

Table 1: AI Development Hardware Comparison for Research Labs (2026 Edition)

Category	Speed & Value	Inference / LLM	Pro Reliability	Maximum Infra	Memory
Representative Product	RTX 4090 Desktop PC	NVIDIA DGX Spark	RTX 6000 Ada Workstation	DGX Station/Server	Mac Studio
Chip / GPU	GeForce RTX 4090	Grace Blackwell	RTX 6000 Ada	H100 / B200 etc.	Apple Silicon
VRAM	24 GB	128 GB	48 GB	288 GB+	Up to 128 GB
Compute Stack	CUDA (standard)	CUDA (standard)	CUDA (standard)	CUDA (standard)	MLX
Lab Noise Level	Loud during seminars	Desk-side conversation OK	Tolerable under desk	Requires separate room	Silent — no fans
Training Inference					
Key Strength	Fastest compute Rapid paper prototyping	CUDA + large memory Loads huge LLMs easily	24/7 stable operation Very low error rate	Shared across teams Multi-user parallel training	Extremely quiet Inference + training
Critical Weakness	VRAM bottleneck, high heat OOM errors on large models	Slower training speed Narrow bandwidth, poor for training	High upfront cost Less cost-efficient than consumer GPU	High maintenance & setup cost Power & space requirements	No CUDA support Latest paper models
Support	Domestic BTO vendor	NVIDIA official support	3–5 yr on-site warranty	24h professional support	Apple support
Est. Price (incl. tax)	From approx. ¥650K	From approx. ¥750K	From approx. ¥1.7M	From approx. ¥10M	From approx. ¥1.5M

* Prices and specifications are estimates based on market research as of 2026.