GPU capacity partnership and pilot

Hourly compute with flexible onboarding from existing ML tools

> **We are designing a service that connects unused GPU capacity with ML teams needing affordable short term compute, with an execution layer and onboarding from their current MLOps tools. We are looking for an infrastructure partner and a real customer workload for a first pilot.**

Suppliers and customers

GPU providers, cluster owners and companies with reserved capacity may have paid capacity available between jobs or outside peak periods. They could offer approved windows to external teams and recover value from hours that would otherwise remain unused.

Customers would be ML teams with an existing training, fine-tuning or inference workload and a temporary need for GPUs. They would book the required hardware and duration, with clear pricing and service terms.

Proposed service

Suppliers would choose which resources to offer, the availability window and minimum acceptable price. Buyers would compare suitable offers and reserve individual GPUs or multiple nodes, where the provider supports the configuration.

Market references include SF Compute for GPU capacity resale and Touchmark for reserved inference. Our proposed service adds the execution and onboarding workflow on partner infrastructure.

- Confirmed bookings would cover an agreed window; interruptible offers would be identified separately.
- Quotes would include GPU time, setup, storage, data transfer and platform charges.
- The supplier would retain control over external access, resource allocation and security requirements. Resale of reserved capacity would require the relevant permissions.

Flexible onboarding and execution

Customers would keep their existing development and experiment workflow. We would start with the connection method that fits the first workload:

- Connect selected experiments through an integration with tools such as W&B, MLflow or ClearML, using available configuration and artifacts and returning logs and metrics.
- Bring a repository or container with a launch command, resource requirements and authorized data access.
- Use a guided job launch, model endpoint or direct VM access, according to the provider's capabilities.

Our execution layer would check software and driver compatibility, prepare data near the selected GPUs and manage the launch. Training support would include checkpoints and results; inference support would include batch processing, model serving and planned traffic spikes.

Pilot proposal

The first pilot would test a complete transaction and workload on partner infrastructure: agree the capacity and price, prepare the job, run it and reconcile usage and payment. The partner and customer would set the hardware, duration and security requirements, including access for development.

- Agree one initial customer workload and the inputs required to run it.
- Connect an agreed pool of resources and test onboarding, isolation and metering.
- Complete a paid workload, review the economics and attempt a repeat booking.

Contributions and commercial terms

- Infrastructure partner: provide agreed capacity, access boundaries and operational support; confirm that external use and any resale are permitted.
- Our team: develop the booking and execution integration, help source customer demand, prepare workloads and handle the agreed customer workflow.
- Customer: provide the workload, environment, authorized data access, budget and acceptance criteria.

We propose a fee on completed bookings, with any partner revenue share agreed in advance. An onboarding or software agreement could be an alternative where the provider prefers that model.

Success measures

We would measure capacity sold, supplier revenue, the customer's full cost, onboarding time, execution success and willingness to book again. Benchmarking would support resource selection and verify that time, performance and model quality meet the workload's needs.

A first pilot would establish the operational and commercial case for further integration. Reserved inference with defined model, throughput and service terms could follow where the partner supports it.

Founder experience

Grigorii Penkin has practical experience with ML experimentation and GPU workload execution. In a previous role, his team used SkyPilot before adopting an in-house execution approach. He has also evaluated dstack. This experience informs the focus on preparing existing code, environments and data for execution on external GPUs.

Introductions requested

The project is at the concept and prototyping stage. We would welcome introductions to GPU providers or authorized capacity holders with available resources, ML teams seeking temporary compute, and people responsible for GPU infrastructure or ML platforms.