economy.

# CONCLUSION

THE FIRE IS LIT

The next decade of computing will not be built on the assumptions of the last one. Centralized hyperscalers cannot supply the world with compute on terms the world can accept. The hardware exists — distributed across millions of devices in millions of locations — but the coordination layer does not, and the trust mechanism does not, and the settlement layer does not. Until now.

Pyronet supplies what was missing. A permissionless marketplace where any GPU can earn. A cryptographic verification layer where every workload is provable. A per-second settlement engine where every party pays only for what they use. All of it composable, all of it on-chain, all of it owned by the network that creates the value.

---

**"THE DECENTRALIZED GPU ECONOMY."  
"FUELING THE FIRE OF OPEN COMPUTE."**

---

The protocol is live. The supply is forming. The demand is unprecedented. The settlement layer is the fastest the world has ever seen. The next era of compute will not be rented from a single vendor — it will be coordinated across an open network of independent operators, verified by mathematics, and settled in real time on-chain.

**This is the compute layer the AI era requires.  
This is the network that delivers it.  
This is Pyronet.**

