

IONET!!! Project

[Waht is io.net](#)



- **io.net** is a company that aims to build the world's largest **AI** compute **DePIN** (decentralized physical infrastructure network). They provide an enterprise-grade decentralized computing network that allows machine learning engineers to access distributed cloud clusters at a much lower cost compared to centralized services.

The company believes that compute power is the "**digital oil**" of this generation and is crucial for driving technological advancements. Their vision is to establish **IO** as the currency of compute, enabling access to compute resources and treating them as valuable assets.

Modern machine learning models often require parallel and distributed computing to optimize performance and handle large datasets. However, traditional cloud service providers have limited availability, lack hardware options, and can be expensive.

io.net addresses these challenges by aggregating **GPUs** from underutilized sources such as independent data centers, crypto miners, and other hardware networks. They create a Decentralized Physical Infrastructure Network (DePIN) that offers engineers accessible, customizable, cost-efficient, and easy-to-implement access to massive amounts of on-demand computing power.

io.net is 10x – 20x more efficient than current Cloud offerings

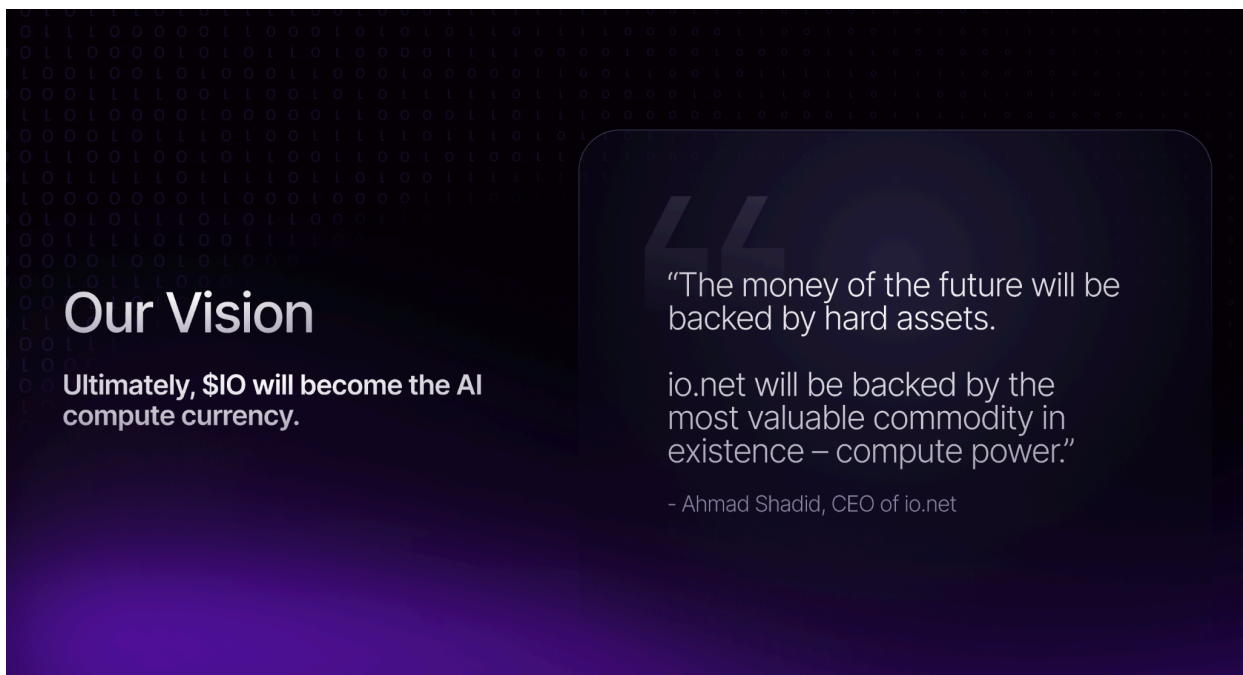
PROVIDER	GPU	PRICE / HOUR	IMAGE GENERATION TIME (SEC)	STABLEDIFFUSION IMAGES PER \$1
io.net	RTX 3090 x8	\$1.60	5s (8x Images)	28,800
	RTX 3090	\$0.20	5s	3,600
CoreWeave	A100	\$2.61	0.88s	1,567
Google Cloud	A100	\$3.67	0.88s	1,114
aws	A100	\$4.10	0.88s	997

With **io.net**, teams can scale their workloads across a network of **GPUs** with minimal adjustments. The system handles orchestration, fault tolerance, scaling, and supports various tasks such as preprocessing, distributed training, hyperparameter tuning, reinforcement learning, and model serving. It focuses on serving general-purpose computation for Python workloads, with a particular emphasis on **AI/ML** workloads.

"Emphasis on AI/ML workloads" means that io.net puts a significant focus on providing computing resources and services specifically tailored for artificial intelligence (AI) and machine learning (ML) tasks. AI and ML workloads often require intensive computational power, parallel processing, and specialized frameworks and libraries.

-The core functions of **io.net's** offering include batch inference and model serving, parallel training, parallel hyperparameter tuning, and reinforcement learning. They provide solutions that allow for efficient parallelization of tasks, leveraging distributed computing libraries and advanced techniques to optimize performance.

- io.net originated from the **Solana** Hackathon in February **2023** and the Solana **Austin** Hacker House.



Our Vision


Ultimately, \$IO will become the AI compute currency.

“The money of the future will be backed by hard assets.

io.net will be backed by the most valuable commodity in existence – compute power.”

- Ahmad Shadid, CEO of io.net

Ignition Rewards



Ignition Program and \$IO Token Airdrop

Solana-based decentralized network

The "Ignition" program is a key driver of io.net's recent growth. This initiative incentivizes users by rewarding them with \$IO tokens for contributing GPU power to the network. The airdrop of these tokens is scheduled for April 28th, 2024, offering a lucrative opportunity for participants. The program underscores io.net's commitment to decentralization and rewarding community contributions. The project's recent funding round, in which it raised \$30 million at a \$1 billion valuation, further demonstrates investor confidence in its potential.

io.net is poised to become a dominant player in the AI computing sector. Its decentralized model, competitive pricing, and focus on accessibility make it a compelling solution for businesses and researchers alike.

io.net Is Launching a Massive Airdrop on Solana in April

- io.net, a decentralized network on the Solana blockchain, addresses the growing AI computational demands.

- The network recently raised \$30 million at a \$1 billion valuation, showcasing investor confidence.
- To incentivize participation and community building, io.net launches its "Ignition" program, offering an airdrop of its native token (IO) on April 28th, 2024.
- The airdrop is attractive due to strong investor confidence, a defined end date, and timely market positioning.
- Participation in the airdrop involves becoming a GPU supplier and completing quests on the Galxe platform.
- To become a GPU supplier, set up your node, connect your GPU, and start supplying GPU power.
- Complete Galxe quests by connecting your wallet, exploring and completing tasks related to io.net, and tracking your progress.
- io.net aims to democratize access to computing resources and drive AI innovation with its unique offering and well-structured airdrop program.

The program consists of three separate pools of rewards:

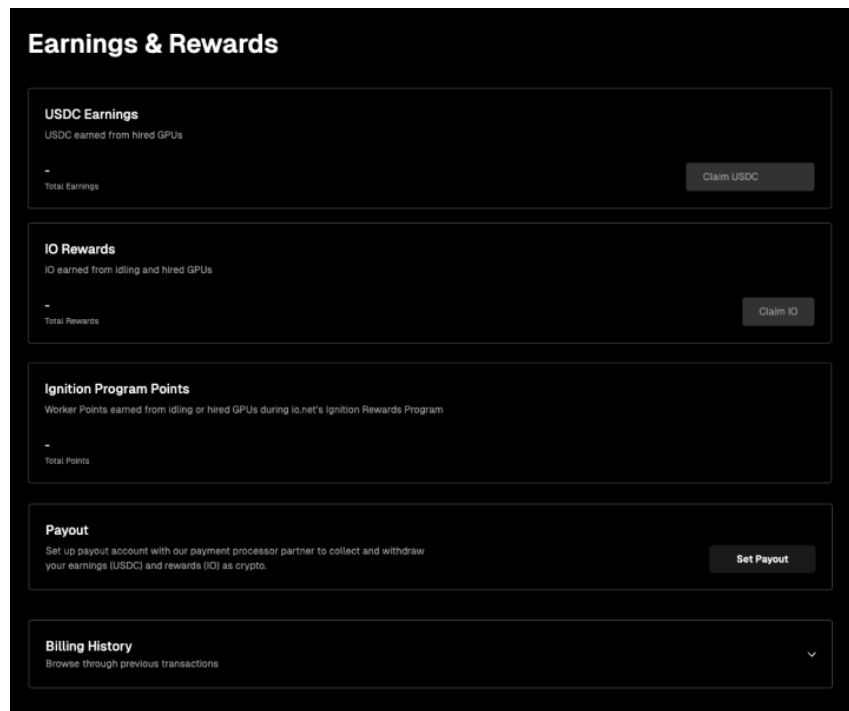
1. Worker Rewards: This pool calculates the score for each node based on their activity from **November 4, 2023**, to **April 25, 2024**. The program records the number of points accrued by each node during this period. At the end of the rewards period, the points earned by each user account will be summed and converted to rewards, which will be distributed proportionally based on their share of points.

2. Community Questing Rewards: Details about this pool are not provided.

3. Discord Role Rewards: Details about this pool are not provided.

The text also includes important notes and factors that contribute to the calculation of rewards:

- A valid Solana wallet must be associated with the user account to receive rewards.
- Self-service withdrawals for **USDC** earnings are currently under maintenance.
- Rewards calculations are specific to the **IO** rewards program and do not include **USDC** earnings rates for hired **GPUs**.



- Factors used for reward calculation include the ratio of job hours done, bandwidth, **GPU** model, and uptime. Bandwidth is classified into different tiers based on download and upload speeds, while **GPU** models are assigned specific scores.

Overall, the Ignition Rewards Program aims to incentivize participation and growth within the io.net Cloud and the IO Network by offering rewards to users based on their contributions and

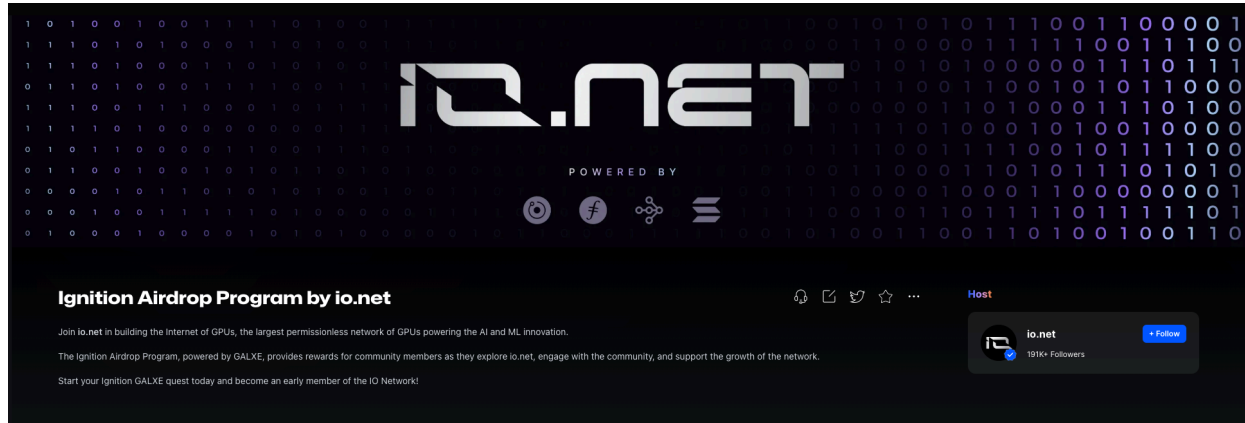
The rewards available in Galxe aim to provide users with reward points.

- **Uptime:** It refers to the total time that a node (server) has been active and available since **November 4, 2023**. The uptime increases when the node is active and available.

- **Viewing Points:** It refers to the ability to view the earned points. Users can view their points at the following link: <https://cloud.io.net/worker/earnings-and-rewards> around April 1, 2024. The current user interface will be updated to display USDC earnings, IO rewards, and Ignition Program points separately.

- **Community Questing Rewards:** It is a program managed by **Galxe** that allows users to earn points by completing tasks in the campaigns mentioned at the following link: <https://galxe.com/io.net/campaign/GCD5ot4oXPA>. **Galxe** rewards are separate from worker rewards and Discord role rewards.

- **Discord Role Rewards:** It is a program managed by the Galxe community team, which requires users to provide a valid Solana wallet address in Discord. Users will receive roles based on their contributions to rewards, content creation, and other activities.



- **Galxe Support:** If you encounter any issues with the Galxe platform, you can contact Galxe support for assistance.

PRODUCTS

IO Cloud



The IO Cloud platform is designed to deploy and manage decentralized GPU clusters on-demand. Its goal is to provide users with flexible and scalable access to GPU resources without the need for expensive hardware investments or infrastructure management. The platform utilizes a decentralized network of nodes called **IO workers**, which allows ML engineers and developers to have a cloud-like experience.

IO Cloud serves as the environment for running AI/ML applications and seamlessly integrates with the **IO-SDK**. It offers a comprehensive solution for scaling AI and Python applications, focusing on high performance and ease of use. The platform simplifies the deployment and management of **GPU/CPU** resources while providing limitless computing power.

The core component of IO Cloud is the Clusters, which are fully meshed self-healing GPUs. These clusters form a distributed and decentralized network capable of executing Python-based ML code. The system is easily scalable to accommodate varying demands.

Key features of IO Cloud include:

1. Becoming the world's largest and most cost-effective on-demand GPU cloud for AI/ML training and inference.
2. Smooth integration with the IO SDK, creating a unified, high-performance environment for AI projects.
3. Offering affordability with costs up to 90% cheaper per TFLOP (floating-point operations per second).
4. Global distribution of GPU resources, acting as a content delivery network (CDN) for ML serving and inference, reducing latency for end-users.

5. Powered by the RAY distributed computing Python framework, which has been used by OpenAI to train models like GPT-3 and GPT-4 across 300,000 servers. This provides an easy framework for scaling Python applications at any scale.

6. Future access to IO Models Store and advanced inference features such as serverless inferences, cloud gaming, and pixel streaming.

Overall, IO Cloud aims to provide a scalable and cost-effective solution for AI/ML workloads, leveraging decentralized GPU clusters and powerful distributed computing capabilities.

- IO Cloud Development Principles:

- **Ease of use:** Focus on a smooth user experience with intuitive and user-friendly control panels that provide users with instant access to the information they need. Setup is made easy, allowing users to get up and running in a matter of seconds.

- **Real-time:** Provide real-time interactions and updates through a specialized Python API layer specifically designed for the client portal.

- **Customizability:** Offer a wide range of customization options to ensure that users can tailor their environment precisely to their preferences. This means spending less time on system setup and more time on actual tasks. Whether it's Ray, Ludwig, Hugging Face, or any other stack, we've got you covered with just a few clicks.

- **Security:** The portal is supported by multiple layers of protection to safeguard user data and operations. A strong firewall, strict access control, and a customizable design are provided to minimize potential risks and isolate different functions to enhance security.

IO Worker



IO Worker is a platform designed to simplify **GPU** node operations management for users through a user-friendly web app. It aims to meet the increasing demand for AI model computations and utilize the unused computing power of decentralized **GPUs**. The platform offers features such as user account management, **real-time monitoring** of computing activities, temperature and power consumption tracking, installation assistance, wallet management, security measures, and profitability calculation.

The platform's development principles revolve around providing real-time updates, stringent security, streamlined operations, and **user-friendly interfaces**. Key principles include:

- 1. Full Financial Control:** Real-time access to earnings and financial metrics for suppliers, ensuring complete transparency and empowering informed decision-making.
- 2. Ease of Management:** Simplified GPU integration process to save time and effort for suppliers, allowing them to focus on their core competencies while the platform manages technical orchestration.
- 3. Efficient Resource Use:** Optimization of GPU utilization through smart task allocation and workload management, ensuring consistent revenue flow for suppliers.
- 4. Real-time Monitoring:** Provides miners with real-time monitoring of **GPU** usage, offering transparency and clarity on earnings and hardware utilization.
- 5. Security & Stability:** Implements stringent security protocols to protect hardware and earnings, ensuring consistent uptime and performance through active fault monitoring.

Overall, IO Worker strives to provide a comprehensive solution for managing GPU node operations efficiently, catering to the unique needs of users and suppliers in the AI landscape.

IO Explorer



IO Explorer provides users with comprehensive insights into the network's **GPU Cloud** operations, similar to how blockchain explorers offer transparency into blockchain transactions. It aims to empower users by offering complete visibility into network activities, vital statistics, data points, and rewards transactions while ensuring privacy and security.

App Walkthrough:

- **Explorer Homepage:** Provides supply insights, verified suppliers, active hardware quantity, and live market pricing, driven by demand.
- **Clusters Page:** Displays deployed clusters and real-time metrics on bookings. Users can access detailed views and track transactions.
- **Devices Page:** Shows connected devices and real-time metrics on bookings. Users can view detailed device information and track transactions.

IO Explorer Development Principles:

1. Intuitive Navigation: Designed with a user-centric approach for easy navigation, ensuring even complex tasks are simple. Clean UI and logical flow facilitate quick access to necessary information.

2. Real-time Cluster Monitoring: Offers instant insights into cluster status, health, and performance, crucial for dynamic **GPU** clusters.

3. Security & Control: Provides granular access control, allowing users to set permissions to ensure only authorized personnel can make changes or access sensitive data.

In summary, IO Explorer serves as a pivotal tool in the ecosystem, offering users a comprehensive view of GPU clusters with intuitive navigation, real-time monitoring, and robust security measures.

IO Architecture

IO.NET's Ecosystem Architecture prioritizes speed, effectiveness, and efficiency, ensuring sustainable, eco-friendly resource use. The architecture is multi-layered, cohesive, and built upon modern technologies for scalability, reliability, and robustness.

1. User Interface Layer: Provides a visual gateway for users with an intuitive design for easy navigation and interaction. Tech Stack includes ReactJS, Tailwind, web3.js, and zustand.

2. Security Layer: Ensures system integrity and safety with Firewall, Authentication Service, and Logging Service. Tech Stack includes pfSense, OAuth, JWT, ELK Stack, and Graylog.

3. API Layer: Acts as a communication bridge with Public API, Private APIs, and Internal APIs for various functionalities. Tech Stack includes FastAPI, Python, GraphQL, RESTful services, gunicorn, and solana.

4. Backend Layer: Manages Providers, Cluster/GPU operations, Customer interactions, Fault Monitoring, Analytics, Billing/Usage Monitoring, and Autoscaling. Tech Stack includes FastAPI, Python, Node.js, Flask, solana, and IO-SDK (a fork of Ray 2.3.0).

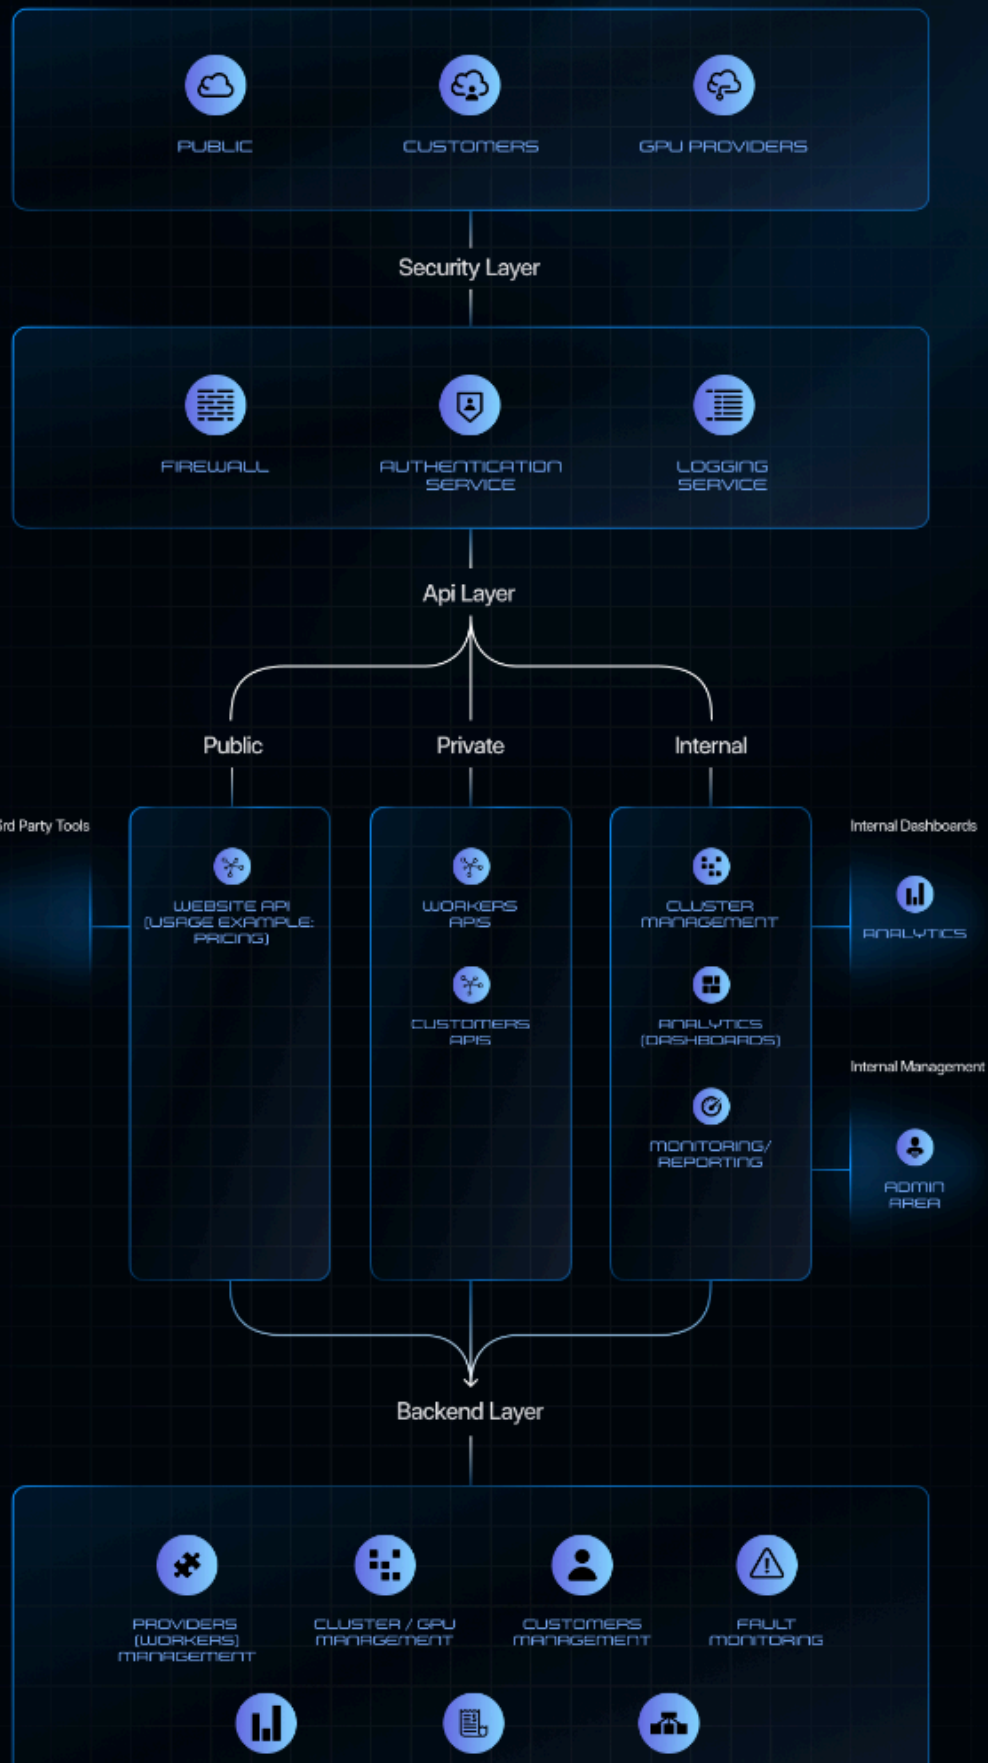
5. Database Layer: Stores structured data and caching for temporary, frequently accessed data. Tech Stack includes Postgres for Main storage and Redis for Caching.

6. Message Broker/Task Layer: Orchestrates asynchronous communications and task management for smooth data flow and efficient task execution. Tech Stack includes RabbitMQ for Message Broker and Celery for Task Management.

7. Infrastructure Layer: Houses GPU Pool, manages deployments, executions, ML tasks, and provides Data Storage solutions. Tech Stack includes Kubernetes, Prefect, Apache Airflow, Ray, Ludwig, Pytorch, Keras, TensorFlow, Pandas, Amazon S3, Hadoop HDFS, Docker, Grafana, Datadog, Prometheus, NVIDIA DCGM.

8. IO-SDK: Specialized fork of Ray, enabling dynamic task execution, rapid data sharing, and dynamic auto-scaling. Compatible with leading ML frameworks like PyTorch and TensorFlow, ensuring scalability and performance.

These layers, powered by the mentioned tech stacks, form a robust and scalable architecture for the IO.NET Portal, meeting the demands of modern users and remaining future-proof.



IO Tunnels

IO Tunnels utilize reverse tunnel technology to bypass firewall and NAT restrictions, enabling secure access to remote resources. This document explains the concept of reverse tunnels, their functionality, and how io.net utilizes them to simplify access for engineers to AntMiners Clusters.

Understanding Reverse Tunnels:

- Reverse tunnels establish a secure connection from a client to a remote server by opening an inbound connection on the server side.
- This method allows engineers to access remote resources behind NAT routers and firewalls without complex network configurations.

How Reverse Tunnels Work:


- The remote server (IO Worker) initiates a connection to the client (engineer's machine) through an intermediate server (io.net server).
- The io.net server listens for incoming connections from both the client and the remote server.
- Data exchange between the client and the remote server occurs through the tunnel, simulating a direct connection.

IO.net Implementation:

- **IO Workers** establish a connection to the io.net server, creating a reverse tunnel.
- The engineer's machine connects to the io.net server as an intermediary, routing traffic between the engineer and the IO Worker securely.
- This setup allows engineers to manage IO Workers securely without complex network configurations or firewall/NAT issues.

Benefits of Reverse Tunnels:

- **Simplified Access:** Engineers can access IO Workers without network restrictions, port forwarding, or VPNs.
- **Enhanced Security:** Reverse tunnels provide a secure communication channel, ensuring data privacy.
- **Scalability:** IO.net can manage multiple IO Workers simultaneously, enhancing efficiency.
- **Flexibility:** Reverse tunnels work across various platforms, ensuring compatibility with different operating systems and environments.



IO Network

1. Mesh VPN Networks:

- Mesh VPN networks connect nodes in a decentralized manner, allowing direct connections between each node.
- Advantages include robustness, scalability, lower latency, and better load distribution compared to traditional hub-and-spoke VPN architectures.

2. Implementation of io.net Network:

- **Enhanced Performance:** Minimizes latency, optimizing application performance and user experience.
- **Improved Resilience:** Decentralized architecture ensures operational continuity despite individual node failures.
- **Seamless Scalability:** Can accommodate more nodes as the user base expands, ensuring consistent performance.
- **Distributed Computing:** Facilitates efficient distributed computing by allowing direct connections between nodes.

3. Decentralized Architecture and Privacy:

- **No Single Point of Failure:** Absence of central concentrator enhances security and operational continuity.
- **Anonymity:** Data travels along multiple paths, making it difficult to trace its origin or destination.

-Traffic Obfuscation: Techniques like packet padding enhance privacy by obfuscating traffic patterns.

4. Network Access Control and Monitoring:

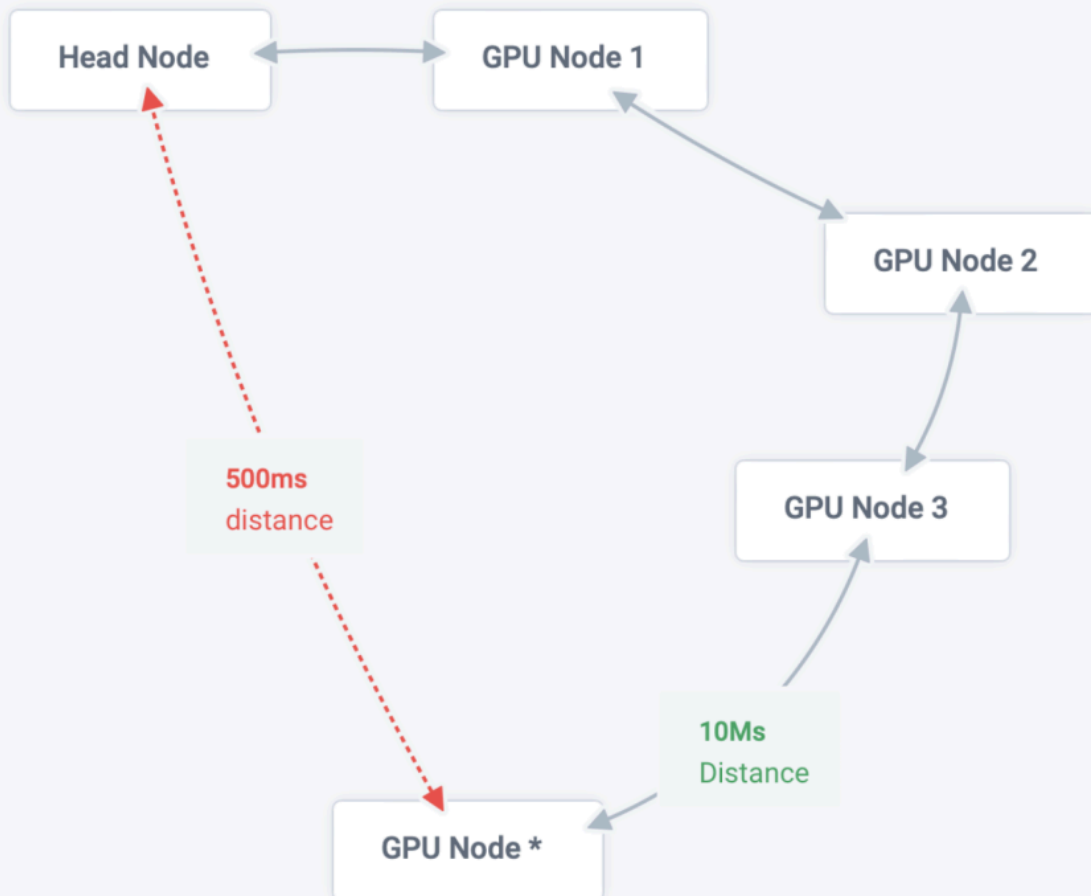
- **Access Control Lists (ACLs):** Restrict communication between nodes to ensure data security and privacy.
- **Regular Auditing and Logging:** Maintain logs of network activities to identify and address potential vulnerabilities or breaches.

Overall, IO Network leverages mesh VPN technology to create a highly efficient, resilient, and scalable networking backend for io.net, ensuring optimal performance, security, and privacy for users.

Cluster Mesh Network - Data traffic

All to All TLS enveloped tcp traffic

Means when {Head Node} want to communicate a computing task to {GPU Node 3} traffic will be transmitted via {Node 1} then {Node 2} if the travel will be faster based on the sum of total ping, if hops were made.

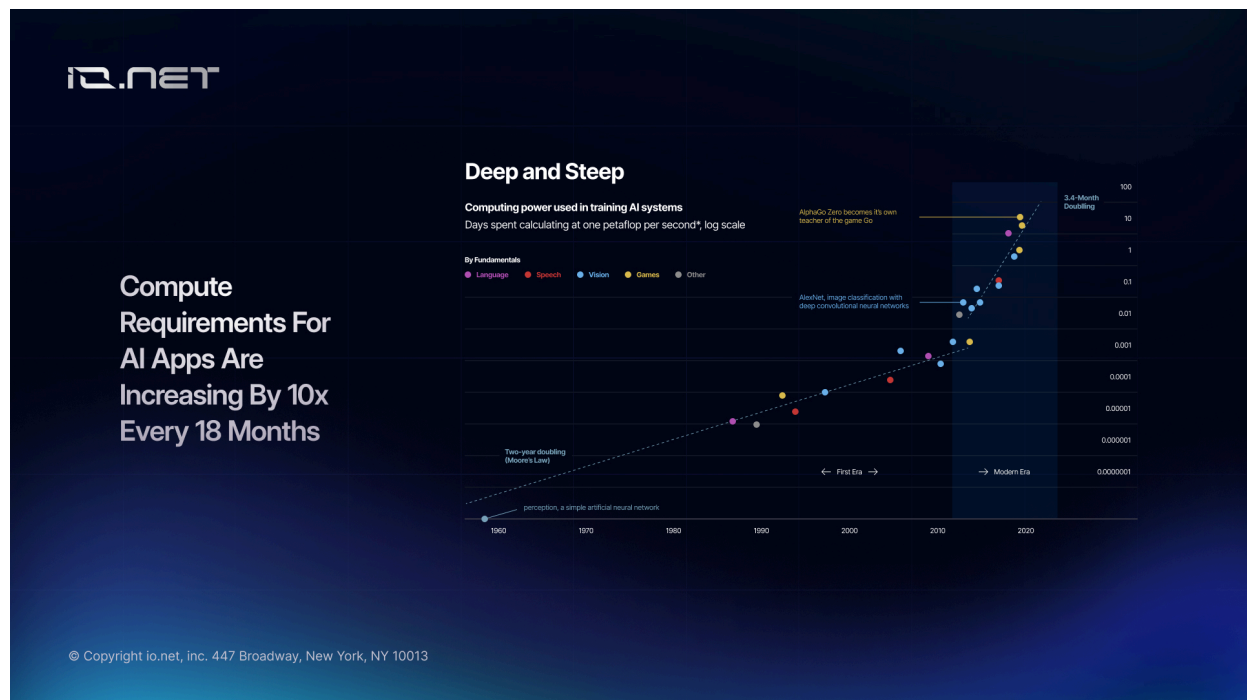


COMPANY

IO Inception

1. Before June 2022:

- io.net focused on developing institutional-grade quantitative trading systems for US stock and cryptocurrency markets.
- Challenges included constructing infrastructure for a robust backend trading system with significant computational power.



- Trading strategies bordered on high-frequency trading (HFT), requiring real-time monitoring of tick data for over 1,000 stocks and 150 cryptocurrencies.
- System dynamically backtested and adjusted algorithm parameters for each asset in real-time, maintaining low latency for order execution across various platforms.

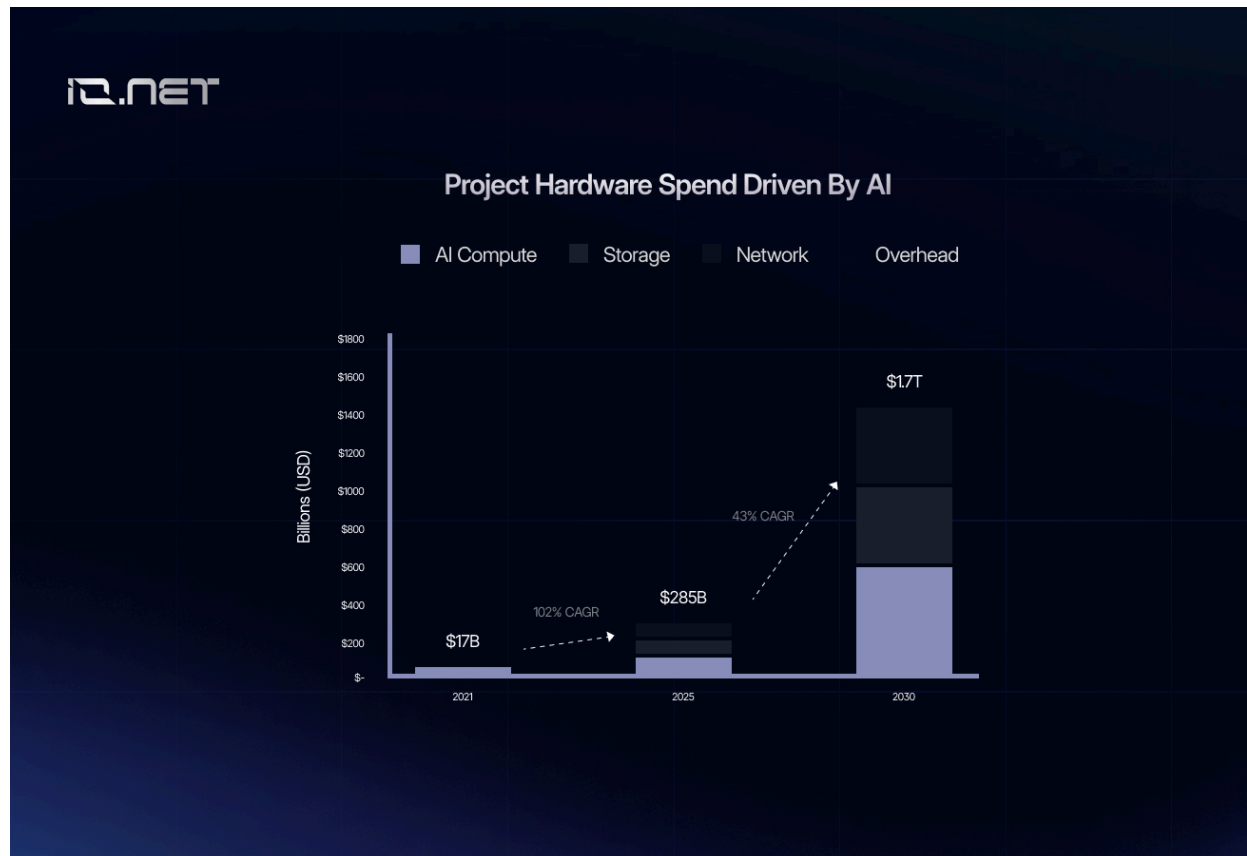
2. Discovery of Ray.io:

- Ray.io, an open-source library used by OpenAI for distributed **GPT-3/4 training** across CPUs and GPUs, revolutionized infrastructure management.
- Integration of Ray into backend reduced build time drastically, from over six months to less than 60 days.

Overall, the integration of Ray.io played a crucial role in optimizing io.net's backend infrastructure, enabling them to meet their computational demands more effectively and reducing the time and resources required for development.

3. Price Issues:

- Faced challenges with high cost of GPU on-demand cloud providers, hindering deployment of GPU and CPU workers for substantial computing power.
- High GPU prices, such as **NVIDIA A100** card, posed significant financial burden.



4. Need for Distributed Computing:

- Distributed computing increasingly necessary due to exponential growth in compute demands, especially in machine learning (ML) applications.
- Moore's Law's decline and rising ML application demands drive need for distributed applications to bridge performance gap.
- Deep learning application demands rapidly increasing, including training, tuning, and simulations, surpassing Moore's Law's growth rate.

- Distributed computing essential to bridge gap between application demands and hardware capabilities in big data and AI domains.

At io.net, development of innovative tools and distributed systems like Ray aims to facilitate transition into this new era of distributed computing.

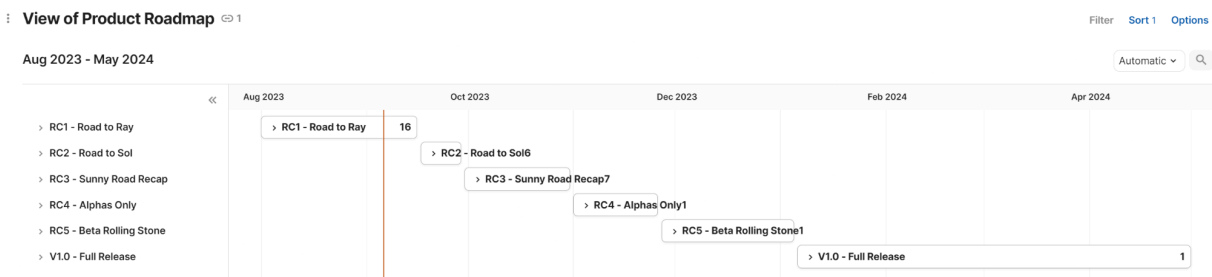
Roadmap

- **RC1 - Road to Ray:**
 - Boosts core functionality on the Client side, including setting up and deploying the cluster, deploying required jobs, viewing performance details, and downloading results.
 - Introduces basic functionalities for Suppliers, such as setting up machine workers within the Portal and connecting them with the Portal instance.
 - Includes middle layer orchestration to facilitate these functionalities.
- **RC2 - Road to BreakPoint:**
 - Enhances Portal capabilities for Clients and Suppliers with detailed user management, additional views, and options for units provided/consumed.
 - Adds payment and payout options through Solana payment gateway.
 - Introduces the Explorer concept for ecosystem monitoring.
- **RC3 - Sunny Road Recap:**
 - Uses feedback from RC1 and RC2 to improve overall ecosystem functionalities.
 - Potential enhancements include adding support for selected jobs, improving machine miner parametrization, and enhancing threading through the orchestration layer.
 - Includes overall Portal updates based on feedback.
- **RC4 - Alphas Only:**
 - Focuses on addressing feedback from limited access partners and community to improve ecosystem stability.
 - Expands available machine miner builds for various use cases and includes standard images for data centers.
 - Enhances Explorer concept and functionalities for better ecosystem network management.
 - Implements customer support features.
- **RC5 - Beta Rolling Stone:**

- Constantly tweaks platform mechanisms based on limited invites bringing in fresh users.
- Optimizes supported types of distributed computing power and introduces categorization in the Portal.
- Implements in-app advisor engines and auto-correction options.
- **V1.0 - Full Release:**
 - Shifts to standard release architecture with constant updates and customer-influenced roadmaps.
 - Works on decentralization of the Ecosystem, making it self-hosted and self-replicated.

The roadmap aims to progressively enhance functionality, stability, and user experience, leading to the full release of the platform.

Expected timeline for our RC's candidates bellow:



GUIDING PRINCIPLES

1. Security:

Modular Infrastructure: Utilizes distinct, modular blocks to compartmentalize responsibilities and minimize risks.

Robust Firewall and Authentication Layer: Implements strong firewall and authentication measures to control access and detect unauthorized activities.

Extensive Logging Layers: Logs every action to track and investigate any suspicious activities for manual verification.

2. High Load:

Modular Design: Breaks down the system logic into smaller, manageable pieces for efficient scaling.

Operational Queuing: Implements queuing to prevent overload on any single component and ensure smooth data flow, especially during peak loads.

Comprehensive Monitoring and Alerting: Maintains detailed logs and monitors system health, performance, and resource utilization in real-time, triggering alerts for swift action and mitigation.

3. Billing:

Automated Cost Monitoring/Alerting: Integrates real-time cost monitoring tools to track GPU usage and associated costs, preventing overuse and staying within budget.

Extensive Notification Capabilities: Implements a comprehensive notification system to alert users and administrators through various channels based on message urgency and nature.

Forecasting: Utilizes predictive analytics tools to forecast future GPU usage and costs, enabling better budgeting, resource allocation, and planning.

10 Backend Principles:

1. Modularity & Specialized Layers:

Dedicated Functionalities: Architects backend with distinct API layers, ensuring each functionality has its dedicated layer for enhanced security, scalability, and maintainability.

Fail-Safe Design: Segregates services interacting with third-party providers to introduce an additional layer of fail-safe design, minimizing disruptions.

2. Task Management & Queuing:

Seamless Operations: Integrates RabbitMQ and Celery for robust task management and queuing, facilitating efficient task processing and preventing system overloads.

3. Monitoring & Autoscaling:

Proactive Management: Leverages Python, Postgres, and Redis for scalable and efficient backend operations. Implements specialized monitoring layers for billing and autoscaling to adjust resources in real-time based on demand.

4. Security & Reliability:

Shielded Operations: Fortifies backend with multiple protective measures to ensure data integrity and operations are never compromised. The modular design minimizes risks by isolating different functionalities.

Supported Devices

GPU Support:

- Lists various NVIDIA GPU models supported on the IO Network, including GeForce GTX, GeForce RTX series, and Tesla series.
- Includes specific models like GeForce RTX 3050, RTX 3060, RTX 3080, RTX 3090, and various Tesla models.
- Mentions that the list may be expanded in the future to support new models from different manufacturers.

- CPU Support:

- Lists supported CPU models from Apple and AMD for the IO Network.
- Includes Apple M1, M2, M3 series, as well as AMD Ryzen Threadripper and Threadripper PRO series.
- Indicates that new models from different manufacturers are actively being supported and users should stay tuned for updates.

Additional Information:

- Provides links for more information on Ignition Rewards rates and Earnings rates.
- Mentions that the probability of device hiring and rewards are multiplied by each 1GB invested in Download/Upload.

New Partnership

[Nimble Network and io.net](#)



A strategic partnership has been formed between **Nimble Network** and **io.net** to transform the landscape of artificial intelligence development.

Nimble Network offers a decentralized framework facilitating the efficient and scalable integration of AI models and data, while **io.net** provides access to distributed GPU compute resources.

- The partnership aims to attract developers and scientists through joint marketing efforts and to provide **Nimble's tools on io.net's platform**. Initial trials will optimize integration, followed by full tool integration and leveraging of independent resources to enhance development capabilities.

Important

- io.net, a decentralized network on the Solana blockchain, addresses the growing AI computational demands.

- The network recently raised \$30 million at a \$1 billion valuation, showcasing investor confidence.

- To incentivize participation and community building, io.net launches its "Ignition" program, offering an airdrop of its native token (IO) on April 28th, 2024.

- The airdrop is attractive due to strong investor confidence, a defined end date, and timely market positioning.

- Participation in the airdrop involves becoming a GPU supplier and completing quests on the Galxe platform.

- To become a GPU supplier, set up your node, connect your GPU, and start supplying GPU power.

- Complete Galxe quests by connecting your wallet, exploring and completing tasks related to io.net, and tracking your progress.

- io.net aims to democratize access to computing resources and drive AI innovation with its unique offering and well-structured airdrop program.

AMA News and updates 10.04.2024

- Currently, there are more than 500,000 GPUs on the network, the total cost of the infrastructure is more than \$ 3 billion.

- Galxe issue have been fixed and you no longer need to join the general chat to complete the task.
- Uptime issues will be fixed soon.
- Major partnerships are in the process of development, which will be announced before TGE.
- Tokenomics is designed in such a way as to benefit everyone, including the investor, the team and the community
- Tokenomics is expected to be presented in 5 days.
- Public events were held last week, including a meeting of the Vietnamese community and an event in Turkey.
- More public events are planned in the near future.

Regarding the team,

there is ambiguity in many people because they have not disclosed their previous work except for a few **62 people**. The **CEO** has previous work

Join @shadid_io at #SuperAI Singapore. Ahmad is the Founder and CEO of @ionet, providing ML engineers with scalable clusters of geo-distributed GPUs at a lower cost than centralised cloud providers. Meet the teams building AI infrastructure this June.

pic.twitter.com/4YyfoTmyBn— SuperAI (@superai_conf) April 9, 2024

Ahmad Shadid

Founder and CEO
io.net



SUPERAI

SINGAPORE
5 - 6 JUNE 2024

@superai_conf | twitter.com | heute um 15:34 Uhr

I asked the support to ask and I will receive full information about the team tomorrow and add it