

System Base Labs

A Carbon-Neutral Company 

SBL

Today's AI Startup. Engineering the Intelligence of Tomorrow



Raw Compute Power

Multi-GPU clusters delivering
petaFLOPs of AI training throughput

Authored by Aleiman Shankar Rao
Founder & CEO, System Base Labs
www.systembaselabs.com



AI-First
Technology



Ethical AI



GPU Farms



Shankar AI



Blockchain +
Biomedical



Education

At the heart of Shankar AI's intelligence lies a roaring engine of performance—multi-GPU clusters engineered to deliver petaflops of training throughput.

Designed for next-gen AI workloads, these compute nodes are stacked with the latest NVIDIA H100s, tensor cores, and parallel-processing pipelines, ready to take on billion-parameter models with ease.

What is Shankar AI?

Shankar AI is a next-generation Augmented AI Engineering platform designed to deliver human-centric intelligence, adaptive automation, and scalable decision systems across industries. Architected on System Base Labs' green compute infrastructure and integrated with OpenAI's advanced LLMs, Shankar AI bridges high-performance computing with interpretable AI.

What is Raw Compute Power?

Raw compute power refers to the sheer computational capacity available to process large datasets and complex AI models.

Measured in FLOPs (Floating Point Operations per Second), it's a critical metric in evaluating how fast and efficiently AI tasks can be performed.

How We Achieve It

We leverage cutting-edge GPU technology like NVIDIA H100s, interconnected through high-bandwidth networks, and orchestrated using Kubernetes and Slurm to optimize resource allocation. These clusters are fine-tuned for energy efficiency and parallel processing, enabling performance at scale.



AI-First
Technology



Ethical AI



GPU Farms



Shankar AI



Blockchain +
Biomedical



Education

Seamless Workload Orchestration

What is it: The automated management and scheduling of compute tasks across multiple GPUs.
How we achieve it: We use orchestration tools like Kubernetes, Docker, and workload managers to automate job distribution.

Advantage: Maximizes resource utilization, reduces idle time, and ensures efficient training pipelines.

Ultra-Deep Neural Network Training

What is it: Training AI models with hundreds or even billions of parameters.

How we achieve it: Our clusters support parallel and distributed training frameworks like PyTorch DDP and TensorFlow MultiWorkerMirroredStrategy.

Advantage: Enables faster convergence, more accurate models, and the ability to tackle complex problems in medicine, finance, and more.

Continuous Fine-Tuning at Scale

What is it: Iterative model updates to adapt to new data and improve performance.

How we achieve it: Our pipeline integrates CI/CD for machine learning with version control, data refreshes, and automated retraining cycles.

Advantage: Keeps AI models fresh, relevant, and accurate without downtime.

Raw speed meets raw intelligence. Welcome to a powerhouse built for tomorrow's thinking.



AI-First
Technology



Ethical AI



GPU Farms



Shankar AI



Blockchain +
Biomedical



Education