

SGI with NVIDIA GPU Accelerators Complete Solutions

Delivering Application-Specific Acceleration

Features

Plug & Play GPU Solutions

Hybrid Solutions

High Performance Density

Accelerating Results with GPU Compute Solutions

SGI leads the industry in delivering application-specific acceleration, dating back to the Geometry Engine™ which accelerated graphics applications in the 1980s. SGI then co-developed the SGI Tensor Processing Unit (TPU), followed by RASC™ technology, FPGA's that were tightly-coupled to our shared memory architecture. With RASC technology, SGI created the world's largest single system image server with accelerators, to solve the most challenging life-sciences problems. With a full team of application experts, SGI has a unique position to help customers solve problems with GPU computing technology and has services and support personnel ready to help customers port and debug specific applications. SGI GPU solutions are integrated with SGI software and SGI InfiniteStorage™ to provide complete solutions for customer workflows. Both scale-up and scale-out GPU solutions are available to tackle any type of problem in scientific research, product development and homeland security.

Workgroup to Enterprise

SGI Rackable™ Servers: Leveraging the winning combination of the latest Intel® Xeon® Processor architecture and the NVIDIA® Tesla® NVIDIA Tesla GPU Accelerators, these servers deliver top value and performance. Rackable servers are fully managed and factory integrated with SGI Management Center and Performance Suite software for ease of administration and performance tuning.

SGI UV™ 30: A quad-socket Intel® Xeon® server, featuring rich memory and I/O in a compact, 2U footprint, the SGI UV 30 is a great solution for many scientific and engineering

workloads. It can be used as a head node or fat node for a scale-out cluster, or standalone. When paired with NVIDIA Tesla GPU Accelerators, the resulting compute and memory density is available for large models and complex optimization.

Supercomputer

SGI UV™ 3000 and UV™ 2000: Customers trying to solve the world's toughest computational challenges independent of the typical limits of CPU, memory and I/O inherent in most twin-socket or even quad-socket designs will find that the SGI UV platform will exceed their needs. The UV platform brings GPUs to a new class of solutions in chemistry, homeland defense, fluid dynamics and biosciences. The GIS Federal searchable geospatial database system GAIA, used by the Army INSCOM, incorporates numerous complex and distinct sources of information, which can be quickly sorted and displayed using easily understood visualization tools. Now running on a 10 terabyte (TB) SGI® UV™ 2000 system, with 16 NVIDIA® Tesla® K20X graphics processing unit (GPU) accelerators, GAIA is able to render complicated geospatial features for intelligent real-time insight, scaling at several orders of magnitude higher in speed and capacity than other compute databases. Through the increased capacity, scale and performance experienced with SGI's UV architecture, GAIA is able to more accurately pinpoint potentially dangerous activity at a given point or route on the map giving military personnel the opportunity to change course of action.



SGI UV™ 300 and UV™ 30EX: Newly enhanced SGI UV 300 and SGI UV 30EX servers are designed for data-intensive, I/O heavy workloads such as data analytics, visualization, and real-time streaming. Featuring Intel® Xeon® E7-8800 v3 processors and a NUMalink topology with ultra-low latency, these servers provide a greater memory to processor ratio and incorporate the latest NVIDIA GPU accelerators.

SGI ICE™ XA and ICE X: For customers who want to manage large scale-out HPC environments that include GPUs, the SGI ICE X platform offers the ability to integrate service nodes containing GPUs into dual-plane, high-bandwidth, low-latency Infiniband network topology. With the assistance of the SGI Professional Services team, SGI has implemented some of the largest hybrid clusters in the world by combining NVIDIA GPUs in service nodes with the SGI ICE platform.

Accelerating Customer Results

SGI GPU solutions accelerate customer results in a wide-range of scientific and engineering disciplines. With validation and integration done by SGI engineering, and systems built and tested in the SGI manufacturing facility, SGI GPU solutions arrive at the customer site ready to plug in and do real work.

“With the assistance of NVIDIA and Kepler series GPUs the Swinburne University of Technology supercomputers from SGI have proven to be excellent research tools in areas of astronomy ranging from simulations of the dynamical evolution of the Universe to the processing of data collected from radio telescopes,” said Dr. Jarrod Hurley, manager of Swinburne’s supercomputer.

Services and Support

SGI has a team of GPU experts who have ported code to both CUDA and OpenCL and are available on-site to accelerate applications in a wide range of technical disciplines. SGI Professional Services is available to integrate hybrid clusters either at the factory, so it reaches your floor ready for immediate availability, or at your site.

SGI GPU Compute Solutions at a Glance

sgi.com/products/gpu

Solution	Vertical "U"	Sockets	DIMM Slots	Coprocessor
SGI® UV™ 3000	10U blade enclosure	Intel® Xeon® E5-4600 v3 series	8	1x NVIDIA® Tesla® per MG blade/16x NVIDIA® Tesla® per SSI
SGI® UV™ 300	5U	Intel® Xeon® E7-8800 v3 series	96	4x NVIDIA® Tesla® per 5U chassis/16x NVIDIA® Tesla® per SSI
SGI® UV™ 30EX	5U	Intel® Xeon® E7-8800 v3 series	96	4x NVIDIA® Tesla®
SGI® UV™ 30	2U	Intel® Xeon® E5-4600 v3 series	48	2x NVIDIA® Tesla®
Rackable C2108-GP5	2U	Two Intel® Xeon® E5-2600 v3 series	24	8x NVIDIA® Tesla®
Rackable C2112-GP2	2U	Two Intel® Xeon® E5-2600 v3 series	24	2x NVIDIA® Tesla®
Rackable C1104-GP1	1U	Two Intel® Xeon® E5-2600 v3 series	16	3x NVIDIA® Tesla®
SGI® ICE™ XA	10.5U blade enclosure	Intel® Xeon® E5-2600 v3 series	8	4x NVIDIA® Tesla® per IP-139CS blade service node
SGI® UV™ 2000	10U blade enclosure	Intel® Xeon® E5-4600 v2 series	8	1x NVIDIA® Tesla® per MG blade/16x NVIDIA® Tesla® per SSI
Rackable C1104G-RP5	1U	Two Intel® Xeon® E5-2600 v2 series	8	3x NVIDIA® Tesla®
Rackable C2110G-RP5-P	2U	Two Intel® Xeon® E5-2600 v2 series	8	3-4x NVIDIA® Tesla®
Rackable C2108-RP2	2U	Two Intel® Xeon® E5-2600 v2 series	24	2x NVIDIA® Tesla®

Global Sales and Support: sgi.com

©2013-2015 Silicon Graphics International Corp. All rights reserved. SGI and the SGI logo are registered trademarks or trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries. All other trademarks are property of their respective holders. 19082013 4235 16102015

