ONE STOP SYSTEMS

GPUltima

The OSS GPUltima is a petaflop compute accelerator in a single rack. The OSS GPUltima supports up to 128 double-wide interconnected accelerators. Consuming only 56 kilowatts of power, the platform uses 94% less power than other petaflop solutions. In addition the single-rack solution occupies 96% less space than other solutions. The OSS GPUltima employs 64 OSS PCIe Gen3 x16 adapters and cables between the OSS servers and the OSS High Density Compute Accelerators (HDCA) each containing 16 NVIDIA K80 GPU's and 32 Mellanox EDR 100Gbs InfiniBand adapters and cables connecting the GPUs through a 36-port InfiniBand switch. One Stop Systems' PCIe expansion systems provide high speed connectivity between the CPU and IO cards.

Features

- Single 19" rack
- Power input and distribution
- 8 OSS compute accelerators
- 16 dual socket servers
- 128 NVIDIA dual GPU cards
- Infiniband Switch
- Ethernet Switch

Specifications

Rack	○ 42U tall rack with extra depth ○ Also available in 24U, 44U and 48U tall versions ○ Optional doors, etc.
High Density Compute Accelerator	○ Supports 16 double-wide GPU cards ○ Four PCIe over cable connections to host servers ○ 6,000 watts n+1 power ○ High-speed push-pull fan configuration and channeled airflow
Server	○ Supermicro 1U Server ○ Dual socket with Intel Xeon "Broadwell" CPUs ○ Two 3.5" hot-swap disk drives ○ Four Gen 3 x16 PCIe expansion slot plus 2 x8 slots ○ Redundant 2,000 watt power supplies
Infiniband Switch	○ Mellanox 36 port Infiniband switch ○ EDR 100Gb/s, QSFP connectors ○ 1U form factor
Ethernet Switch	○ Cisco 48 port switch ○ 1U Form Factor
GPUs	○ NVIDIA K80 ○ Contains 2 GPU components ○ 4,992 cores per card ○ 8.74 Tflops per card (boost enabled) ○ 24GB memory ○ PCIe Gen 3 x16 interface to backplane ○ 300 watts

PCIe Over Cable Interface Card	○ PCIe Gen 3 x16 interface ○ Two cards used per server
Infiniband Interface Card	○ EDR 100Gb/s, QSFP connectors ○ Single port version, dual port also available ○ Two cards used per server
Power Distribution Unit Option 1	○ Tripp-Lite Monitored PDU ○ 12.6kW power ○ Input: 208V 3 phase, 60A ○ Power monitoring via display and Ethernet ○ System utilizes 8 PDUs ○ 100kW total power ~ 80% over-provisioned
Power Distribution Unit Option 2	○ Tripp-Lite Monitored PDU ○ 27.6kW power ○ Input: 380/400V 3 phase, 63A ○ Power monitoring via display and Ethernet ○ System utilizes 4 PDUs ○ 110kW total power ~ 97% over-provisioned
Cables	○ Power cables – 59 ✓ 32 for servers ✓ 24 for compute accelerators ✓ 1 for Ethernet switch ✓ 2 for Infiniband switch ○ PCIe cables – 32 ✓ 32 for servers to compute accelerators ✓ Gen 3 x16 – 1 meter long ○ Ethernet cables – 48 ✓ 32 for servers to Ethernet switch ✓ 8 for compute accelerators to Ethernet switch ✓ 8 for PDUs to Ethernet switch ○ Infiniband cables – 32 ✓ 32 for servers to Infiniband switch
Optional Liquid Cooling	CoolIT Direct Contact Liquid Cooling Model DCLC Custom solution providing special heat sinks with liquid cooling onto each GPU card Tubing to front of each GPU canister Plumbing within rack External pump and heat exchanger

Architecture