

# AI / GPU Computing: 4U8GPU-A100-4189A

## 4U 8xGPU-A100, 6xU.2/SATA, 32x8-Channel DDR4, 3<sup>rd</sup> Gen. Scalable Xeon



### Key Features:

- 4U 8 x Nvidia® HGX A100 GPU 40GB/80GB, NVLink, NVSwitch
- Dual Socket P+ for 3<sup>rd</sup> Gen. Intel® Scalable Xeon Processors
- Front 6 x 2.5" NVMe U.2 4.0 / SATA / SAS Hot Swap System Drives
- Rear 4 x 2.5" NVMe U.2 4.0 SSD for GPU Direct Mapping Storage
- 2 x M.2 NVMe / SATA for OS Mirrored Drives
- 32 DDR4 up to 12TB 8-Channel Memory, supports Optane® Memory
- 10 PCIe 4.0 x16 Low Profile Slots
- OCP 3.0 PCIe 4.0 x16 for Networking Module, IPMI 2.0
- 4 x 2200W Titanium Redundant Power Supply
- GPU Cluster for AI, Deep Learning, HPC, etc.

### Specifications:

#### Processor

- Dual 3<sup>rd</sup> Generation Intel® Scalable Xeon
- Socket P+ (LGA 4189), supports 270W, Up to 40 Core 80 Threads

#### Chipset

- Intel® C621A

#### GPU

- Nvidia® HGX A100 8-GPU 40GB/80GB, NVLink, NVSwitch

#### Memory

- 32 DDR4 DIMM Sockets
- Supports up to 12TB ECC DDR4 3200MHz RDIMM
- 8-channel memory bus per processor
- Supports Intel® Optane™ persistent memory 200 series

#### Storage

- 6 x 2.5" NVMe U.2 PCIe 4.0 / SATA How Swap Bays
- Supports NVMe (16GT/s)/SATA (6Gb/s)/SAS (12Gb/s) RAID-0/1/5/10
- 2 x M.2-20110/2080 Sockets for NVMe 4.0 x4/SATA

#### Network

- OCP 3.0 (PCIe 4.0 x16) for Network Module
- Network Modules are available with 2-4 ports of 1G/2.5G/10G/25G/100G RJ45/SFP+/SFP28/QSFP28
- RJ45 Gigabit LAN for IPMI 2.0 by AST2600 (KVM)
- Additional LAN or InfiniBand by PCIe add-on cards

#### Expansion Slot

- 10 x PCIe 4.0 x16 (Low Profile 2 Front, 8 Rear)
- 1 x OCP 3.0 (PCIe 4.0 x16) Slot for Ethernet Network Module

#### Other I/O

- 1 x RJ45 IPMI (Front)
- 1 x DB15 VGA, 2 x USB 3.1 Gen2 (Front)

#### Internal I/O

- 1 x RS-232 COM and 1 TPM header
- 2 x USB 3.1/3.2 Header, 2 USB 3.2 Gen2 Type-A

#### BIOS

- AMI
- Watchdog Supported

#### Chassis

- 4U Rackmount
- 6 x 2.5" NVMe U.2 4.0/SATA/SAS Front Hot Swap Drive Bays
- 4 x 2.5" NVMe U.2 4.0 Rear GPU Storage Drive Bays
- Dimension: 6.9" (174mm) H x 17.6" (446mm) W x 35.4" (900mm) D
- System Weight: 170 lbs. (77kg) Net, 230 lbs. (105kg) Gross

#### Power Supply

- 4 x 2200W Redundant Titanium Level
- 1200W: 100-127Vac / 2200W: 220-240Vac Input

#### OS Supported

- Windows Server, Linux

#### Applications

- CPU Cluster for AI, Deep Learning, HPC.

#### Environments

- Operating Temperature: 0 to +35 C
- Non-operating Temperature: -40 to +60 C
- Vibration - Operating: 9.8 m/s<sup>2</sup>(1.0G) 5 to 500Hz
- Vibration - Non operating: 49 m/s<sup>2</sup>(5.0G) 15 to 500Hz
- Shock - Operating: 3,920 m/s<sup>2</sup> (400G) 2ms
- Shock - Non-Operating: 8,820 m/s<sup>2</sup> (900G) 1ms
- Operating Relative Humidity: 8% to 90% (non-condensing)
- Non-operating Relative Humidity: 5% to 95% (non-condensing)
- MTBF: 1,000,000 Hours
- FCC, CE, UL, RoHS Compliance