



## AMAX High Performance GPGPU Clusters

Dense, Scalable and Efficient  
Supercomputing Solutions for Every Application

# Applications and Introduction

## Tesla™ Parallel High Performance Computing (HPC) Solutions

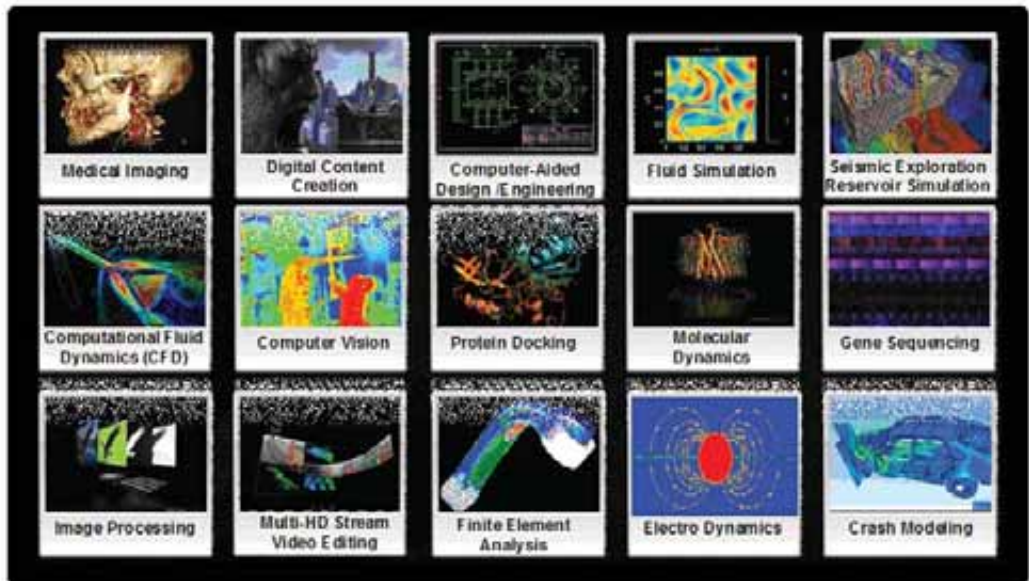
Accelerate your Applications with AMAX's Massively Multi-threaded, Parallel Computing Architecture

AMAX utilizes the revolutionary **NVIDIA® Tesla™**, the world's first teraflop, many-core processor technology, to deliver parallel high performance computing (HPC) solutions to solve the world's most important computing challenges - more quickly and accurately. With **448 cores per processor** and a complete suite of developer tools (C Compiler, debugger, performance profiler, and optimized libraries), AMAX's highly scalable and energy efficient Tesla™ HPC solutions provide a powerful platform for accelerating your applications.



When compared to the latest quad-core CPU, Tesla 20-series GPU computing processors deliver equivalent performance at 1/20th the power consumption and 1/10th the cost. Each Tesla GPU features hundreds of parallel CUDA cores and is based on the revolutionary **NVIDIA® CUDA™** parallel computing architecture with a rich set of developer tools (compilers, profilers, debuggers) for popular programming languages APIs like C, C++, Fortran, and driver APIs like OpenCL and DirectCompute.

### Applications:



# Cluster Management Software

## Cluster Management Options

AMAX provides customers with a flexible range of open source commercially supported cluster management solution options that can be easily tested and pre-integrated as a part of a complete package to include HPC professional services and support. The comprehensive features of AMAX's cluster management solution help lower total cost of system ownership, increase productivity, and provide a better return on your investment.

AMAX's cluster management software provides the necessary power and flexibility to monitor essential system metrics from a single point of control. It reduces the time and resources spent administering the system by improving software maintenance procedures and automating repetitive tasks. It also offers a web-based management interface that is easy-to-use providing customers with powerful tools to manage the software stack, including workload manager, resource management tools, remote control and advanced power management functions.

In addition, AMAX's cluster management software supports multiple networking topologies, diskless configuration and network failover to achieve maximum reliability, performance and high availability. It features root file systems for instant provisioning of rapid, standard Linux installs on large diskless systems allowing them to boot 64 to 6,400 nodes at the same time.

## Key Features & Benefits:

- High Performance Monitoring -- Provides real-time analysis of over 120 essential system metrics from each node.
- Automated System Management -- Takes pre-emptive actions when set system thresholds have been exceeded. System thresholds can be defined on any set of the standard metrics.
- High Performance Provisioning -- Through multicast technology, an image can be copied to every system node in the same amount of time it takes to copy an image to a single node.
- Version Controlled Image Management -- Tracks the changes to the OS or kernel over time.
- Heterogeneous IPMI 2.0 Support -- Makes it easy to configure IPMI, provide remote console access, power management, and environmental monitoring; including temperature, fans, voltages and other metrics provided by IPMI.
- Easy to Use -- An advanced Java-based GUI streamlines workflow by putting all common tasks into easily accessible locations.
- Simplifies IT Management & Administration -- Offers robust capabilities including: Network management, server management, cluster management, storage management



# Cluster Management Software Specifications

## Cluster Management Technical Specifications

### Cluster Management

- Supports partitioning a cluster into multiple logical computers
- Maps logical computers (clusters) onto servers (nodes)
- Supports multiple independent OS configurations
- Manages and monitors logical computer (clusters) status
- Provides cluster status to management system
- Integrated job scheduling and management
- Manages and monitors operating system instances (nodes)
- Provides node status to management system

### System Management

- Management of overall system configuration
- Supports redundant management servers with automatic failover
- Designed to anticipate and tolerate failures
- Supports enterprise level availability requirements with 24 hr MTTR

### Server Management

- Automatic discovery of server hardware
- Remote server control (Power On/Off, Cycle)
- Remote server initialization (Reset, Reboot, Shut Down)
- Support for scalable fast diskless or data-less booting for large node count systems
- Supports server redundancy and failover
- Provides server status to the management system

### Network Management

- Automatic discovery of interconnect hardware
- Supports multiple interconnect fabric topologies
- Supports redundant paths and networks
- Supports load balancing and failover
- Provides network status to the management system

### Storage Management

- Supports scalable root file systems for diskless or data-less nodes
- Supports multiple global storage configurations
- Supports high BW to secondary storage for data and check pointing
- Provides server status to the management system

### Graphics User Interface

- Web based

### Reliable, Available, Serviceable (RAS)


- Redundant management network (GbE and 10GbE) with failover
- Redundant high bandwidth network (InfiniBand™) with failover
- Redundant management servers (two levels) with failover
- Redundant root file system with failover
- Built-in multi-generation configuration management for software
- Redundant storage via RAID

# Intel Xeon 5600 Series (Westmere-EP) CPU

Employing Intel's latest technology, the **Intel Xeon 5600 Series (Westmere-EP)** provides up to 35% improvement in virtualization with increases in threads and cache and latency, making AMAX's GPGPU cluster more powerful and efficient than ever!

## Intel Xeon 5600 Series (Westmere-EP): Next Generation Xeon Processor

- Up to Six cores & 12 threads per processor
- Up to 12MB L2 cache
- Support for low power DDR3 1.35V DIMMs
- Intel® Turbo Boost Optimization
- Trusted Execution Technology® for secured trusted system boot
- Faster encryption with new instructions for the Advanced Encryption Standard (AESNI)
- Further reduction in latency for virtualization for direct I/O
- Higher frequencies on servers for 130w systems
- Lower power LV skus targeted down to 40w TDP



# Tesla

## PETASCALE COMPUTING WITH TERAFLOP PROCESSORS

Based on the new NVIDIA CUDA™ GPU architecture codenamed “Fermi”, the Tesla™ C2050 / C2070 computing processors are designed from the ground up for high performance computing. It supports “must have” features for technical and enterprise computing including C++ support, ECC memory for uncompromised accuracy and scalability, and 8X the double precision performance compared to the Tesla 10-series GPU computing products. Compared to the latest quad-core CPUs, Tesla 20-series GPU computing processors deliver equivalent performance at 1/20th the power consumption and 1/10th the cost.

## GPU powered by the massively parallel CUDA architecture

Transform your workstation 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

Choose C, C++, OpenCL, DirectCompute, or Fortran to express application parallelism and take advantage of the “Fermi” GPU’s innovative architecture.

## Scale to multiple GPUs and harness the performance of a small server cluster

Solve large-scale problems faster than a small server cluster on a single workstation with multiple GPUs.

## IEEE 754 single and double precision floating point units

Achieve up to 600 Gigaflops of double precision performance for faster and more accurate results.

## ECC Support

Offers protection of data in memory to enhance data integrity and reliability for applications. Register files, L1/L2 caches, shared memory, and DRAM all are ECC protected

## Up to 6GB GDDR5 compute memory

Allows faster access to larger data sets.

## NVIDIA® Parallel DataCache™ Technology

Accelerates algorithms such as physics solvers, ray-tracing, and sparse matrix multiplication where data addresses are not known beforehand.

## NVIDIA® GigaThread™ Engine

Maximizes throughput with faster context switching, concurrent kernel execution, and improved thread block scheduling.

## Asynchronous transfer capability

Turbo charges system performance by overlapping data transfers with computation.

## High Speed, PCIe Gen 2.0 Data Transfer

Fast and high-bandwidth communication between CPU and GPU.

# Tesla

## SCALABLE SOLUTIONS

- Scalable from one to thousands of GPUs
- Available as computing processor and 1U rack-mount GPU computing system.

## SOFTWARE DEVELOPMENT TOOLS

- C language compiler, debugger, profiler, and emulation mode for debugging
- Standard numerical libraries for FFT (Fast Fourier Transform), BLAS (Basic Linear Algebra Subroutines), and CuDPP (CUDA Data Parallel Primitives)

CUDA™  
CODE

## TESLA ARCHITECTURE

- Massively-parallel many-core architecture
- 448 CUDA cores per GPU
- Integer, single-precision and doubleprecision floating point operations
- Hardware Thread Execution Manager enables thousands of concurrent threads per GPU
- Shared data memory enables processor cores to collaborate on shared information at local cache performance



# Benchmarks

## AMAX High Performance GPGPU Clusters

AMAX is a leader in the HPC field, specializing in GPGPU Cluster Solutions. AMAX is one of only a few vendors in world that is backed by two giants in the cluster market space: Nvidia and Intel. AMAX is an Nvidia Tesla Preferred Provider and is also part of Intel's Cluster Ready Program. Together, this combined partnership allows AMAX to offer HPC clusters that are optimized to your environment.

	1 Rack	2 Racks	4 Racks	8 Racks	12 Racks	24 Racks
# of CPUs	84	168	336	672	1,008	2,016
# of CPU Cores	504	1,008	2,016	4,032	6,048	12,096
# of GPUs	84	168	336	672	1,008	2,016
# of GPU Cores	37,632	75,264	150,528	301,056	451,584	903,168
Maximum Host System Memory Capacity (TB)	4.03	8.06	16.13	32.26	48.38	96.77
Maximum Tesla GPU Memory Capacity (GB)	504	1,008	2,016	4,032	6,048	12,096
Single Precision GPU Peak Performance (TF)	86.52	173.04	346.08	692.16	1,038.24	2,076.48
Double Precision GPU Peak Performance (TF)	43.26	86.52	173.04	346.08	519.12	1,038.24

Scientists and engineers have made the transition to 'many-core' computing because their problems have reached a complexity that required them to look for new ways to boost their work. The revolutionary scalar, single, and double precision floating-point performance enables solving a wide range of high-performance computing applications, whose complexity has outstripped the CPU's ability to solve them.



# Configurations

## Specifications:

AMAX offers several configurations of the Tesla preconfigured GPGPU clusters that are customizable to your needs. Depending on your application and requirements, our cluster specialists will work with you to create the most optimal configuration to suit your needs.

	ServMax PGC-4204	ServMax PGC-4208
Configurations	4 Tesla S2050/S2070 Preconfigured Cluster	8 Tesla S2050/S2070 Preconfigured Cluster
GPUs	16 Tesla 20-Series C2050/C2070 GPUs	32 Tesla 20-Series C2050/C2070 GPUs
CPU Servers	16 Dual-socket Quad-core Xeon E5620 2.4 GHz	32 Dual-socket Quad-core Xeon E5620 2.4 GHz
Head Node Server	2 Dual-socket Quad-core Xeon E5620 2.4 GHz	2 Dual-socket Quad-core Xeon E5620 2.4 GHz
Memory	48GB DDR5GPU memory w/ S2050, & 96GB DDR5GPU memory w/ S2070	96GB DDR5GPU memory w/ S2050, & 192GB DDR5GPU memory w/ S2070
Storage	Up to 32TB SATA drives	Up to 64TB SATA drives
Network Connectivity	QDR Infiniband on each node	QDR Infiniband on each node
Network Switch	Cisco 1U 20-port Gigabit Ethernet Switch	Cisco 1U 20-port Gigabit Ethernet Switch
Infiniband Switch	1U 36 4X QDR ports Infiniband switch	1U 36 4X QDR ports Infiniband switch
Rocks+ Cluster Rolls Software for CUDA	Rocks+ software with 1-year or 3-year support	Rocks+ software with 1-year or 3-year support

## Tesla™ Parallel High Performance computing (HPC) Specifications

- Up to 17x lower cost, 21x lower power consumption, and 60% less space than traditional 1U servers
- 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 fabric
- Real time InfiniBand diagnostics



## Complete Cluster Distribution

- Free IPMI management suite
- Seamless standard and custom application integration & cluster installation
- Firmware upgrades
- Multiple Interconnect Options: InfiniBand, Fibre channel, Ethernet (Gigabit, 10GbE)

## Clustered File Storage (from Terabyte to Petabyte)

- Hardware design & software stack
- Lustre / Open source file system (Redundancy across system nodes)

## Complete Cluster Assembly and Set Up Services

- CUDA compatible clustering software and developer tools with up to three years of support
- Configuration of cluster nodes and the network
- Installation of applications and computers to offer a true replacement of the customers' IT department

## Installed Software

The AMAX Tesla preconfigured cluster comes with preinstalled with

- Redhat Enterprise Linux 5.x, 64-bit
- CUDA 3.0 Toolkit and SDK
- Clustering software (Rocks+ Roll)

## GPU Software Development Tools

- Optional C-based software development tools and various libraries for GPUs.
- Optional NVIDIA® RealityServer® 3.0 3D web services software suite