Solution Brief

SGI® High-performance Solutions

SGI Maximizes Productivity and Minimizes TCO

Bringing High-performance Open Systems to the Datacenter

From its inception, SGI has been about one thing: BIG DATA. Whether crunching it, storing it or visualizing it, SGI systems have always excelled at helping people turn mountains of data into knowledge.

Where there is data, there tend to be databases, so it is no coincidence that SGI systems are extremely well suited to running the data management software people use to manage overwhelming amounts of data. Whether the application involves online transaction processing or data warehousing/business intelligence, industry standards-based SGI platforms offer outstanding performance and total cost of ownership for users with demanding database challenges.

SGI® Altix®: A Server Architecture Optimized for Databases

Generally speaking, there are three ways to make a database run faster. First, you can put more of the data in main memory, reducing the number of disk reads and writes. Second, you can improve I/O and memory bandwidth, reducing the time needed for each disk access that is required. And third, you can add more processors, to run the database software faster and to enhance performance of the applications sitting atop the database. Designed to efficiently handle massive datasets and complex computational problems by enabling unsurpassed scal-ability along each of these dimensions, SGI Altix is uniquely well suited to hosting high-performance database applications.

SGI® NUMAflex™

The SGI Altix server family is based on SGI's ultra-high performance NUMAflex shared-memory architecture. This architecture provides unmatched scalability and the ability to expand system resources independently to efficiently address specific processing requirements.

SGI Maximizes Productivity and Minimizes TCO

For example, a single Altix node, running one copy of the Linux® operating system, can scale up to 608GB or 24TB of shared memory respectively for mid-range Altix® 450 servers and high-end Altix® 4700 servers. Effectively, Altix can accommodate as much of the database in memory as the software itself will permit. Altix is unmatch-ed in terms of I/O and memory bandwidth, having demonstrated sustained I/O bandwidth of over 7 gigabytes per second and achieved a world-record 1 terabyte per second result on the STREAM[1] Triad memory bandwidth benchmark. All this bandwidth means each I/O operation requires much less time to complete, for a dramatic boost in bottom line performance. Altix is also unmatched in processor
scalability, supporting from 1 to hundreds of processors per node, all operating against a pool of global shared memory. This unique ability to independently
scale compute resources in an ultra-high performance architecture results in a cost-effective data- base platform with the headroom to accommodate virtually any processing needs.

Intel® Itanium® 2 Processor Architecture

SGI Altix servers are built on the Dual-Core Intel Itanium 2 processor, which features the Explicitly Parallel Instruction Computing (EPIC) architecture. As a database engine, Itanium 2 processors offer many advantages over other processors in the market, including proprietary RISC processors. EPIC features a very large (currently up to 24MB) cache located on the processor itself, which allows Itanium 2 processor-based systems to hold database records[2] in cache for the entire duration of a transaction, enabling the I/O portion of the transaction to take place at speeds faster than memory access.

The design characteristics of Itanium 2 also make it ideal for performing the complex floating point and integer calculations required for database security. Encryption and decryption are used to authenticate database users and for data protection. The greatly expanded arithmetic capabilities of Itanium 2 (enabled by multiple floating point, memory, integer, and branch execution units, and other advances) permit encryption/decryption on the fly, in real time.

The overall efficiency of the Altix server also provides unmatched application server consolidation. Tests have shown that even when doubling resources Altix delivers double the performance[3].

SGI InfiniteStorage

The NUMAflex architecture ensures that the processors will never be starved – *if* storage performance is high enough to keep up with the system architecture. SGI InfiniteStorage solutions are designed to do just that. InfiniteStorage solutions range from SATA to fiber channel arrays and include award-winning NAS, SAN and Data Lifecycle Management technologies.

Altix – A Breakthrough Database Solution

Reduced Total Cost of Ownership

With Altix, you can scale system resources such as memory, I/O, and processors independently of one another. For example, a user who needs 2 processors and 128GB of memory to run their application optimally can configure an Altix server to those precise specifications. With most other systems on the market, one would have to purchase far more than 2 processors to get 128GB of memory. By allowing users to buy only the resources their workflow actually demands, Altix slashes hardware cost. And since database software is often priced per processor, Altix can slash software costs as well.

Best of all, Altix allows you to start small and grow as your requirements grow, greatly reducing or eliminating the need for “forklift upgrades.” The Altix server can easily be expanded or upgraded with additional memory, processors or I/O, to best address processing needs. And the various system resources are all connected not by a contention-prone bus, but by the 6.4GB/ second NUMAflex fabric, putting an end to back-plane bottlenecks.

State-of-the-Art Algorithm Optimization:
Pushing the Performance Envelope

To further boost the performance of frequently executed algorithms (and database queries), SGI recently

released the Reconfigurable Application Specific Computing (RASC) subsystem for Altix that allows Field Programmable Gate Arrays (FPGAs) to attach directly to the architecture's 6.4 GB/s interconnect fabric. This technology can provide substantial performance improvement for database applications that spend a majority of their run time on a set of specific algorithms. Some examples include text search and pattern recognition routines, data mining, and very large database solutions that require sub-second response times.

Configuration Flexibility

While many users run their database with a single copy of the operating system ("single system image"), vendors like IBM® Information Management and Oracle® are increasingly emphasizing clustered versions of their database software for enhanced availability and scalability. Altix is unmatched in its ability to scale "up" and it can scale "out" just as easily. The SGI® Altix® XE product line compliments the Altix server product line, and is based on the Intel® Xeon® Dual and Quad-core processor architecture. Altix XE provides customers with an advanced platform on which to deploy scale-out database configurations, delivering top value and price/performance with advanc-ed features that include a super-fast 1333MHz front-side bus, up to 32GB of memory per compute node, fully-buffered DDR2 memory and ultra-dense architecture that optimizes performance density for the data center. Altix XE clusters are available in a choice of two packages, designed to optimally meet diverse customer needs.

The SGI Altix XE1200 cluster offers advanced extensibility, with a rich set of expansion and I/O options to address the broadest range of compute requirements. Based on an innovative new board design, the SGI Altix XE1300 supports up to 16 processor cores per 1U package, delivering industry-leading performance density and cost efficiency. Altix XE clusters are the ideal platform for Oracle® Database 10g Real Application Clusters (Oracle10g RAC) and IBM DB2® Universal Database with Data Partitioning Feature.

The SGI Altix and Altix XE server line provides customers with the broadest set of options, enabling customers to build and deploy database configurations that optimally address current and future requirements, with "fat" or "thin" cluster nodes to efficiently address specific workload requirements. For example, an Altix XE cluster can be configured with hundreds of nodes, with node size ranging from single socket (two or four processor cores) to four sockets (eight or sixteen processor cores) per node. Or, like NASA Ames, you can configure an Altix cluster with twenty nodes, with each node containing 512 processors, multiple terabytes of shared memory, and a single copy of the Linux operating system.

While some database applications lend themselves to cluster platforms, others are more suited to a shared-memory approach. Shared-memory systems are better able to handle environments where many concurrent users must access multiple applications that share a single data repository. A shared-memory approach also enables easier system and software implementation and management.

Because of its unique shared-memory scalability, Altix delivers efficient management of large data sets, access to massive data sets, and execution of more instructions per cycle, all of which boosts overall performance.

Data Availability and Protection

Data protection, data integrity, and business continuance are essential to every user. SGI Altix supports a range of leading solutions like those from EMC® Legato and SteelEye Technology® Inc.

SGI Altix Database Solutions

Data Warehousing/Business Intelligence

These critical applications turn vast pools of data into knowledge or "insights" that can be used to make decisions. The ability to make better, faster decisions shortens overall "time to insight." Altix is well suited to the task. Data warehousing and business intelligence applications are extremely read-intensive, so the more data that is held in memory (rather than on disk), the better performance, and the shorter the time to insight. With its massively scalable shared memory, Altix can accommodate even the largest data warehousing problems at maximum theoretical speeds.
The scalability of Altix also makes it the perfect platform for running demanding data mining and business intelligence applications such as SAP® Analytics, IBM DB2 Intelligent Miner®, and SAS® Enterprise BI Server.

Altix and IBM

Altix is an outstanding platform for running DB2 Universal Database Server because the DB2 architecture – like the Altix architecture – supports essentially unfettered scalability. Altix is also extremely well suited to run-

ning IBM DB2 Entity Analytic Solutions (EAS), one of the most demanding and powerful database applications in the world. EAS is used by commercial and homeland security customers to uncover hidden relationships that signal fraudulent or improper behavior. The EAS data repository is a real time operational data warehouse that accepts multiple data sources, processes the data at high ingest rates, "resolves" entities, uncovers hidden relationships among entities, and becomes smarter as more records are added to the database. The ability to rapidly ingest and process large amounts of data, and efficiently execute complex, iterative queries on that data makes Altix an outstanding platform for the demanding requirements of this type of innovative database solution.

Altix and Oracle

SGI Altix servers provide a unique opportunity for Oracle users to scale up, scale out, or scale up and out as required to meet their specific workflow objectives. Some problems lend themselves to tall, single system image machines with large pools of shared memory, compute power, and I/O bandwidth. Other problems are more suited to a grid approach using the popular Oracle Real Application Clusters technology. Either way, users can scale their system – or their grid nodes – to whatever size best fits the task. In some cases, a mixed architecture that employs clusters of "fat" nodes may outperform either a single system image or standard "thin" node cluster strategy. Whatever the requirements, Altix offers the scalability and flexibility to meet the challenge.

Altix and SAP

With Linux gaining momentum worldwide, customers continue to adopt Linux as their preferred platform for enterprise application deployment. Enterprises that rely on SAP solutions can cost-effectively consolidate their applications on the award- winning SGI Altix server for unsurpassed scalability, performance, and robustness on Linux. The powerful combination of SGI Altix servers and SAP applications allows SAP users to leverage the same scalable and high-throughput SGI® servers already adopted by many of the world's leading automakers, energy giants, pharmaceutical companies, and manufacturers. SAP applications running on SGI systems can enable the "real-time enterprise" where IT bottlenecks evaporate and critical business decisions can be made on the fly at a lower TCO than customers see today.

Altix and Databases: A Natural Fit

A handful of market trends are redefining database requirements across industries and organizations: migration from proprietary operating systems to Linux, consolidation of database servers, and pursuit of the "real-time enterprise" in which both raw data and knowledge ("insights") are immediately available to everyone who needs them. The SGI Altix family of high performance servers provides a unique solution to this challenge. Whether you are running an operational system, a data warehouse, or a mixed environment, Altix delivers critical business benefits: real-time response rates, fast time to insight, low total cost of ownership, and extreme flexibility – all on an open, industry standards-based platform.

---

[1] SGI Altix ranks the highest among all posted systems in the STREAM Triad Top 20 list of shared-memory systems. Details on all STREAM Triad results are available at http://www.cs.virginia.edu/stream/top20/Bandwidth.html.

[2] The Intel® Itanium® Architecture Comes of Age can be found at http://www.intel.com/pressroom/kits/events/enterprise_server/itanium_ecosystem.pdf.

[3] SPECjbb2005 competitive results, ranked here by order of total throughput metric, were accessed from http://www.spec.org on Aug. 23, 2006. Main configuration details shown above, but complete data are given at http://www.spec.org/jbb2005/. SPEC and SPECjbb are registered trademarks of Standard Performance Evaluation Corporation.

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

North America +1 800.800.7441
Latin America +55 11.5509.1455
Europe +44 118.912.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000

3929 [01.25.2007]

J15148