BOSTON ANNA PASCAL XL
PERFECT FOR HIGH PERFORMANCE COMPUTING & DEEP LEARNING

BOSTON ANNA PASCAL XL

Featuring up to eight NVIDIA® Pascal™ P100 GPUs, connected via NVLink™ in a 4U rackmount server, the Boston ANNA Pascal XL is perfect for high performance computing & deep learning. Dual Xeon E5-2600 v4 processors and a huge 3TB DDR-4 registered ECC memory capacity to make this the perfect choice for intensive processing tasks, like neural network training, or scientific simulation.

A compliment of 16 x 2.5" hot pluggable disk bays allow for a range of high performance storage options, including SAS, S-ATA, or even up to eight NVMe based devices for extreme IO requirements.

There are several networking choices with dual 10 Gigabit Ethernet ports on board, plus add on card options to include four 100GbE, EDR InfiniBand or Omni-Path ports with non-blocking PCI-Express connectivity.

Eight SXM2 slots support the latest generation of NVIDIA® GPUs, including the Pascal based P100. NVIDIA® NVLink™ connectivity of 80GB/s per link ensures that there is plenty of bandwidth between GPUs for even the most IO bound applications, reducing latency and the CPU as a bottleneck for transfers.

KEY FEATURES
- Deep Learning & HPC GPU Accelerated Compute Solution
- 8 x SXM2 connected NVIDIA P100 GPUs a in 4U
- NVIDIA® NVLink™ connectivity between GPUs as standard
- Dual Intel Xeon E5-2600 v4 processors, up to 22 cores each
- Up to 3TB DDR-4 Registered ECC Memory
- EDR InfiniBand or Omni-Path interconnect options
- 16 x SAS/SATA or 8 x NVMe HDD / SSD support

SUITABLE USAGES
- Deep Learning
- Advanced Neural Network Analysis
- High Performance Computing
- Molecular Dynamics
- Computational Fluid Dynamics

Powered by four ground-breaking technologies with discontinuous jumps in performance, NVIDIA® Tesla® P100 GPUs enable lightning fast nodes to deliver the highest absolute performance for HPC and deep learning workloads with infinite computing needs.

@bostonlimited | sales@boston.co.uk | www.boston.co.uk | +44 (0) 1727 876 100

BOSTON Servers | Storage | Solutions | Powered by SUPERMICR®

Today's leading deep learning models typically take days to weeks to train, forcing data scientists to make compromises between accuracy and time to deployment. The Boston ANNA Pascal XL, features 8 x NVIDIA® Tesla® P100 GPU accelerators for deep learning training, purpose-built to dramatically reduce training time. Running Caffe and Torch on the Tesla P100 delivers the same model within days versus weeks on CPU based compute systems.

Boston Ltd are also offering customers a dedicated optimised clo ud platform for their deep learning analytics solutions - removing the initial capital expenditure of a physical on premise system, and enabling you to immediately gain insight through processing and analysing your data in the cloud. Our cloud platform enables agile sys-admin teams to quickly deploy scalable, production-ready deep learning environments in its own public cloud.

BOSTON ANNA PASCAL XL BENEFITS

HIGH PERFORMANCE

The Boston ANNA Pascal XL is perfect for high performance computing & deep learning. Dual Xeon E5-2600 v4 processors and a huge 3TB DDR-4 registered ECC memory capacity to make this the perfect choice for intensive processing tasks, like neural network training, or scientific simulation.

DEEP LEARNING

Today's advanced deep learning neural networks, as found in the Boston ANNA Pascal XL, use the power of GPUs to learn at speed, with accuracy and to scale. These are all driving forces of the artificial intelligence and AI computing revolution.

INTERCONNECTED

One of the most interesting features of the Boston ANNA Pascal XL is the groundbreaking NVIDIA® NVLink™ interconnect system. Designed to solve the challenges of exascale computing, NVLink is a high-bandwidth, energy-efficient interconnect that enables ultra-fast communication between the GPUs.

BOSTON ANNA PASCAL XL SPECIFICATIONS

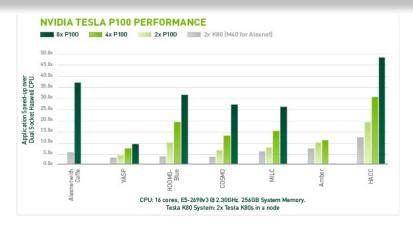|	
FORM FACTOR	4U Chassis
CPU	Dual HSW/BDW CPUs
MEMORY	24 DDR4 DIMMs
DRIVES	16x 2.5" HS SATA/SAS bays
GPU	8 Pascal w/ 80GB/s NVLink
EXPANSION SLOTS	4 x16 PCIe 3.0 slot; 2 x8 slot
POWER SUPPLY	2x 2000W Titanium PWS

Customised configurations are also available - contact us to discuss your requirements

NVIDIA TESLA P100 PERFORMANCE

8x P100 4x P100 2x P100 2x K80 (M40 for Alexnet)

CPU: 16 cores, E5-2698v3 @ 2.30GHz. 256GB System Memory.
Tesla K80 System: 2x Tesla K80s in a node

BOSTON
Servers | Storage | Solutions

Powered by
SUPERMICRO