

SYNAPSE POWER AI LLC

SYNAPSE POWER

WHITEPAPER



Table of Content

Introduction	03	Security & Sustainability	11
Company Overview	04	Team & Expertise	12
Problem & Solutions	05	Partners	13
Our Technology	06	Community Share Pool	14
Node Leasing & Rental	10	Future Roadmap	15





Introduction

Synapse Power is a decentralized infrastructure platform delivering high-performance GPU computing for artificial intelligence (AI), machine learning (ML), and data-intensive workloads. As global demand for computational resources continues to grow exponentially, Synapse Power addresses emerging infrastructure limitations by providing a scalable, transparent, and high-availability GPU compute layer distributed across regulatory-compliant data centers.

Built by Synapse Power AI LLC, a United States-based technology company with cross-continental engineering teams, the platform is designed to function as a foundational layer for global AI scalability. Headquartered in the US, Synapse Power AI LLC focuses on infrastructure-grade reliability, energy-efficient data center deployment, and ongoing technical innovation in distributed compute orchestration.

As an infrastructure-native organization, Synapse Power AI LLC does not rely on speculative or purely financial operating models. Its core value proposition is centered on real-time workload execution and verifiable compute utilization. Through proprietary orchestration systems and a decentralized telemetry architecture, the platform provides uninterrupted access to GPU resources tailored to enterprise-grade workloads, including model training, inference, and large-scale data processing pipelines.

Synapse Power is committed to expanding access to advanced compute infrastructure for startups, research institutions, and enterprises worldwide, while maintaining strict standards for data protection, regulatory compliance, and sustainable infrastructure practices. Synapse Power is not positioned solely as a platform, but as a long-term initiative to decentralize global compute access, supported by a team with deep technical expertise and operational discipline required to sustain the next phase of AI infrastructure evolution.

Company Overview

By bridging the gap between AI demand and compute availability, Synapse Power AI LLC empowers developers, researchers, and enterprises with scalable access to high-performance GPU resources.



Vision

To be the global backbone of decentralized AI infrastructure, enabling equitable and secure access to GPU computing power for every innovator.

Mission

To deliver scalable, transparent, and sustainable GPU infrastructure that empowers the AI revolution through advanced technology, global reach, and trust-first engineering.

Core values



Integrity

Every compute cycle is verified and performance-audited. Transparency is not optional—it's built-in.



Resilience

We design for uptime, redundancy, and modular scalability, ensuring mission-critical operations remain uninterrupted.



Sustainable Innovation

Our commitment to green technology, efficient hardware lifecycles, and renewable energy sources is embedded in every deployment.



Problem & Solutions

The Problem:

The rise of artificial intelligence has generated an insatiable demand for computational power, particularly in the form of GPU acceleration. However, the current compute landscape is fragmented, centralized, and increasingly unsustainable:

- **Limited Access:** High-performance GPUs are often reserved for large enterprises or specific regions, leaving startups, researchers, and developers underserved.
- **Inefficient Infrastructure:** Legacy cloud platforms and centralized compute systems often lead to resource wastage, latency issues, and poor scalability.
- **Lack of Transparency:** Users often pay for allocated resources, not actual performance, with minimal visibility into workload distribution and efficiency.
- **Environmental Burden:** Data centers with inefficient cooling and outdated hardware significantly contribute to global energy consumption and carbon emissions.

The Solution:

Synapse Power introduces a decentralized GPU infrastructure platform that redefines how compute resources are accessed and utilized:

- **Decentralized Access:** Our globally distributed data centers in USA provide democratized access to high-performance GPUs with minimal latency and maximum reliability.
- **Telemetry-Driven Efficiency:** Each node is equipped with real-time monitoring tools to optimize resource allocation, manage thermal loads, and ensure efficient energy use.
- **Performance-Based Transparency:** Users interact with a dashboard that reflects real-time usage metrics, workload validation, and compute output for every session.
- **Sustainable Design:** From energy-efficient cooling systems to hardware lifecycle policies, Synapse Power is built for long-term sustainability without compromising performance.

Our Technology

At the core of Synapse Power lies a robust, modular, and intelligent technology stack purpose-built to deliver dependable, high-throughput compute infrastructure. Each component is engineered to ensure maximum performance, minimal latency, and seamless integration for AI developers and enterprise workloads.

1. **Compute Nodes:** Synapse Power nodes are built using industry-leading GPUs, including NVIDIA H100, A100, RTX 4090, and RTX 5080. Each node is optimized for large language models (LLMs), computer vision, scientific simulations, and inference tasks. All nodes operate in data centers with redundant power, Tier-3+ availability, and enterprise-grade network infrastructure.
2. **Node Virtualization & Isolation:** We employ container-based (Docker) and virtual machine (KVM) isolation to allow multi-tenant workloads per physical GPU while maintaining performance and security. This ensures users can scale applications with confidence.
3. **Smart Job Orchestration:** A telemetry-based job dispatcher dynamically assigns compute tasks based on real-time thermal loads, available TFLOPS, GPU memory, CUDA bandwidth, and workload priority. The orchestration engine is designed to self-optimize across clusters to eliminate bottlenecks and idle cycles.
4. **Telemetry & Real-Time Monitoring:** All GPU sessions are backed by transparent telemetry that includes usage metrics (TFLOPS, memory, power), thermal profiles, fan load balancing, uptime tracking, and session-based processing benchmarks. This ensures accountability, diagnostics, and historical traceability.
5. **Decentralized Scheduler Layer:** Unlike legacy infrastructure, our compute queue operates using a decentralized logic layer that ensures high availability and automatic failover. The architecture provides workload resiliency, redundant job replication, and autonomous scaling to meet rising demand without manual intervention.
6. **Scalable Deployment Architecture:** Our infrastructure supports horizontal node scaling through rapid deployment protocols, uniform configuration templates for efficient provisioning, and VPN-secured interconnectivity across regions for synchronized task management.
7. **API and Developer Tooling:** Synapse Power includes a comprehensive suite of APIs and command-line tools enabling developers to deploy, manage, and monitor compute tasks programmatically. These tools support pipeline automation and give access to detailed historical performance logs.
8. **Compliance & Infrastructure Governance:** All systems follow international standards such as ISO/IEC 27001 and GDPR. We implement jurisdictional compliance with data sovereignty regulations and conduct continuous auditing to align with enterprise IT governance frameworks.

Synapse Power is not a cloud mimicry—it is an entirely re-architected compute backbone tailored to the realities of today's AI evolution and tomorrow's decentralized infrastructure needs.

Token Information

Name : Synapse Power
Symbol : XNAP
Total Supply : 1,000,000,000 XNAP
Blockchain : Binance Smart Chain (BEP-20)
Decimals : 18
Type : Utility Token



XNAP Token Smart Contract Overview

The XNAP token is implemented as a fixed-supply ERC-20 compatible smart contract designed to support long-term ecosystem sustainability, transparent supply mechanics, and controlled integration with decentralized exchanges.

Token Parameters

```
string public constant name = "Synapse Power";
```

```
string public constant symbol = "XNAP";
```

```
uint8 public constant decimals = 18;
```

```
uint256 public constant MAX_SUPPLY = 1_000_000_000 * 1e18;
```

XNAP has a fixed total supply of 1,000,000,000 tokens, minted once at deployment. No minting functions exist, ensuring supply immutability.

Supply and Burn Mechanics

XNAP includes an on-chain burn mechanism that reduces circulating supply based on network activity. Burns are applied automatically depending on transaction type.

```
uint256 public constant TRANSFER_BURN_BP = 10; // 0.1%
```

```
uint256 public constant BUY_BURN_BP = 30; // 0.3%
```

```
uint256 public constant SELL_BURN_BP = 30; // 0.3%
```

- Standard transfers: 0.1% burn
- Liquidity buys: 0.3% burn
- Liquidity sells: 0.3% burn

To prevent excessive supply reduction, each transaction burn is capped.

```
uint256 public constant MAX_BURN_PER_TX = 5_000_000 * 1e18;
```

Burned tokens are permanently removed from total supply:

```
_totalSupply -= burnAmount;
```

```
emit Transfer(from, address(0), burnAmount);
```

XNAP Token Smart Contract Overview

Liquidity Detection and Trading Logic

The contract detects interaction with the PancakeSwap liquidity pool to determine transaction context.

```
if (from == pancakePair) {
    burnAmount = (value * BUY_BURN_BP) / BP_DENOM;
} else if (to == pancakePair) {
    burnAmount = (value * SELL_BURN_BP) / BP_DENOM;
} else {
    burnAmount = (value * TRANSFER_BURN_BP) / BP_DENOM;
}
```

This allows XNAP to apply predictable behavior for buys, sells, and wallet transfers once trading is live on PancakeSwap (BSC).

Access Control Model

XNAP uses a two-role control structure to reduce governance risk.

```
address public owner;
address public controller;
```

- Owner
 - Transfers ownership
 - Sets the controller (one-time)
- Controller
 - Configures the PancakeSwap liquidity pair
 - Manages excluded addresses

```
function setController(address _controller) external onlyOwner {
    require(controller == address(0), "controller already set");
    controller = _controller;
}
```

Once set, the controller cannot be replaced, reinforcing long-term operational discipline.

Security and Design Principles

- Fixed supply with no minting functions
- No rebasing or reflection mechanics
- Transparent on-chain burns
- Minimal logic surface to reduce attack vectors

```
function burn(uint256 amount) external {
    _balances[msg.sender] -= amount;
    _totalSupply -= amount;
}
```

All supply changes are verifiable on-chain.



Node Leasing & Rental

Synapse Power offers a flexible and transparent GPU leasing model that enables individuals, researchers, and enterprises to access powerful compute nodes on demand. This service is underpinned by a performance-centric infrastructure and structured to accommodate a wide range of workloads, from short AI inference sessions to prolonged model training cycles.

Our node configurations are based on real-world performance metrics such as TFLOPS output, CUDA capability, memory bandwidth, and thermal efficiency. Customers can choose from varying levels of compute packages tailored to different technical requirements, whether for training deep neural networks, running scientific simulations, or managing real-time inference tasks. Each package is defined by deterministic compute power, ensuring fairness in performance delivery.

A core differentiator of our rental system is the real-time usage metering. Unlike conventional flat-rate rentals, Synapse Power monitors actual compute consumption through advanced telemetry, giving users precise insights into how their sessions utilize allocated resources. This metering data is accessible through the Synapse Power dashboard, promoting full transparency and enabling users to optimize workload scheduling and cost efficiency.

Session types include both time-based and workload-based options. Time-based sessions provide fixed durations of compute access, while workload-based models adapt dynamically based on real-time usage, enabling efficient use of resources across varied compute intensities. Whether running burst compute for quick AI tasks or sustained jobs for extensive data processing, Synapse Power provides the flexibility to scale on demand.



Security & Sustainability

At Synapse Power, security and sustainability are foundational principles that guide every layer of our infrastructure. As we build the future of decentralized AI compute, we ensure that data integrity, operational safety, and environmental responsibility are embedded within our systems from the ground up.

From a security perspective, Synapse Power enforces a multi-layered defense architecture. All compute nodes are housed within physically secured, Tier-3+ compliant data centers equipped with biometric access control, 24/7 surveillance, and redundant power systems. Digital infrastructure is safeguarded by industry-standard firewalls, continuous threat monitoring, and real-time anomaly detection protocols. End-to-end encryption secures all job submissions, telemetry streams, and API interactions, while role-based access controls (RBAC) ensure that system privileges are tightly managed and auditable. Our infrastructure undergoes periodic penetration testing and third-party audits to validate resilience against evolving cybersecurity threats.

Operationally, Synapse Power follows a zero-trust framework that mandates verification at every access point. All nodes are independently monitored with real-time logging and behavior-based alerts that detect anomalies across workloads and user sessions. Backup and disaster recovery protocols ensure data durability and rapid fault recovery across all regions, minimizing the risk of compute disruptions.

On the sustainability front, Synapse Power is committed to energy-efficient operation and environmentally responsible practices. All active data centers leverage renewable energy sources wherever possible, including hydro and wind power in regions such as in the US. Our infrastructure incorporates advanced thermal management systems—dynamic fan control, airflow optimization, and component-level heat mapping—to minimize cooling loads and reduce total energy consumption.



Team & Expertise

The strength of Synapse Power lies in the multidisciplinary expertise and global vision of its founding team. Our leadership and engineering divisions are composed of professionals with deep backgrounds in high-performance computing, cloud architecture, artificial intelligence, network security, and data center operations. With decades of combined experience, our team brings both the technical acumen and strategic foresight necessary to build and operate a world-class GPU infrastructure platform.

At the core of our engineering group are infrastructure specialists who have designed scalable systems for major hyperscale providers, AI labs, and financial institutions. These individuals lead the deployment and optimization of GPU clusters, ensuring each node meets enterprise-grade standards for uptime, throughput, and workload isolation. Our platform architects are veterans in container orchestration, virtualization, and performance telemetry—ensuring that our compute layer remains lean, intelligent, and responsive to real-time demand.

On the operational front, our data center strategists and DevOps engineers collaborate to maintain regional uptime targets, manage redundancy protocols, and oversee thermal and energy efficiencies at each physical site. With a focus on both automation and precision, they ensure consistent performance delivery even under extreme loads or geographic failover scenarios.

Our leadership includes executives with backgrounds in global infrastructure deployment, venture-backed startups, and open-source innovation. United by a shared mission to decentralize compute access for the benefit of science and innovation, they guide Synapse Power's roadmap with both commercial insight and a deep commitment to technical excellence.



Partners

Synapse Power strategically collaborates with global leaders across multiple industries to ensure the integrity, scalability, and innovation of its decentralized compute infrastructure. These partnerships reinforce the technical foundation of the platform while enabling seamless integrations and access to high-quality hardware, facilities, and ecosystem extensions





Community Share Pool

As part of our commitment to transparency and decentralized participation, Synapse Power has developed the Synapse Power DApp—a decentralized application designed to host and manage the Community Share Pool. This digital layer empowers community members to track compute activity, monitor infrastructure performance, and engage with the broader Synapse Power ecosystem in a verifiable, on-chain environment. The Community Share Pool acts as a gateway for contributors and stakeholders to access real-time operational data from our infrastructure network. The DApp provides visibility into node uptime, workload distribution, telemetry logs, and active session metrics across all regions. By publishing performance data on-chain, Synapse Power ensures a tamper-proof and transparent reporting mechanism that reflects the true state of infrastructure operations.

Participants in the share pool can view their contribution records, historical uptime ratios, and scheduled access entitlements based on infrastructure usage and participation tiers. This system supports transparent community governance by making usage data accessible and audit-ready, encouraging long-term engagement and accountability. The Synapse Power DApp also serves as a foundation for future decentralized governance, enabling stakeholders to vote on infrastructure upgrades, expansion proposals, and resource allocation priorities. This design ensures that Synapse Power remains community-aligned while scaling globally as a neutral, trusted compute layer for AI.

By bridging transparent infrastructure metrics with decentralized community oversight, the Synapse Power DApp embodies the next generation of participatory AI infrastructure.

Roadmap



2025

Q2 Launches its GPU node leasing platform in the USA. This initial deployment integrates real-time telemetry into infrastructure monitoring and activates the internal orchestration engine with smart job routing and thermal management.

Q3 Synapse Power DApp goes live, offering public access to workload monitoring and Community Share Pool interaction. The platform expands to many countries, while CLI and API tools become available for developers.

Q4 Synapse Power introduces on-chain compute metrics and a governance layer for community participation, enabling community-driven node onboarding and enterprise dashboards for performance auditing. In parallel, Synapse expands its physical infrastructure with the launch of a Sweden-based data center to support regional scalability and reliability.


2026

Q1 Synapse will launch decentralized cloud storage, release a Telegram mini app running on Synapse infrastructure, and deploy enhanced monitoring with predictive analytics. A developer sandbox will support AI and blockchain testing, alongside enterprise onboarding for large-scale compute clients. In parallel, the XNAP mainnet will go live with XNAP launching on PancakeSwap.

Q2 Synapse will expand into the UAE, Singapore, Australia, and Canada, deploy green-certified compute hubs powered by renewable energy, launch a Node Optimization Framework to improve GPU efficiency, release Synapse SDK v1, and form strategic partnerships with AI labs and enterprise clients.

Q3 Synapse will develop ZKP modules for verifiable compute, launch an AI Blockchain for decentralized training and model sharing, and introduce a developer integration framework with SDKs and smart contracts. The platform will also enable secure decentralized data sharing, cross-chain job execution, and an on-chain AI model registry for verification.

Q4 Synapse will introduce an AI Workload Marketplace connecting researchers and enterprises, launch the Synapse Power Grants Program to support academic and open-source projects, and develop an AI Agent Network for distributed inference and collaboration. The ecosystem will be governed through an AI Governance DAO, alongside community incentive programs designed to reward contributors and developers.



Join Our Community

www.synapsepower.io

general@synapsepower.io