# **B** FarmGPU

### Decentralized. Sustainable. Secure.





### **Executive Summary**

• FarmGPU is a neocloud provider of GPU compute for startups and enterprises

Own High-Demand GPUs, Earn Sustainable Returns

- Data Center provider and launch partner for Silicon Network by Berkeley Compute, Inc. offering tokenized GPUs
- Initial offering with RTX 4090 inference optimized rack mounted workstation





## The Problem & Market Opportunity

#### The Problem



High costs:

Al developers pay up to 10x more for GPU compute at major CSPs

Vendor	H100 \$/hr
FarmGPU	\$3
AWS	\$12.25

Vendor Lock-in and Centralized Infrastructure: Traditional cloud providers force customers into closed ecosystems, restricting flexibility and increasing costs

#### Limited Access to Entry-Level GPUs:

Mainstream cloud providers often overlook the needs of smaller players and researchers who require affordable, entry-level GPU resources.

#### Market Opportunity



 Explosive Growth in AI Compute Demand:
On-demand GPU compute is fastest growing cloud segment



Decentralized & Neocloud Advantage: Market pricing, decentralized compute, lean operations



# Own High-Demand GPUs, Earn Sustainable Returns:

Enabling a new asset class for GPU ownership and strong, passive income



>\_

### FarmGPU Cloud Platform

- AI Data Center Blueprint: Transforms existing data centers into optimized AI facilities (45-130kW racks, 800G networking, liquid & immersion cooling).
- On-Demand Compute Platform: Instant access to powerful NVIDIA GPUs – 70% lower cost than hyperscalers – on RunPod Secure Cloud.
- Storage for AI: Accelerates training with GPU-direct storage & BlueField DPUs; enables scalable inference via vectorDB integration.





### **GPU Revenue Performance**

#### Utilization

- Platform: RunPod Secure Cloud
- GPU: RTX 4090
- Ave rate \$0.53/hr/GPU (+ local storage)
- \$383/GPU/month, \$4599/GPU/year
- Outlook strong demand for 2025 even with 5090 transition, used prices increasing



#### 肖 GPU Cloud Revenue Model

Calculate the total cost of ownership, expected revenue, and return on investment

	GPU Model		Consumer	Overview  Overvider Cos	sts 🛛 🛢 GPU Stats		
	RTX 4090 (\$2,600)						
	Server Cost (\$)	GPUs per Server		9 Initial Investment		Annual Revenue	
	6000	4					
Iry model here!	Total initial investment: \$4,100 per GPU	J (\$2,600 + \$1,500 server	cost)	\$4,100		\$1,720.037	
calculator.farmgpu.net/	Contract Duration (Years)						
	3 Years					S Payback Period	
	🛢 Rental Type		Secure ~	🕒 Total ROI 🚯	~ IRR ()	26.6 months	
				58.5%	34.0%	0	36 mont
	4 Usage Distribution 6	•					
	Idle		4%				
	Spot		4%	Cumulative Revenue			
	On-Demand		92%	\$6,000			
				\$4 500			
	Effective Hourly Rate:		\$0.649/hr			Ini	
				\$3,000			
	8 Revenue Sharing 6			\$1,500			
	Platform Fee		20%	\$0			
	Owner Share (of remaining	)	50%	Omo 2mo 4mo	6mo 8mo 10mo 12mo 14mo 1	16mo 18mo 20mo 22mo 24mo 26mo 28mo 30mo 3	32mo 34mo 36mo
	Provider Share (of remainin	ng)	50%				
	Annual Owner Revenue:		\$1720/vear				

#### E FarmGPU

### Server Configurations





Gigabyte G493-ZB2-AAP1 Dual AMD EPYC 9654 96-Core Processor (192 vCPU) 8 x H100 NVL (PCIe 5.0 x16) or 8x 6000 Ada (PCIe 4.0 x16)

24x 64GB DDR5 at 4800MT/s

#### MSI G4101-01

AMD EPYC 9254 24-Core Processor (48 vCPU)

4 x GPU (PCIe 4.0 x16)

12x DDR5 16GB at 4800MT/s



#### Lenovo H200 HGX



SMCI B200 HGX

# Our Data Center



- FarmGPU main site is hosted in Conscious Capital Data Center in Rancho Cordova (next to Solidigm campus)
- Tier 3 data center, fully redundant power
- 99.99% availability
- Equipped with liquid cooling, immersion cooling
- High power racks ready for latest GPUs
- 50k square ft, 7MW





www.runpod.io

Globally <mark>distributed</mark> GPU cloud for your AI workloads

Powerful & Cost-Effective GPUs for Every Workload

Spin up a GPU pod in seconds

Choose from 50+ templates ready outof-the-box, or bring your own custom container.

FarmGPU is RunPod Secure Cloud partner for us-west region

### **ON-DEMAND COMPUTE**

	cure Cloud 🗸 🗮 Ne	etwork Volume 🗸 🕥 Any 🕥	~ 註
Filter GPUs by VRAM			
Any 16 24		140	
AMD			
MI200Y	\$2.00 <i>/</i> /		
MISOUX	<b>\$3.99/hr</b> 3.39/hr		
192 GB VRAM	8 max		
VVIDIA Latest Gen			
NVIDIA Latest Gen -	<b>\$2.99/hr</b> 2.49/hr	RTX 4090	<b>\$0.69/h</b> r 0.59/hr
NVIDIA Latest Gen H100 SXM 80 GB VRAM	<b>\$2.99/hr</b> 2.49/hr 8 max	RTX 4090 24 GB VRAM	<b>\$0.69/h</b> i 0.59/hi 8 max
Latest Gen H100 SXM 80 GB VRAM 125 GB RAM + 16 vCPU	<b>\$2.99/hr</b> 2.49/hr 8 max High	RTX 4090 24 GB VRAM 30 GB RAM - 8 vCPU	<b>\$0.69/h</b> r 0.59/hr 8 max Higt
Latest Gen H100 SXM 80 GB VRAM 125 GB RAM - 16 vCPU L40S	\$2.99/hr 2.49/hr 8 max High \$1.03/hr 0.88/hr	RTX 4090 24 GB VRAM 30 GB RAM - 8 vCPU RTX 6000 Ada	\$0.69/hr 0.59/hr 8 max Higt \$1.03/hr 0.88/hr
Latest Gen H100 SXM 80 GB VRAM 125 GB RAM + 16 vCPU L40S 48 GB VRAM	\$2.99/hr 2.49/hr 8 max High \$1.03/hr 0.88/hr 8 max	RTX 4090 24 GB VRAM 30 GB RAM - 8 vCPU RTX 6000 Ada 48 GB VRAM	\$0.69/hn 0.59/hi 8 max High \$103/hn 0.88/hi 8 max
VVIDIA     Letest Gen       H100 SXM     80 GB VRAM       80 GB VRAM     16 vCPU       L40S     48 GB VRAM       48 GB VRAM     18 vCPU	\$2.99/hr 2.49/hr 8 max High \$1.03/hr 0.88/hr 8 max Medium	RTX 4090       24 GB VRAM       30 GB RAM + 8 vCPU       RTX 6000 Ada       48 GB VRAM       62 GB RAM + 14 vCPU	\$0.69/hn 0.59/hi 8 max Higt \$1.03/hn 0.68/hi 8 max Medium
Latest Gen H100 SXM 80 GB VRAM 125 GB RAM + 16 vCPU L40S 48 GB VRAM 62 GB RAM + 16 vCPU L4	\$2.99/hr 2.49/hr 8 max High \$1.03/hr 0.88/hr 8 max Medium \$0.43/hr 0.37/hr	RTX 4090 24 GB VRAM 30 GB RAM - 8 vCPU RTX 6000 Ada 48 GB VRAM 62 GB RAM - 14 vCPU RTX 2000 Ada	\$0.69/hn 0.59/hi 8 max Higt \$1.03/hh 8 max 8 max Medium \$0.28/hi 0.21/hi
VVIDIA     Letest Gen       H100 SXM     80 GB VRAM       80 GB VRAM     16 vCPU       L40S     48 GB VRAM       62 GB RAM - 16 vCPU     14       24 GB VRAM     10 vCPU	\$2.99/hr 2.49/hr 8 max High \$1.03/hr 0.88/hr 8 max Medium \$0.43/hr 0.37/hr 10 max	RTX 4090       24 GB VRAM       30 GB RAM - 8 vCPU       RTX 6000 Ada       48 GB VRAM       82 GB RAM - 14 vCPU       RTX 2000 Ada       16 GB VRAM	\$0.69/hn 0.59/hr 8 max Higt \$1.03/hr 0.88/hr 8 max Medium \$0.28/hr 0.21/hr 0.21/hr

# Tokenized Ownership

# **Silicon**

- 1. NFT-Backed GPUs: Partnering with Silicon Protocol to tokenize GPU assets on the blockchain.
- 2. Lower Entry Barrier: Acquire fractional ownership—no need to buy entire GPUs outright.
- 3. Enhanced Liquidity: Easily trade or sell your GPU tokens on secondary markets.
- 4. Portfolio Diversification: Blend a cutting-edge tech investment with stable, revenue-generating infrastructure.

# The Next Generation of AI Infrastructure Ownership

Own High-Demand GPUs, Earn Sustainable Returns

High-Growth Market: AI workloads driving demand



Sustainable Returns: up to 30% IRR, 40% annual ROI



Transparent Model: Tokenized ownership on blockchain



Turnkey Management: We handle data center operations, you reap the benefits.



Sustainable Operations: Liquid and immersion cooling

# How It Works:

# Simple, Secure, and Profitable from Day One

			Tumor 0
FarmGPU ©			
	⊖ GPUs	ළ GPUs Rented	\$ Total profit / day
128 96 RTX 4090	1024 768 RTX 4090	959 703 RTX 4090	\$5,645.48 \$3,700.22 RTX 4090
¢1 011 57	3 02		
Total earnings	5.02	Jan 23 - Dec 24 This week	This month All time
	02 Sept 2023 \$56,678		
			Sept 2024 Dec 2024
Summary Servers Fir			
			Filter # =

- 1. Purchase GPU NFT: One-click purchase.
- 2. We Host & Manage: We install, configure, and maintain your GPUs in our tier-3 data center.
- 3. Earn Ongoing Revenue: GPUs are rented out for AI workloads; you receive consistent income (tracked on-chain).
- 4. Monitor in Real Time: Our dashboard shows performance metrics, uptime, and revenue anytime, anywhere.







# MEET THE TEAM







JM Hands CEO Lloyd Ross COO Pat Harrington CBO



Gary Lamb iBridge Cloud Technologies, Networking



Nilesh Shah Emerging technology

#### Advisors



Steven Kaplan Author The ROI Story, ex-Nutanix

## **Rating GPU Clouds**

- Security
- LifeCycle and Technical Expertise
- Slurm and Kubernetes
- Storage
- NCCL/RCCL Networking Performance
- Reliability and Service Level Agreements (SLAs)
- Automated Active and Passive Health Checks and Monitoring
- Consumption Models, Price Per Value, and Availability
- Technical Partnerships



Source: <u>SemiAnalysis</u>

## Energy will be limiting factor for AI

With AI models requiring ever greater investment to improve "there is a capital question of at what point it stops being worth it to put the capital in," Zuckerberg said. "But I actually think that, before we run into that, we're going to run into energy constraints."

"Tokens per watt per dollar"—the sweet spot where energy, compute power, and intelligence meet—will be a game-changing formula for driving GDP growth said Satya Nadella

Decentalized, small regional data centers, and edge - O- compute will be required for AI to continue growing



# GAMING GPU = ENTRY LEVEL AI GPU

GPU	H100 NVL PCIE	RTX 4090
Price	\$26,400	\$1,600
FP32 TFLOPS	60 TFLOPS	82.58 TFLOPS
Power	350-400W	450W
Memory Size / Type	94GB HBM3	24GB GDDR6X
Memory Bandwidth	3.94 TB/s	1.01 TB/s



