

BUILT FROM THE GROUND UP FOR THE INTERSECTION OF AI & HPC

NVIDIA® GPU Solutions



NumberSmasher 4U Tower/GPU Server

Balanced CPU:GPU Configuration

A balanced, versatile configuration with 2 CPUs and 4 GPUs suitable for many HPC applications. These 4U GPU servers offer additional PCI-E slots and optional hardware RAID support for up to 8x drives. Support for a wide array of GPUs and PCI-E devices is also available.



Navion 2U with 4 NVIDIA A100 with NVLink GPUs

NVLink Interconnect

Compute faster with four NVLink-connected NVIDIA A100 GPUs for AI training. With two AMD EPYC Processors and four NVIDIA A100 GPUs, this server offers a compelling pairing of leadership-class CPUs and GPUs. It delivers high-bandwidth transfers via PCI-E Gen4 interfaces and is available with high-speed fabric + NVMe.



NumberSmasher® 1U NVIDIA PCI-E GPU Server

Cost-Effective Density

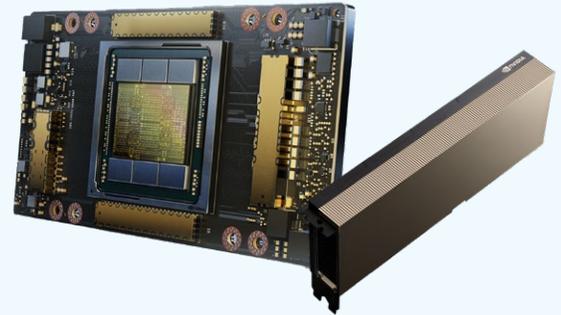
Customize to your workload with this powerful, versatile configuration. Supporting up to 4 PCI-E GPUs, plus high-speed fabric and storage, these 1U servers offer a compact footprint. With several CPU, GPU, and storage configurations available, our experts optimize these servers for your critical applications.



Navion 4U 8 GPU Server, Octopus™ Servers

Highest Capacity

Scale up to 195 FP32 TFLOPS, 1.5 PFLOPS of TF32 for AI with 8 or 10 NVIDIA A100 GPUs. Ideal for AI, highly accelerated GPU apps & GPU-Direct Peer-to-Peer. Optional InfiniBand for clustering, RAID, NVMe storage and Quadro® RTX professional graphics. Support for full PCI-E Gen4 link speeds in our Navion 4U 8 GPU Server.



NVIDIA® A100 GPU Accelerator

Unprecedented Acceleration at Every Scale

- 9.7 TFLOPS FP64 | 19.5 TFLOPS FP32 | 156 TFLOPS TF32
- Leading HPC, AI Training, or AI Inference Performance
- 3rd Gen NVLink™ interconnect: max 600GB/s per GPU
- 40GB/80GB memory, 1.5TB/2TB/s Memory Bandwidth
- New TF32 Instructions for Training + FP64 Tensor Cores
- Socketed SXM4 (NVLink) and PCI-E 4.0 Form Factors

NVIDIA® DGX A100™ with NVIDIA A100 GPUs

The Universal System for AI Infrastructure

- Nearly 5 PFLOPS of AI training performance
- 8 NVIDIA A100 with NVLink GPUs
- 3rd Gen NVIDIA NVLink Interconnect (600GB/sec/GPU)
- AI Training, Inference, or HPC
- Deep Learning & AI Appliance: Ready to run frameworks, Easy to Use Docker-interface



Software Integration Experts

- NVIDIA CUDA® Toolkit installed and configured
- Intel, NVIDIA HPC SDK with OpenACC, or GNU Compilers
- Linux (CentOS, RHEL, Ubuntu, more) or Windows Server
- Microway Cluster Management Software (MCMS™)
- Legendary Service and Technical Support



Schedule Number:
GS-35F-0431N



ELITE
PARTNER

Microway Scalable NVIDIA GPU Clusters



NVIDIA GPU Clusters

Microway's robust, NVIDIA GPU-based clusters offer high-speed interconnects, NVIDIA datacenter GPUs, MCMS Remote Cluster Management and Monitoring Tools or the OpenHPC Stack, plus the NVIDIA CUDA® and HPC SDKs.

Microway provides these fully integrated Linux clusters at very competitive prices. Users worldwide pushing the limits of technology in life sciences, universities, commercial and government research count on our expertise and attention to detail.

Delivering Innovative HPC Solutions Since 1982



	NVIDIA GPU Cluster Options
Processor	Intel® Xeon® Scalable Processors with 4-28 Cores AMD® EPYC™ Processors, 16-64 Cores
Processor Sockets	1P, 2P, 4P
Compute Nodes	NumberSmasher (Xeon) & Navion (EPYC) 1U, 2U, 4U GPU Servers NVIDIA DGX AI Systems
GPU Options	NVIDIA® A100, Tesla® V100 or Quadro® RTX
Visualization Node Option	NumberSmasher or Navion 1U, 2U, 3U, 4U GPU Servers
Memory per Node	2P: up to 3 TB DDR4 , 4P: Up to 6 TB DDR4
Disk/Media Bays	1U: 2x2.5" Hard Drives, 2U, 4U: 8x3.5" Hard Drives; OctoPuter: 24x 2.5" Hard Drives
	Diskless compute node configurations available
Cluster Interconnect	Mellanox® ConnectX®-6 HDR or EDR InfiniBand 200GigE, 100GigE, 50/25GigE, 10GigE or Gigabit Ethernet
Management Network	IPMI v 2.0 and Redfish
Storage Options	Gigabit or InfiniBand Storage NAS, Nexsan, DDN ExaScaler Lustre, BeeGFS, or Panasas®
Operating Systems	Linux: CentOS, Red Hat, Ubuntu, more; Windows® Server
Compilers	CUDA, NVIDIA HPC SDK, Intel and GNU
Cluster Software	MVAPICH2, OFED, and OpenMPI; SLURM, PBS Professional, or IBM Spectrum LSF
Monitoring & Management Software	Microway Cluster Management Software (MCMS) MPI Link-Checker™ and InfiniScope™ Bright Cluster Manager or IBM Spectrum Cluster Manager
Cabinets and Infrastructure	APC® NetShelter™ Cabinets, APC PDUs and UPS Protection
Services	Optional onsite installation, Factory Pre-Installation Service (Power & network pre-wiring, pre-installation of rails) Optional complete-rack air-ride shipment (roll-off, power-on)
Hardware Warranty	2 years with advanced replacement parts or return-to-factory Optional extended warranty term

www.microway.com/nvidia

Call our technical sales team at **508.746.7341**