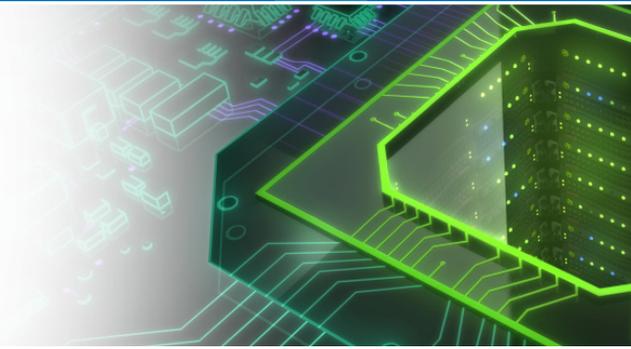


Delivering
Application-Specific
Acceleration

SGI® GPU Compute Solutions



KEY FEATURES

Solutions from the Desktop
to the Supercomputer

SGI high performance
GPU solutions optimized
to address the most
computationally-intensive
environments

Solutions for finding oil,
calculating complex
financial derivatives,
identifying a key nugget
of information for
intelligence-gathering, or
designing a new chemical
compound

Hybrid solutions that
combine traditional
microprocessors and
accelerators

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 the backing of a team of application experts, SGI is in a unique position to help customers solve problems with GPU computing technology. SGI has services and support personnel ready to help customers port and debug specific applications.

Desktop High-performance Computing

SGI® Octane III: SGI Octane III offers the capabilities of a high-performance cluster with the portability and usability of a workstation. It is available as a single or dual-node graphics workstation with support for the fastest NVIDIA® Quadro® professional graphics and compute GPU cards. Octane III is optimized for use with multiple display solutions for advanced visual computing scenarios.

Workgroup to Enterprise

SGI® Rackable™ Servers: Leveraging the winning combination of the latest Intel® Xeon® Processor architecture, NVIDIA graphics and Tesla™ 10 and 20 -series solutions, these servers deliver top value and performance. Rackable servers also offer the Design-to-Order model for great flexibility in a variety of chassis and motherboard formats.

Supercomputer

Altix® UV: 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 Altix UV platform will exceed their needs. The Altix UV platform brings GPUs to a new class of solutions in chemistry, homeland defense, fluid dynamics and biosciences. The Center for Remote Data Analysis and Visualization at the University of Tennessee recently installed Altix UV 1000 system with 128 CPUs, 4 TB of main memory and 8 NVIDIA GPUs to enhance the capabilities of the National Science Foundation (NSF) to 'see and understand' large volumes of data produced on the NSF's TeraGrid.

Altix® ICE: For customers who want to manage large scale-out HPC environments that include GPUs, the SGI Altix ICE 8400 platform offers the ability to integrate service nodes containing GPUs into dual-plane, high-bandwidth, low-latency InfiniBand networking topologies. 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 with the Altix ICE platform.



SGI® GPU Compute Solutions

Delivering Application-Specific Acceleration

www.sgi.com/products/gpu

Accelerating Customer Results

SGI has been working with the Irish Center for High-End Computing (ICHEC) on accelerating applications for GPUs. Founded in 2005, ICHEC was established as a national high performance computing (HPC) provider with offices in Dublin and Galway, Ireland. "Our focus has been placed firmly on developing new algorithms capable of exploiting the considerable power of GPUs. For example, recent results from our port of the DL_POLY MD application suggest that this technology indeed has the potential to make HPC accessible to a broader audience. We expect similar success with our ongoing port of Quantum Espresso," said James Slevin, director of ICHEC. "By partnering with SGI, ICHEC will provide business, scientific and academic customers with a high level of consultancy to support the development of effective computing solutions."

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

Solution	Vertical "U"	Sockets	DIMM Slots	NVIDIA GPU options
Octane III	Desktop	Up to 18S	12 per 'node', up to 9 nodes	NVIDIA Tesla C1060, C2050, C2070 (up to four GPUs, two per workstation)
Rackable C1103-TY12	1U	2S Intel Xeon 5600	12	NVIDIA Tesla, M2050, M2070, M2090, S2050, NextIO vCore Express S2070, S2090
Rackable 1104-2Y12	1U	2 x 2S Intel Xeon 5600	12 per node; 2 nodes	NVIDIA S2050, NextIO vCore Express S2070, S2090
Rackable C2112-4Y14	2U	4 x 2S Intel Xeon 5600	12 per node; 4 nodes	NVIDIA Tesla S2050 (up to two), NextIO vCore Express S2070, S2090
Rackable C3108-TY11	3U	2S Intel Xeon 5600	18	NVIDIA Tesla C1060, S1070, C2050, C2070, S2050, NextIO vCore Express S2070, S2090
Altix UV 10, Altix UV 100, or Altix UV 1000	3U enclosure or 18U enclosure	Up to 256S	Up to 16 TB memory	NVIDIA Tesla S2050 (up to eight GPUs in a single system image), NextIO vCore Express S2070, S2090
Altix ICE 8400		1 x 2S Intel Xeon 5600 or AMD Opteron 6100 up to 1000s	12 per blade	Rackable C3108 as GPU service node

Corporate Headquarters

46600 Landing Parkway
Fremont, CA 94538
tel 510.933.8300
fax 408.321.0293
www.sgi.com

Global Sales and Support

North America +1 800.800.7441
Latin America +55 11.5185.2860
Europe +44 118.927.8000
Asia Pacific +61 2.9448.1463