



## High Performance Parallel Computing Solutions - Powered by NVIDIA® Tesla™

AMAX's GPU parallel computing solutions are powered by the NVIDIA® Tesla™ 20-series GPU computing platforms, based on the new Tesla T20 "Fermi" CUDA massively parallel architecture. They support "must have" features for the technical and enterprise computing space, including ECC memory for uncompromised accuracy and scalability, and 8X the double precision performance compared with previous generation of GPU products.

Compared to the latest x86 CPUs, AMAX's GPU parallel computing solutions deliver equivalent performance at 1/20th the power consumption and 1/10th the cost, enabling end-customers to solve the world's most important computing challenges more quickly and accurately across many industries.

### Parallel HPC Clusters



- Up to 86.52 Teraflops of single precision performance and 43.26 Teraflops of double precision performance per cluster
- Up to 504GB dedicated GPU memory
- Up to 84 12-Core AMD or 6-Core Intel processors on host systems
- Up to 252TB hot swap shared storage on host systems
- QDR InfiniBand interconnect
- Preinstalled with CUDA 3.0 Toolkit and SDK, Clustering software

### Personal Supercomputers



- 1,792 CUDA Cores, 4.12 Teraflops single precision & 2,060 Gigafllops double precision floating point performance (peak) Intel Xeon DP PSC, PSC-2n
- 1,344 CUDA Cores, 3.09 Teraflops single precision & 1,545 Gigafllops double precision floating point performance (peak) Intel Core i7 PSC, PSC-1i
- 896 CUDA Cores, 2.06 Teraflops single precision & 1,030 Gigafllops double precision floating point performance (peak) Intel Xeon DP PSC, PSC-5n

### 1U/4U GPU Servers



- 896 CUDA Cores, 2.06 Teraflops single precision & 1,030 Gigafllops double precision floating point performance (peak), Intel Xeon DP 1U GPU server with 3 x 3.5" hot swap drive bays, Xn-1113G / Xn-1113M
- 896 CUDA Cores, 2.06 Teraflops single precision & 1,030 Gigafllops double precision floating point performance (peak), Intel Xeon DP 1U GPU server with 6 x 2.5" hot swap drive bays, Xn-1114G
- 3,584 CUDA Cores, 8.24Teraflops single precision & 4,120 Gigafllops double precision floating point performance (peak) Intel Xeon DP 4U GPU server, Xn-4101G

### Tesla S2050 1U GPU Computing Server



- Seamlessly fits into enterprise server clusters and scales to solve the most complex computing problems.
- 1,792 CUDA Cores, 4.12 Teraflops single precision & 2,060 Gigafllops double precision floating point performance (peak)

### Tesla S2050 Host Servers



- Seamlessly connects to a Tesla S2050 server and scales to solve the most complex computing problems.

### KEY FEATURES:

- GPUs powered by the massively parallel CUDA architecture, transforming a GPU solution to perform like a small cluster at 1/20th the power consumption and 1/10th the cost.
- CUDA programming environment with broad support of programming languages and APIs, including C, C++, OpenCL, DirectCompute, or Fortran
- IEEE 754 single and double Precision floating point units, delivering up to 8X the double precision performance compared with previous generation of GPU products
- ECC Support, meeting a critical requirement for datacenters and supercomputing centers deploying GPUs on a large scale with uncompromised computing accuracy and reliability.
- NVIDIA® Parallel DataCache™, accelerating algorithms such as physics solvers, ray-tracing, and sparse matrix multiplication where data addresses are not known beforehand.
- NVIDIA® GigaThread™ Engine, maximizing the throughput by faster context switching, concurrent kernel execution, and improved thread block scheduling.
- Asynchronous Transfer, turbocharging system performance by executing data transfers, even when the computing cores are busy.

### Applications:

