White Paper

SGI® Altix® Server Platform: Innovation Leadership for Real World High-Performance Computing

Table of Contents

1.0 Vanquishing the Challenges of High-Performance Computing

The demand for high-performance computing (HPC) is at an all-time high, fueled by a combination of lower hardware costs and the need to develop ever greater compute capacity to solve real-world problems. As performance affordability has grown, HPC has expanded beyond the realm of technical and scientific computing to embrace a wide variety of business computing needs as well. In both the scientific/technical and business domains, organizations face similar needs for exponentially increasing computing capability. At the same time that organizations seek to construct larger models using bigger datasets to achieve finer results granularity, business competition and pressing global issues dictate the need for ever faster time to discovery and time to market. The ability to implement and employ HPC capabilities rapidly and effectively has become crucial to many an organization's survival.

1.1 Overcoming the Roadblocks

As organizations seek to establish and scale their HPC capabilities, however, they find themselves grappling with issues both new and exacerbated, compared to their experience with conventional computing systems. These issues span the computing lifecycle, presenting formidable challenges for system set-up and configuration, ongoing system management, application development, scalability, and job control. To tame the power of HPC effectively, organizations must learn new skills and adopt new tools. Furthermore, the complexity of the hardware itself requires unprecedented attention to RAS (reliability, availability, serviceability). The growth in computing resources also places demands on the data center environment, pushing the limits of power, cooling, and space.

In order to be effective, organizations must select the HPC architecture that best matches their needs – high-performance computing encompasses several different system architectures, spanning both capability and capacity workflows. Different types of HPC applications optimize best under quite different architectures, as summarized in Figure 1.

To realize the full potential of HPC, organizations must meet the challenges:

- Choosing the right architecture at the right price for the specific need
- Minimizing set-up and configuration time
- Managing system complexity, particularly in regards to scalability
- Handling ongoing administration and application development complexities
- Minimizing impact on the data center environment (power/cooling/space)
- Maximizing reliability, availability, and serviceability

1.2 Technological Leadership in HPC

With over 25 years of experience tackling the most demanding compute and data-intensive problems, SGI® provides leadership in HPC technical innovation through its SGI® Altix® family of servers. The SGI Altix server platform encompasses a set of architectures to meet the diverse needs of high-performance computing, from capability to capacity, with systems that focus on both performance and economic requirements.

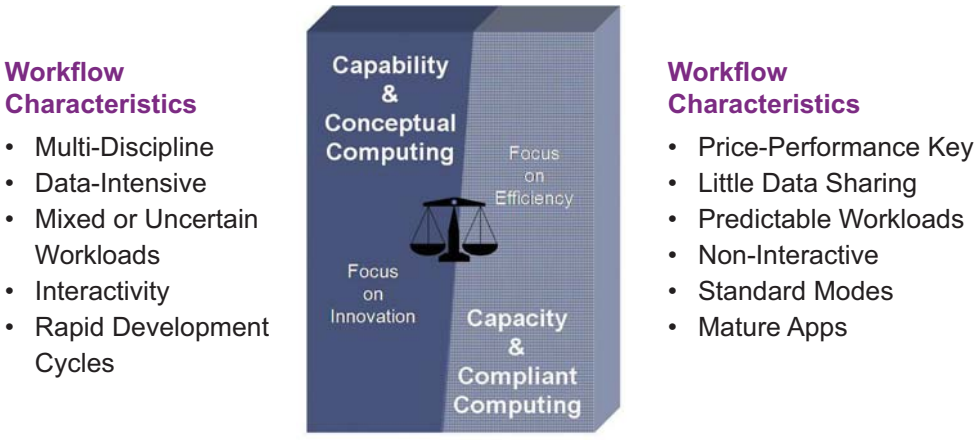

Figure 1. HPC Architecture/Application Characteristics

Figure 2. SGI Altix Server Portfolio

The Altix server family comprises the most advanced HPC platforms commercially available, providing solutions that maximize performance while minimizing overall cost. Altix servers meet the challenges of high-performance computing head on, with technology that addresses all key issues, including scalability, manageability, reliability, and data center constraints. Altix installations occupy prominent positions on the Top500 list of supercomputers worldwide. The same technology that drives these supercomputers also underlies SGI's HPC solutions for more modest workflows. SGI's technological leadership brings the power of high-performance computing to the real world.

Besides its expertise in developing HPC hardware, SGI brings its leadership in Linux® to bear on the Altix server line. SGI's proficiency in HPC software is evident in the comprehensive software suites that accompany its systems, as well as in the knowledge and experience of its Professional Services and Support organizations. SGI has long been a major contributor to Linux, developing key HPC enhancements that have been incorporated into the standard. SGI's expertise lifts open-source Linux to a level of reliability and functionality formerly obtainable only through proprietary Unix environments. The SGI Industrial Strength Linux Environment (SGI ISLE) offers a new standard for Linux capability in high-performance computing.

1.3 Systems Designed to Meet the Challenges

The SGI Altix 450 mid-range and SGI Altix 4700 supercomputer servers (collectively referred to as "SGI Altix 4X"), based on industry-standard Dual-Core Intel® Itanium® Series 9000 processors, have been designed to solve the world's toughest problems. They incorporate cutting-edge innovations such as unprecedented Linux scalability, global addressable memory up to 128TB, single system images that span up to 1024 core processors, and SGI® NUMAlink™, the fastest interconnect on the HPC market. The Altix 4X line of servers offers all of this, plus features that enhance RAS and minimize data center costs, in a modular, "plug-and-play" blade package. For scale-up, capability-focused computing, Altix 4X provides peerless power and versatility.

On the capacity side of the HPC spectrum, SGI offers cluster systems built on industry-standard Dual-Core and Quad-Core Intel® Xeon® processors: the SGI® Altix® ICE and SGI® Altix® XE platforms. Representing state-of-the-art in performance density, reliability, and efficiency, the Altix ICE platform is an integrated blade cluster system designed for the most demanding scale-out workloads. This innovative new product line from SGI employs a blade architecture designed by SGI specifically to meet the unique needs of the HPC market. SGI Altix ICE forges a radically new approach to high-performance computing. It delivers breakthrough scalability, manageability, reliability, and price/performance.

SGI Altix XE rounds out the Altix family, with a choice of economical cluster-targeted servers that lead the industry in price/performance and low TCO (total cost of ownership). Altix XE extends SGI's technological innovations to the value-focused segment of the HPC market, bringing the benefits of advanced, no-compromise high-performance computing to every budget.

SGI designs all its Altix product lines with a single goal in mind: to build systems offering the highest possible performance at the lowest overall cost. When SGI looks at cost, it considers the entire product lifecycle, including purchase, installation, application development, maintenance, data center impact, scalability, and upgradeability. To meet the constraints of today's data centers, the Altix family incorporates key innovations in the areas of power and cooling efficiencies and performance density. To accelerate discovery and keep a lid on TCO, all Altix systems ship with a comprehensive, configurable, standards-based software suite. To ensure future-proof scalability and upgradeability, Altix systems incorporate a variety of modular features. To maximize reliability and availability, the Altix product lines employ numerous technological innovations in key areas such as board and interconnect design, component redundancy, memory, and serviceability. To minimize start-up costs and ensure quick productivity, Altix systems arrive at the customer site factory-integrated and tested, with "power-up-and-go" functionality. SGI supplements its HPC servers with an extensive range of integrated storage solutions and industry-leading Professional Services and Support organizations.

With its Altix family of servers and clusters, SGI takes the lead in hybrid computing, offering a full range of systems to fulfill the precise requirements of any organization, whether the need be for capability or capacity computing, or for some mix of both. SGI simplifies the process of management and application development within a hybrid environment with a common software environment spanning the Altix family, including open-source Linux and the SGI ProPack software suite of optimization and enhancement tools.

Offering unmatched technological leadership, SGI designs systems to meet the real world challenges of high-performance computing:

- Solving the world's toughest problems – Systems to support vital design and research requiring the precise coordination of thousands of processors working in parallel.
- Addressing the hidden challenges of clusters – Scalability, system administration, application development, and managing overall costs.
- Easing the IT burden – Designing for data center efficiency and for the long term.
- Accelerating discovery through innovation – Providing cutting-edge technology to maximize the capabilities of HPC systems.

2.0 Solving the World's Toughest Problems

SGI offers the world's most advanced supercomputing platforms for both capability- and capacity-focused high-performance computing. In contrast to some specialty one-off systems from other vendors, SGI's supercomputers are production-built, off-the-shelf systems. Still, they consistently permeate the Top500 list of supercomputers. SGI systems are in service around the world, enabling universities, government research institutes, and businesses to pursue solutions requiring the highest performance and the greatest degree of reliability available.

Altix servers can be found operating in the most demanding environments worldwide, solving the world's toughest problems:

- *Massive Systems on Top500 Supercomputer List*

  A 14,336 core, 28TB SGI Altix ICE system, purchased for the **New Mexico Computing Applications Center**, hit the number 3 spot on the Top500 supercomputer list in November 2007. This achievement came less than six months after initial Altix ICE product introduction. The system, one of the world's largest dedicated to non-confidential projects, will fuel scientific and engineering breakthroughs for both private industry and public research institutions.

  At **NASA Ames Research Center**, a 4,096 core, 4TB Altix ICE system was up and running in only eight days. This system, ranked number 34 on the 11/07 Top500 list, serves NASA's Aeronautics Research Mission Directorate by supporting research in fundamental aeronautics, paving the way to transforming America's air transportation system and supporting the design of future air and space vehicles. "Whenever we can find ways to either shorten the time it takes for researchers to see results or to increase the level of fidelity of our simulations, we can speed the pace of discovery," said Juan J. Alonso, Director of the Fundamental Aeronautics Program at NASA. "The SGI Altix ICE system makes thousands more processor cores available for research, and will be especially useful for running simulation codes that are optimized for distributed cluster architectures."

- *World's Largest Shared Memory System*

  NASA Ames has also recently installed a 2,048-core, 4TB Altix 4700 supercomputer, creating the world's largest Linux single system image (SSI), spanning all 2,048 cores. The system generates 13.1 teraflops/second of compute power.

- *DoD's "HAWK" Takes Off On Schedule, Powered by 9,000+ Cores*

  The largest and most powerful computer at the **US Department of Defense (DoD)**, ranking 21st on the Top500, is an Altix 4700 supercomputer, powered by 9,216 Itanium 2 processor cores and featuring 20TB of globally addressable memory. The system, dubbed "HAWK," was operational within five weeks from the date it arrived at the Aeronautical Systems Center Major Shared Resource Center (ASC MSRC) at Wright-Patterson Air Force Base. Within its first few weeks of operation, it had already proved its worth running projects to study low-temperature turbulence, ocean modeling codes for enhanced weather prediction, airflow through turbine engines, and research aimed at developing next-generation materials for submarines and other military vehicles.

- *World Record for Oracle E-Business Suite Performance*

  An SGI Altix 450 server recently achieved the performance and scalability record for real-world, multi-tier applications running Oracle Database 10g. Running the **Oracle E-Business Suite 11i (11.5.10) Benchmark**, also known as the Oracle Applications Standard Benchmark (OASB), a 16-core, 128GB Altix 450 system outperformed the previous record holder by 70%, achieving an average response time of 0.453 seconds for 3,000 online users. The results also demonstrated that Altix 450 can reliably scale to accommodate more users and workloads with negligible impact on performance. As the benchmark load more than doubled from 1,400 to 3,000 concurrent users, the average response time of the Altix 450 slowed by just 2.5 percent. "Businesses and organizations are under a lot of pressure to respond quickly to change and keep up with demanding, real-time data-intensive computing requirements," said Boyd Davis, general manager, Server Product Marketing, Digital Enterprise Group, **Intel Corp**. "SGI systems powered by the latest Dual-Core Intel Itanium processor 9100 series deliver amazing scalable performance, mainframe-class RAS and overall great value for today's mission-critical enterprise solutions." The results were independently audited and verified on Oct. 18, 2007[i].

3.0 Addressing Cluster Challenges

Clusters, with their economical approach to high-performance computing, have been the main force driving the rapid expansion of HPC, with an installation growth rate of approximately 50% per year over the last five years. The ongoing acceleration of higher performance for lower cost is reducing HPC cost-of-entry and helping to bring the benefits of parallel computing to an ever-greater range of applications and organizations. Clusters' inherent modularity means that businesses can start small and build their HPC environments as their need grows.

As organizations that have been using clusters for some time will attest, however, initial hardware cost represents just a small portion of the total cost of ownership (TCO). The immediate advantage of low cost-of-entry, particularly evident for relatively small installations of just a few servers, can become eclipsed over time by ancillary costs. Several factors, in particular, can lead to high TCO costs:

- Cluster configuration and set-up effort
- System complexity, leading to difficulties in maintenance and scaling
- Ongoing management and application development complexities
- Data center environmental impact

3.1 SGI Altix XE: Superior Technology for Mainstream HPC

The Altix XE line of cluster solutions brings low TCO to high-performance computing. With SGI's innovative total lifecycle approach to product design and support, TCO continues to track low, even as HPC requirements scale over time. Altix XE clusters are proving their value in many ways:

- *Advanced Clusters for Customers Across a Wide Range of Industries, Research Fields, and Geographies*

  University of Minnesota has installed a 2,048-core SGI Altix XE1300 cluster for research in physical, biological, medical, mathematical and computing sciences, in addition to use in engineering studies and academic-industry collaboration.

  South Australian Partnership for Advanced Computing (SAPAC) is running a 544-core SGI Altix XE1300 cluster for use in computational chemistry research.

  The Biophysics Institute at Federal University of Rio de Janeiro has purchased a 112-core SGI Altix XE cluster for genomics research on HIV and malaria and to serve as a second leg of the university's grid effort.

  International Truck uses a 40-core Altix XE cluster to perform full-vehicle engineering analysis of new truck designs.

• *Hybrid Systems to Meet Diverse Needs*

The North German Group for High- and Highest-Performance Computers (HLRN) has selected a 25,000-core hybrid Altix ICE and Altix XE system for their new supercomputer complex, 60 times more powerful than their existing HPC resource. Such scalability is unprecedented in cluster-based supercomputing. When complete, the system will provide maximum performance of 312 teraflops/second, making it one of the world's fastest performing systems.

• *Winning the Supercomputing Cluster Challenge*

At the Supercomputing 2007 (SC07) conference in Reno, Nevada, a team from the **University of Alberta** won the first-ever Cluster Challenge with a 64-core, 144GB Altix XE310 cluster. In the Cluster Challenge, teams of undergraduate students assembled clusters on the exhibit floor and ran benchmarks and applications selected by industry and HPC veterans. Using the Altix XE cluster, the Alberta team beat five other teams who ran the same set of computations on competing platforms.

"This was an exciting competition that showed the deep pool of talent from our team," said Paul Lu, associate professor in the Department of Computing Science at the University of Alberta, and one of the coaches of the winning team. "The hard work and preparation of our team and its coaches paid off when the entire convention center lost power briefly on Tuesday. That forced us to completely restart the cluster and the computations. Fortunately, our Altix XE cluster was back up and running in about 10 minutes, which was not the case for all the other clusters. We were thrilled with the power and reliability of the SGI cluster."

• *Most Innovative Cluster Solution*

Also at SC07, **HPCwire**, the leading source for global news and information covering the ecosystem of high productivity computing, announced that its readers had selected SGI Altix ICE as the Most Innovative HPC Cluster Solution for 2007. HPCwire readers also honored Dr. Eng Lim Goh, SGI senior vice president and chief technology officer, with an HPC Community Recognition Award. Dr. Goh is widely known throughout the industry for his insights into how next-generation computer systems must meet the application price/performance demands of customers.

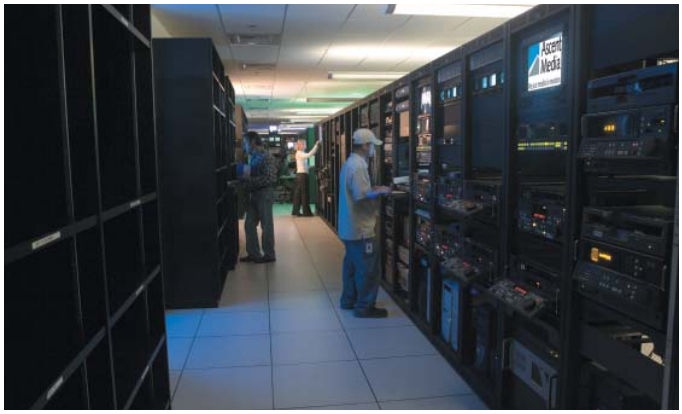

Figure 3. Data Centers Fast Reaching Capacity

4.0 Easing the IT Burden

4.1 Designing for Data Center Efficiency – Power, Cooling, Space

Today's data centers are in a state of crisis. AFCOM's Data Center Institute predicts power failures and power unavailability will halt IT operations at more than 90% of companies within the next four years[ii]. An industry survey by Emerson Network Power further projects that 96 percent of existing data centers will run out of capacity by 2011[iii]. Space is also an issue of great concern in today's data centers. Space availability is particularly critical in high-performance computing, where high-throughput applications can require co-location of large numbers of compute nodes within a single data center. The environmental inefficiencies of traditional clusters (low density, combined with high power and cooling requirements), barely noticeable on a small scale, become prohibitive when scaling larger installations.

The designs of SGI Altix systems take into account the environmental constraints of today's data centers and offer high efficiency solutions to meet the challenges. SGI has long been a technological leader in solutions optimizing power and cooling efficiency, designing a power architecture for the Altix 4X and Altix ICE systems that features highly efficient power supplies and a third-generation water-based cooling system. Together, these innovations represent a major advance in overcoming the constraints of the data center. SGI combines 90% efficient power supplies with high efficiency voltage regulators on the blades, to minimize power losses throughout the system. In addition, Altix systems employ a combination of high-efficiency blowers and optional water-cooled rear doors to deliver impressive cooling efficiency results. While actual performance depends on many site and geographic variables, the water-cooled doors can reduce cooling equipment power consumption by 48% or more.

Figure 4. SGI 3rd-Generation Water Cooling

To deal with the third constraint of the data center, space availability, SGI designs its systems for optimal performance density. The entire Altix family employs numerous innovations to attain maximum compute power in minimal space. Altix ICE, for example, achieves a density of 512 Intel Xeon processor cores per rack. In raw compute power, this equates to 6 teraflops per 42U rack. Despite this level of performance density, the fully loaded rack stays within data center flooring constraints, with a footprint of 246 lb/ft^2.

With innovative approaches to power, cooling and performance density, SGI Altix stands in sharp contrast to most HPC computing platforms. These innovations ensure maximum utilization of scarce data center resources.

4.2 Designing for the Long Term – Reliability, Availability, Serviceability

With the massive increases in system memory capacity and processor count, combined with continuing growth in job sizes and datasets, issues of reliability, availability and serviceability (RAS) are becoming increasingly critical to high-performance computing. Even as SGI pushes system sizes to new levels, the company delivers new levels of RAS to high-end Linux environments. SGI's RAS efforts leverage its unique experience building the world's largest and most robust HPC systems in combination with focused investments to enhance memory reliability for large shared-memory environments.

SGI's focus on reliability starts at the component level, with stringent attention to the design, manufacturing, and testing of components. Reliability in Altix 4X is further enhanced by using components that automatically detect and correct errors, such as Itanium processors with built-in RAS capabilities, system buses and memory with error-correcting code, and cyclic redundancy check technology on all NUMAlink channels between Altix blades. Similarly, Altix ICE employs fully buffered DIMMs to reduce transient errors and an InfiniBand backplane for high signal reliability. In addition, the Altix ICE compute blade is diskless, thereby eliminating one key point of potential failure.

SGI also engineers the overall hardware system to enhance reliability. Altix 4X and Altix ICE racks, for instance, contain redundant, hot-swappable power supplies and blowers. SGI's high-efficiency power and cooling technologies, described earlier, also enhance reliability by ensuring clean power, reduced heat dissipation, and optimal operational temperatures. Altix ICE gains further boosts in reliability by employing blade enclosures that are essentially cable-free, in contrast to the cable clutter endemic to most cluster systems. Finally, the focus on performance density in all Altix systems means less components and less overall potential for failure.

Availability begins with the focus on reliability, but extends to additional features to maximize up-time. SGI Altix servers are designed to boost availability throughout the system. Altix servers incorporate a host of availability features, including independent reboots of partitions – useful for maintaining partial system available during hardware repairs, rolling kernel updates, and restart of test environments. Altix 4X servers also boost availability with memory enhancements, such as flawed page detection/avoidance capabilities and failure analysis tools, and highly redundant fibre-channel storage infrastructures.

Finally, serviceability is a key SGI differentiator in the high-performance computing marketplace. The advanced, modular blade designs of Altix 4X and Altix ICE servers enable administrators to access and replace components easily. To quicken resolution of service issues, all Altix systems ship with the SGI Embedded Support Partner – an integrated suite of services that work together to monitor and manage systems and protect against problems. In addition, the SGI customer service organization offers a broad range of support offerings, up to and including mission-critical 24x7 system support. Dedicated to excellence, the SGI support team consistently ranks among the top in the industry, according to the SatMetrix™ third-party evaluation.

Reliability, availability, serviceability – these vital qualities are inherent in the design of every Altix product.

5.0 Accelerating Discovery through Innovation

The key driver of high-performance computing is its ability to accelerate discovery – whether the goal of discovery be to investigate climate trends, to develop new drugs, or to mine vast data warehouses for real-time business intelligence. SGI has a reputation for continual innovation in technologies and processes, constantly raising the bar of HPC capability. SGI is dedicated to accelerating discovery by delivering the fastest, most capable systems to its customers.

5.1 Systems that Deploy Quickly and Easily

To accelerate discovery, SGI designs systems that deploy and configure easily. To this end, Altix systems undergo intensive integration and testing at the factory, before delivery to the customer. As a result, even the largest systems become productive in a matter of days, rather than months.

- *Power-Up-And-Go Functionality*

  The Altix ICE blade cluster solution is designed for extremely fast deployment. At **Idaho National Laboratory**, a 2,048-core, 4TB Altix ICE cluster system was up and running in just a day and a half. “It was impressive to see how quickly our SGI Altix ICE cluster was up and running,” said Peter Cebull, HPC User Consultant, Idaho National Laboratory. “Even though this new Altix ICE will provide four times the compute capability of the system it’s replacing, it has been extremely easy to deploy and manage. We plan to leverage the ICE platform’s integrated InfiniBand interconnect to scale our applications across more cores and achieve far more detailed simulations.”

  At **NASA Ames Research Center**, a 4,096 core, 4TB Altix ICE system was up and running in only eight days. The system ranks number 34 on the Top500 list.

- *9,000+ Cores, Operational in Five Weeks*

  On the capability side of the supercomputing spectrum, SGI also designs its Altix 4X systems for rapid deployment. Recently, SGI deployed the largest and most powerful computer at the **US Department of Defense**. HAWK, ranking 21st on the Top500, is an Altix 4700 supercomputer, powered by 9,216 Itanium 2 processor cores and featuring 20TB of globally addressable memory. The system was operational within five weeks from the day it arrived at Wright-Patterson Air Force Base.

“With many multimillion-dollar supercomputers staying technologically current for only three to five years, spending 10 to 20 percent of that time on just the start-up deployment process can be costly. Consequently, from scale-out clusters to shared-memory supercomputers, SGI has invested heavily in the development of rapidly deployable solutions,” said Dr. Eng Lim Goh, SGI’s senior vice president and chief technology officer. “It’s one thing to throw thousands of processors at a problem, but it’s something else entirely to deliver a solution that is productive within days, or even hours, after it rolls off the truck.”

5.2 SGI® RASC™: Innovative Application Acceleration with Reconfigurable Computing

SGI RASC (Reconfigurable Application Specific Computing) technology leverages the power of FPGAs (field programmable gate arrays), which use gate array technology that can be reconfigured by the user for optimal performance on a specific algorithm. Users whose applications spend a majority of their run time on a set of specific algorithms can dramatically accelerate time to discovery by custom-configuring the RASC. This reconfigurable technology is particularly beneficial when running data-intensive applications critical to oil and gas exploration, defense and intelligence, bioinformatics, medical imaging, broadcast media, and other data-dependent industries.

- *World’s Fastest FPGA*

  SGI offers the world’s fastest FPGA with its RC100 blade for Altix 4X systems. RC100 blades can lead to orders of magnitude application acceleration, providing the same results as a supercomputer cluster but at a fraction of the cost and power consumption, and with a much smaller form factor. SGI has recently extended the reach of its RASC technology with its new RC200 for Altix ICE and Altix XE cluster systems. The RC200 plugs directly in a Xeon dual-core socket, bringing the power of RASC to cluster-based mainstream HPC.

- *World’s Largest FPGA Supercomputer*

  SGI recently built the world’s largest FPGA supercomputer, featuring 70 FPGAs. This accelerated the performance of a complex **BLAST-n query** by more than 900 times, completing in less than 33 minutes what took a 68-node Opteron-based cluster approximately three weeks to finish[iv]. The application matched 20 nucleotide base pairs against 600,000 queries. SGI configured the system using only off-the-shelf components, including its SGI RASC appliance for bioinformatics—featuring Mitrion™-Accelerated BLAST-n. No hardware or software was modified for the test[v].

5.3 SGI® NUMAflex®: Global Shared Memory for Data-Intensive Computing

SGI offers NUMAflex, the industry's most scalable memory architecture, with up to 128TB of globally addressable memory in an Altix 4700 system, creating what is sometimes referred to as a "supercluster." Even much smaller Altix 450 systems can be configured with up to 854GB of shared memory. These figures stand in stark contrast to clustered architectures, where addressable memory is typically limited to 32GB. SGI manages to achieve this breakthrough at exceedingly low cost, with memory-only blades that plug directly into the Altix chassis. To enable global shared memory, Altix 4X systems use SGI NUMAlink, the fastest interconnect on the market with 6.4GB/s bandwidth and <1 microsecond latency.

With SGI's memory architecture, massive in-core computation and memory-mapped I/O lead to breakthrough results. Entire databases can be held directly in memory. It becomes possible to model complete systems without the need for reduced resolution or precision and without breaking a problem apart or using an emulation. Furthermore, some applications, such as NWChem and Nastran, only scale with system memory, not with increased processor counts.

Many different types of applications benefit from the power of Altix global shared memory capabilities:

- *Bioinformatics* – A UK-based pharmaceutical research institute has developed a genomics search and matching algorithm that produces results up to 1000x faster than the commonly used BLAST application. This algorithm relies on 192GB of memory available on each of two SGI Altix servers with 4 and 16 processors.

- *Massive structural analysis* – **ANSYS** announced recently that it had become the first engineering simulation company to solve a structural analysis model with more than 100 million degrees of freedom (DOF), making it possible for their customers to solve full resolution models of aircraft engines, automobiles, construction equipment and other complete systems. The 111 million DOF structural analysis problem was completed in only a few hours using only a handful of processors on a large memory SGI Altix computer.

- *Large-scale seismic dataset interaction* – **Marathon Oil** has met the challenge to deliver larger and larger seismic data to the interpreting geoscientist. For the first time, geoscientists can visualize and interact with more than 400GB of seismic data in real time. This is more than 4 times the previous record[vi].

- *Rapid response data analysis* – A US government agency is using 32-processor Altix systems equipped with 4TB of memory as the only solution allowing them to handle the massive data mining and search algorithms essential to their work.

- *Memory resident databases* – **Xcelerix** IMDB and database accelerator has demonstrated orders of magnitude performance gains over traditional, disk-resident databases on Altix systems for databases containing up to 500 Million records (120GB).

The largest shared-memory supercomputer in Latin America is an Altix 4700, recently installed at Brazil's **Universidade Federal do ABC** (UFABC), in Sao Paulo. The 136 processor core, 272GB system will be used to run data-intensive and computationally intensive applications in nanoscience, materials science, engineering, chemistry, and physics. "We have some applications, mainly for materials science and chemistry, that perform better on computers that have a very fast interconnect, and SGI is one of the best, if not the best, in the world, with their NUMAlink interconnect," said Gustavo Dalpian, Assistant Professor, Center for Natural Sciences, UFABC. "SGI can give us the speed we needed. Also, the ability of using all the memory in the Altix in a single process, running very large applications that demand a lot of memory, was another important part of our choice."

6.0 Summary – Future-Proof Computing with SGI Altix

The SGI Altix line of servers, with systems ranging from small clusters to some of the world's fastest supercomputers, offers solutions to facilitate the diverse spectrum of HPC requirements. Whether the domain is technical, scientific, or business, SGI delivers products with unsurpassed price/performance, along with crucial technological innovations to ensure that the systems excel across the entire range of application needs, now and well into the future, from delivering economical departmental research capability to solving the world's toughest problems.

7.0 About SGI

SGI (NASDAQ: SGIC) is a leader in high-performance computing. SGI delivers a broad range of high-performance server, storage and visualization solutions along with industry-leading professional services and support that enable its customers to overcome the challenges of complex data-intensive workflows and accelerate breakthrough discoveries, innovation and information transformation. SGI helps customers solve significant challenges whether it's enhancing the quality of life through drug research, designing and manufacturing safer and more efficient cars and airplanes, studying global climate change, providing technologies for homeland security and defense, or helping enterprises manage large data. With offices worldwide, the company is headquartered in Sunnyvale, Calif., and can be found on the Web at sgi.com.

[i] *All OASB results are available at: http://www.oracle.com/apps_benchmark/html/results.html*

[ii] *"Five Bold Predictions for the Data Center Industry That Will Change Your Future", AFCOM's Data Center Institute, Data Center World Conference, Atlanta, GA, March 2006.*

[iii] *Bob Bauer, Emerson group vice president and Liebert Worldwide president, previewed the results of a survey conducted by the Data Center Users' Group during his keynote address at Emerson's AdaptiveXchange 2006™ (pr dated 11-16-06). See http://www.liebert.com/bottom_news2.asp?id=2386.*

[iv] *Results compared to industry-standard Opteron processor-based system measured in internal tests. Running a released version of BLAST, SGI used a test case from Affymetrix comparing approximately 600,000 queries with a query size of 25 base pairs against the Human Unigene and Human ReSeq databases, which is representative of current top-end research in the pharmaceutical industry. Total execution time on a traditional 68-node Opteron-based cluster would be approximately three weeks. On the SGI reconfigurable supercomputer, benchmark input data was split in 169 jobs, which were run in groups of 70, 70 and 29 FPGAs. Total wall clock time for the run was 32m:29.183s, representing a 916X speedup over the 68-node traditional cluster.*

[v] *The tested SGI configuration consisted of 35 dual-FPGA SGI® RC100 RASC blades, a 64-processor SGI Altix 4700 with 256GB of globally addressable memory, and standard SUSE Linux Enterprise Server 10 SP1 (kernel version 2.6.16.46-0.12) running an unmodified release of RASCAL, SGI's RASC Abstraction Layer.*

[vi] *See: www.lgc.com*

Corporate Office
1140 E. Arques Avenue
Sunnyvale, CA 94085
(650) 960-1980
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5185.2860
Europe +44 118.912.7500
Japan +81 3.5488.1811
Asia Pacific +61 2.9448.1463

4080 [03.03.2008] **J15361**