

3U Compute Accelerator with NVIDIA Tesla GPUs

The CA16000 Compute Accelerator with sixteen NVIDIA® Tesla® GPU accelerators is employed in a variety of HPC applications including oil and gas exploration and financial services. Completely integrated with the GPUs most suited for a specific application, it's easy installation and tested reliability makes it superior to alternative products. The CA16000 occupies only 3U of rack space and connects directly to one or four host server(s) through the latest technology PCIe x16 Gen3 connections. Four removable canisters house up to four full-height, full-length, PCIe x16 double-wide GPUs and one half-length, half-height IO card each. The system is powered by three 3000-watt redundant power supplies and includes an IPMI-based system monitor.

PN: OSS-PCIe3-3UZ-20-4-X



Features

- 3U High
- Four removable canisters with four GPUs each
- Optional 5th slot in each canister to support EDR IB or 10/40GbE
- Fully IPMI v2.0-compliant system monitoring capability
- Three 3000-watt power supplies
- Superior Cooling with eight temperature controlled fans
- Up to four PCIe x16 Gen3 cable connections to host server(s)

Specifications

Enclosure	<ul style="list-style-type: none"> • Dimensions: 17"W x 5.25" H x 38"D • Supports 16 full-length, full-height, 2-slot PCIe x16 GPU/PHIs • All 16 boards face the rear of the chassis (no IO bracket access) • Supports four half-height, half-length, single-wide PCIe x16 cards with IO bracket access • Removable front bezel with air filter • Front panel LEDs • Four rear panel PCIe x16 Gen3 cable interfaces • 4 individually-removable rear fans and 4 canister mounted fans • Weight: 92lbs when fully loaded with 16 accelerators
Main Backplane	<ul style="list-style-type: none"> • Four PCIe 3.0 x16 cable inputs to rear of enclosure • Four PCIe 3.0 x16 high-density connectors to each canister • Three PCIe 3.0 switches manage PCIe cross connects from cables to canisters • 2x RJ45 connectors for IPMI v2.0 System Monitor • 1x HD DB-9 serial port for IPMI network configuration • Optional RJ45 for basic SYSMON2 chassis monitor (not required when using IPMI System Monitor) • Supports bus-bar power distribution to the canisters through 8 high-power bladed connectors (2 per canister) • On board IPMI System Monitor & SYSMON2 connectors
Canister Backplane	<ul style="list-style-type: none"> • 4x PCIe 3.0 x16 double-spaced slots in 2 ranks of 2 GPUs each • 4x 8-Pin 12V power connectors for GPU/PHI AUX power cables • 1x PCIe 3.0 x16 single-wide slot for half-height, half-length IO cards • PCIe 3.0 switch
Power	<ul style="list-style-type: none"> • 6000W redundant power subsystem • Three 3U 3,000-watt front removable, hot-swap supplies • Each supply measures 1U (1.65") x 2.7" x 28.5" • 2+1 redundant with full current sharing operation • 3,000W each at 208-277VAC, 15A max input • 1,500W each at 90-124VAC, 15A max input • 15A breaker and IEC C19 power input at rear for each supply • +12V and +12V standby voltage outputs • All +12V power rails shared on copper bus bar delivery system

Power Cords	<ul style="list-style-type: none"> • The HDCA ships with C19-C20 15A PDU type power cords • 240V power cord for PDUs <ul style="list-style-type: none"> ◦ OSS Part number: OSS-CBL-PWR-C20-C19-15A-8 ◦ IEC C19 to IEC C20, Straight, 14AWG, 15A, 8' (2.44m) • Other power cords available on request
System monitoring/alarming	<p>Fully IPMI v2.0 compliant monitoring, control & alarming system</p> <p>Temperature</p> <ul style="list-style-type: none"> • Monitors inlet & exhaust temps • Fan speed auto adjusts by temp • Alarm set-points for over temp <p>Fans</p> <ul style="list-style-type: none"> • Monitors all system fan tachs • PWM fan speed control • Alarms for slow or failed fans <p>Power</p> <ul style="list-style-type: none"> • Monitors supply telemetry • Monitors output voltage rails • Alarms for voltages out of range • Alarms for supply failure <p>Add-in Cards</p> <ul style="list-style-type: none"> • Monitors add-in card I2C SM bus • Alarms for abnormal card telemetry <p>Interface</p> <ul style="list-style-type: none"> • CLI or web GUI • Supports SNMP and RCMP • Remote chassis and canister LED tagging
Air Filter	<ul style="list-style-type: none"> • 30 ppi open cell polyfoam • Die-cut, removable and replaceable
Cooling	<ul style="list-style-type: none"> • Four 80 x 80 x 38mm fans on the rear of the enclosure • One 80 x 80 x 38mm fan on the front of each canister • All fans are 141CFM each in push-pull configuration • All fans PWM monitored and speed controlled by the IPMI system monitor • Rear fans hot-swap from rear of the chassis • Power supplies separately cooled from internal 25mm fans • Optional liquid cooling for approved accelerators available
Operating Environment	<ul style="list-style-type: none"> • Temperature range: <ul style="list-style-type: none"> • Operating: 10°—35°C • Storage: -40°—85°C • Humidity range: <ul style="list-style-type: none"> • Operating: 20% to 80% relative (non-condensing) • Non-operating: 5% to 95% relative (non-condensing) • Altitude range: <ul style="list-style-type: none"> • Operating: 0 to 10,000 ft. • Storage: 0 to 50,000 ft.
Agency compliance	<p>Designed to meet the following agency certifications with testing currently pending:</p> <ul style="list-style-type: none"> • FCC - Part 15 of the FCC Rules, Class A, 47CFR • Canada ICES-003, issue 4, Class A • UL/IEC 60950-1 • Canada: CSA C22.2 No. 60950-1 • Argentina: IEC60950-1 • Japan: VCCI, Class A • Australia/New Zealand AS/NZS CISPR 22, Class A • IEC 60950-1 (CB Certificate and CB Test Report) • CE Mark (EN55022 Class A, EN60950-1, EN55024, EN61000-3-2, EN61000-3-3) • CISPR 22, CISPR 24, Class A • CE Emissions 2004-108EC • RoHS compliance (Directive 2002/95/EC) • CCN NWGQ, NWGQ7

	Tesla M40	Tesla M60	Tesla K80	Tesla P100
Peak Double Precision Performance	.2 TeraFLOPS	3.8 TeraFLOPS	2.91 TeraFLOPS	4.7 TeraFLOPS
Peak Single Precision Performance	7 TeraFLOPS	7.4 TeraFLOPS	8.74 TeraFLOPS	9.3 TeraFLOPS
Number of GPUs	1 Maxwell GM200	2 Maxwell GM204s	2 Kepler GK210s	1 Pascal GP100
Number of CUDA Cores	3072	4096	4992	3584
Memory Capacity and Bandwidth	24GB GDDR5 at 288GB/s	16GB GDDR5 at 320GB/s	24GB GDDR5 at 480GB/s	16GB CoWoS HBM2 at 720GB/s or 12GB CoWoS HBM2 at 540GB/s
Power Consumption	250 W	300 W	300 W	250 W

For a list of qualified servers, go to <http://www.onestopsystems.com/hpc/3u-compute-accelerator-nvidia-tesla-gpus>