

INDUSTRY-READY AI AT A FRACTION OF THE COST

Brought to you by



Powered by **MiPhi**



Train Real-World Models Efficiently

Run large-scale models like DeepSeek70B and Falcon 180B using just 2 to 4 GPUs, thanks to MiPhi's memory optimization.



Up to 90% Lower CapEx

Compared to traditional GPU clusters, aiDAPTIV+s delivers the same performance at up to 90% lower capital expenditure.



Up to 90% Lower OpEx

aiDAPTIV+ systems reduces ongoing power, cooling, and operational costs by as much as 90%, making AI infrastructure sustainably affordable.



Multi-User Concurrency

aiDAPTIV+ enables multiple users or students to train, fine-tune, or run inference at the same time, making it ideal for AI labs, classrooms, or collaborative R&D environments.



Scalable AI Infrastructure

Start with a small team or class (10 users) and scale up to 100+ users as your needs grow.

* Why It Works

MiPhi's aiDAPTIV+ uses an optimized middleware layer that connects your GPUs to MiPhi's high-speed aiDAPTIV+ Cache (2TB–4TB SSD memory). This adds memory capacity beyond the GPU's VRAM, enabling large AI models to run on fewer GPUs.

The middleware automatically slices and manages model data flow, no code rewrites needed.

* Ideal for

- Colleges & Universities building AI labs
- Enterprises running private GenAI workloads
- Government or R&D institutions with sensitive data
- Startups scaling AI workloads cost-effectively

KEY STATS

Real-World Comparison (Based on 70B Model) Total VRAM Memory Required 1.2 TB

Traditional AI Systems		Miphi's Aidaptiv+ System
62 x Nvidia RTX 4000 ADA (20GB each)	GPUs Required	2 x RTX 4000 ADA (with aiDAPTIVCache)
8 servers (up to 8 GPUs/server)	Servers Used	1 Workstation/server
8,000W (GPU only)	GPU Power Consumption	260W (GPU only)
₹2.4 Crores	CapEx Cost (excl. software)	₹25 Lakhs incl ProSuite Software
Limited (one workload at a time)	Concurrent Users	Supports multiple users simultaneously

Note: This comparison is based on training a 70B full-precision (FP32) model — with no quantization, compression, or reduction in accuracy. aiDAPTIV+ is also capable of training ultra-large models like DeepSeek 671B using just a single rack server, something previously possible only in large-scale GPU clusters.

* Academic & Industry Integration Support

MiPhi doesn't just provide the hardware — we also ensure your faculty and students are ready to use it.

Faculty Enablement & Curriculum Development: MiPhi provides training for educators to use the aiDAPTIV+ system effectively and supports curriculum design that aligns with current industry needs. MiPhi also includes access to its free ProSuite tools, making it easier for faculty and students to work with large AI models, manage experiments, and optimize training workflows.

Real-World Project Collaboration: Institutions using MiPhi infrastructure can engage students in real industry projects, building job-ready skills through practical experience.