

CASE STUDY

How one company went from 28% GPU utilization to 73% with Run:ai

A multinational company, focused on innovations in computer vision technology, managed to go from 28% GPU utilization to optimization of over 70%, and achieve a 2X increase the speed of their training models with Run:ai. After expecting to make an additional GPU investment, with a planned hardware purchase cost of over \$1 million dollars, the company can now maintain their current infrastructure, meet the needs of their data science teams, and even improve training times.

Customer Background

A world leader in facial recognition technologies, the company provides AI services to many large enterprises, often in real-time. Accuracy, measured in terms of maximizing performance of camera resolution and FPS, density of faces, and field of view are critically important to the company and their customers.

Customer AI Infrastructure and Team

- On-Premises environment with 24 Nvidia DGX servers and additional GPU workstations
- 30 researchers spread on two continents

After Implementing Run:ai's Platform

70% Average GPU Utilization

> Higher ROI

2X Experiments per GPU

> Better Data Science

Multi-GPU

Training by Default

> Faster Time to Value



Simplified Workflows

> Reduced DS Hassles

- Data science teams simplified GPU utilization workflows and increased productivity by 2X, allowing them to more quickly deliver value with deep learning models
- Removed bottlenecks, resulting in faster training times - shortened by 75% on average
- Gained control and visibility into GPU clusters and saw utilization go from 28% to more than 70% for better budgeting and planning of new hardware needs
- Achieved ability to scale deep learning so new researchers and jobs easily gain access to infrastructure

Challenges

- **Sharing Resources Across Teams and Projects was Unsuccessful:** GPU resources were statically allocated, creating times with bottlenecks and other times with inaccessible, but available infrastructure.
- **Prioritizing and Scheduling Deep Learning Training Tasks was Ineffective:** Researchers lacked the ability to see and manage available resources which was slowing down their jobs.
- **Expensive Investment in Scaling Hardware Led to Increased Costs:** Although the utilization of existing hardware was extremely low, visibility issues and bottlenecks made it seem like additional hardware was necessary.
- **Low GPU Utilization:** Some peaks, but mostly inefficient and underutilized resources.
- **Different Usage:** Profiles 'Build' 'Train' 'Retrain' with very different needs





Solution

Run:ai's platform capabilities enabled the Company achieve:

- Increased GPU Utilization by moving teams from static, manual GPU allocations to pooled, dynamic resource sharing across the organization.
- Increased Productivity for the data science teams using hardware abstraction, simplified workflows, and automated GPU resource allocations.
- Visibility into the GPU cluster, its utilization, usage patterns, wait times, etc., allowed the company to better plan hardware spending.
- Accelerated Training Times, using automated, dynamic allocation of resources which enabled the data science teams to complete training processes significantly faster.

Increased GPU Utilization

Project and node utilization is now visible to all teams.



About Run:ai

Make data science more productive with the compute power to run more experiments and deploy AI initiatives faster. Run:ai helps organizations accelerate their AI journey, from building initial models to scaling AI in production. Using Run:ai's software platform, companies streamline the development, management and scaling of AI applications across any infrastructure (on-premises, edge, cloud). Researchers gain on-demand access to pooled resources for any AI workload. An innovative, cloud-native operating-system helps IT manage everything from fractions of GPUs to large-scale distributed training. Learn more at www.run.ai.