White Paper

SGI® Database Accelerator Solution
Increasing the Scalability of Real-Time Transaction Processing

Table of Contents

1.0 Ever Increasing Demand for Real-Time Transaction Performance

Many government agencies, telecommunications carriers, and financial institutions optimize results by utilizing real-time systems. Military missile systems continuously monitor targets and respond instantaneously to intercept threats. Telecommunications carriers query databases for credit authorizations in real time before delivering on-demand content services to subscribers. Financial institutions rapidly execute trades in volatile capital markets to gain desirable positions in financial equities for clients. While traditional uses for real-time processing remain important, an increasing number of applications within commercial enterprises are now emerging with similar performance requirements (Table 1).

Industry	Performance-Critical Applications
Government	•Military systems •Central intelligence •Aerospace defense systems •Simulation and training •Real-time fraud detection
Networks	•Network security •Fraud detection •Transaction processing •e-Commerce •Web portals •Web services
Telecommunications	•Service authorization •Revenue assurance •Network and QOS management
Financial services	•On-line trading •Smart order routing •Position keeping •Real-time analytics •Real-time fraud detection
Commercial enterprises	•Service oriented architectures •Call centers •Online personalization •Business intelligence •Business activity monitoring •Operations and logistics •Information management •Customer relationship management (CRM) •Supply chain management

Table 1. Examples of industry-specific and commercial enterprise performance-critical applications

Demands for personalized on-line interactions, instantaneous delivery of information, and self-service portals are driving requirements for better response times and pushing the limits of current enterprise architectures. Many organizations also find that current call center and knowledge base systems must deliver faster results in order to facilitate better customer service. In addition, more effective and rapid identification of new revenue opportunities requires higher processing throughput from business intelligence and business activity monitoring applications.

In fact, whole classes of data-intensive applications, including fraud detection, biometrics, and social networking can greatly benefit from near real-time transaction speeds and result in dramatic security improvements for enterprises and government agencies. However, the effective manipulation of massive amounts of data in real-time provides challenge. Queries against large data sets eventually encounter long latencies for disk data retrieval, excluding any possibility of real-time analysis.

1.1 In Pursuit of Real-Time Performance

To accomplish real-time processing for small data sets, many enterprises rely upon custom or in-house solutions. Customized architectures often insert a special data layer that is optimized for rapid query response times. Unfortunately, creating and maintaining the data layer adds complexity to the environment. As another option, commercial object-oriented databases and proprietary systems can help enterprises meet real-time requirements. However, specialized software and hardware solutions often lack flexibility and lead to increased management and maintenance burden.

Seeking an alternative strategy, a number of organizations employ resourceful methods to improve the performance of general-purpose systems and standard relational databases. While installing faster processors and larger memory configurations provides some gains, true performance-critical applications require massive improvements in throughput. Utilizing query results caches, creating large database caches in the application tier, and co-locating applications and database services can help speed database operation. Unfortunately, caches quickly become out-of-date and co-location of applications with the database can cause system tuning conflicts, limiting the effectiveness and applicability of these solutions.

At the root of the problem, disk-based storage remains a barrier to achieving desired performance levels. Indeed, application processing can stall, waiting for key data to be transferred from disk. Employing large memory systems to keep entire data sets resident in memory often provides improvement by significantly reducing data access latency times. However, disk-based database algorithms often are not properly architected for in-memory operation. In a scenario where data is fully cached, traditional database technology performance must implement memory buffer management mechanisms and support multiple data locations, resulting in increased overhead. Despite all practical efforts, the throughput requirements of performance-critical applications often reach well beyond the capabilities of traditional disk-based relational database systems.

1.2In-Memory Database Technology

Forming the foundation of an ideal solution for many performance critical applications, in-memory databases reduce I/O latency by holding all data in memory at runtime and utilizing disk storage only for persistence and recovery functions. In-memory databases are constructed from the ground up to take full advantage of memory-resident data sets. For example, query execution, buffer pool management, and index structures are all architected with memory-resident storage in mind. By removing the need to track data locality between disk and memory storage, the number of machine instructions required for each operation is dramatically reduced, buffer pool management and extra copies of data become unnecessary, and index pages are greatly simplified and reduced in size.

Such performance improvements enable in-memory databases to process tens to hundreds of thousands of transactions per second, even on systems with a small number of processors. In contrast, the performance of traditional disk-based databases only recently broke the 60,000 transactions per second barrier1. Once positioned as a high-end specialized solution, in-memory databases are now viable for wide spread adoption due to lower memory prices, broader adoption of 64-bit operating systems, and the availability of standards-based in-memory database technology from Oracle.

2.0 Introducing the SGI Altix and Oracle TimesTen In-Memory Database Solution

Hosting Oracle TimesTen In-Memory Databases on SGI Altix servers creates an exceptional solution for meeting the demands of performance-critical applications. The Oracle TimesTen In-Memory Database enables real-time database transaction performance in an architecture based on open-standards, easing adoption of this technology by commercial and government organizations. By choosing SGI Altix servers for execution of an Oracle TimesTen In-Memory Database, organizations can take advantage of unparalleled memory scalability, the powerful SGI Altix NUMAflex® architecture, robust availability features, pay as you go upgrade options, and the affordability of a platform based on industry-standard technologies. Combining high performance, scalable SGI Altix servers with the responsiveness of an Oracle TimesTen In-Memory Database creates a solution that improves the economics and flexibility of rapid transaction processing as compared to proprietary and custom real-time architectures.

2.1 Scalable SGI and Oracle TimesTen In-Memory Database Solutions

Traditional real-time systems often support the analysis and querying of less than 100 gigabytes of data. However, many modern enterprises are data rich. On-line business transactions stream continuously, creating databases of information that seem too immense to properly analyze. By utilizing scalable SGI Altix servers to host Oracle Times Ten In-Memory Databases, organizations can meet the transaction speed requirements of a wide range of performance-critical projects (Figure 1). Flexible SGI Altix servers scale down to meet the needs of small projects, and offer unique high-end scalability for applications with large data set sizes. In fact, the SGI Altix 4700 server supports up to 128 TB of globally addressable memory. The ability to affordably process massive amounts of data in real time opens up a wealth of possibilities for enterprises to improve business results, reduce risk, and prevent fraud.

Small configuration

- SGI® Altix® 450 server
- 64 GB Oracle® TimesTen In-Memory Database
- Up to 8 CPUs and 128 GB memory
- SGI® InfiniteStorage 220
- Call center routing
- Customer relationship management
- Real-time billing

Mid-size configuration

- SGI Altix 450 server
- 128 GB Oracle® TimesTen In-Memory Database
- Up to 16 CPUs and 256 GB memory
- SGI InfiniteStorage 220
- Fraud detection
- Knowledge based systems
- Application server cache
- Mediation systems

Large configuration

- SGI® Altix® 4700 server
- 1536 GB Oracle® TimesTen In-memory Database
- Up to 32 CPUs and 3072 GB memory
- SGI® InfiniteStorage 4500
- Fraud detection
- Social networking
- RFID
- Biometrics

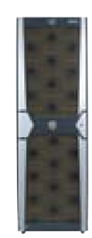

Figure 1. SGI and Oracle TimesTen In-Memory Database sample configurations

2.2 The Performance and Cost Benefits of a Memory-Centric Database Architecture

Until recently, systems that could host an in-memory database of only a hundred gigabytes broke most budgets. Due to lower memory prices, servers equipped to host large-scale in-memory databases are now affordable. In addition, employing commodity processors in combination with the exceptional speed of an in-memory database can provide a valuable price per performance advantage over a less responsive disk-centric architecture. Furthermore, large disk subsystems consume high amounts of energy and floor space. In contrast, utilizing a memory-centric architecture can positively impact the total cost of ownership by minimizing utility and real estate expenses. The results of an in-memory transaction benchmark provided by FedCentric Technologies illustrate the ability of memory-centric database architecture to produce exceptional response times while also providing dramatic energy and datacenter floor space efficiencies (Table 2).

	In-Memory Transaction Benchmark
Processor type	Itanium 2
Number of processor cores	32
Database engine	Oracle TimesTen In-Memory Database
Memory	1024 GB
Total storage	4,096 GB
Server cost	$946,659
Storage cost	$135,226
Average response time	0.00486 to 0.01738 seconds
Performance in transactions per minute (tpm)	32,183,908
Price/performance	$0.06
Kilo Watt hour (kWh) consumption per year	139,858
Annual power cost	$17,655
Server footprint (cubic ft)	51.4
Storage footprint (cubic ft)	1.4
Energy Efficiency (tpm/kW)	2,2017,223.2
Space Efficiency	
(tpm/cubic ft)	609,543.71

Table 2. Performance results and cost metrics for an in-memory transaction benchmark executed on an SGI Altix 4700 server

3.0 Oracle TimesTen In-Memory Database

The Oracle TimesTen In-Memory Database is a unique offering, enabling real-time performance on a standards-based, easy-to-use database engine that is capable of multiuser operation and ready for deployment in a number of scenarios. Utilized as a standalone relational database, as a cache for traditional disk-based relational databases, or for application-tier data management in next generation enterprise architectures, the Oracle TimesTen In-Memory Database helps organizations simplify the implementation of real-time processing projects.

3.1 Instantly Responsive, High Throughput Database Architecture

The Oracle TimesTen In-Memory Database is a memory-optimized relational database that empowers applications with instant responsiveness and very high throughput. Often deployed in the application tier as a cache or as a standalone database, the Oracle TimesTen In-Memory Database operates on data stores that fit entirely in physical memory using standard SQL interfaces. By optimizing data structures and algorithms for in-memory data storage, utilization of an Oracle TimesTen In-Memory Database can drive dramatic performance gains — even when compared to a fully cached relational database. In fact, an Oracle TimesTen In-Memory Database can retrieve a database record in less than 10 microseconds and update or insert a record in less than 30 microseconds (Figure 2). Furthermore, organizations can enhance application performance by embedding Oracle TimesTen In-Memory Database libraries within applications to eliminate context switching and unnecessary network operations.

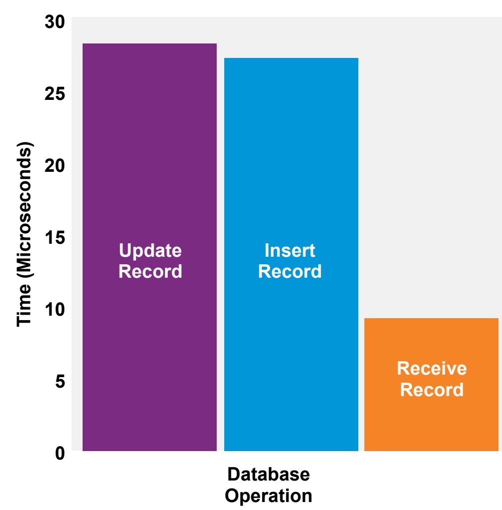

Figure 2. Oracle TimesTen In-Memory Database average response times for database operations on a four CPU system

3.2 Delivering Multiuser Capabilities, Reliability, and Data Integrity

In-memory databases are often believed to be limited to single-user applications and exposed to data loss when a server fails. Designed with data integrity in mind, the Oracle TimesTen In-Memory Database includes capabilities such as row-level locking with committed-read isolation, making it safe for multiuser, multithreaded applications. In addition, internal record versioning eliminates lock contention between readers and writers, providing consistent response times and supporting high levels of concurrency.

In order to maintain data integrity in the face of possible system faults, enterprises can take advantage of a transaction logging feature of the Oracle TimesTen In-Memory Database. Log records can be written to a disk-resident copy of the database asynchronously or synchronously throughout processing. Specific settings for log records are controlled by the application at the transaction level. In cases where data integrity must be continuously preserved, organizations can use memory image checkpointing along with synchronous logging to ensure durability. As an alternative, the asynchronous logging capabilities of Oracle TimesTen In-Memory Database enable extremely fast performance with minimal exposure, providing an option for even higher throughput.

3.3 Simplifying Application Development

Most real-time systems require custom development of hard-coded functions, limiting flexibility and adding solution complexity. Even commercial off-the-shelf products for high-performance transaction processing often utilize proprietary application programming interfaces (APIs) and custom data models to achieve performance goals. Built on open, flexible, industry-standard interfaces and technology, the Oracle TimesTen In-Memory Database increases the viability of broader adoption and the implementation of real-time transaction processing. In fact, database administrators and programmers already familiar with Oracle Databases or SQL can immediately start productive development of real-time applications with queries against Oracle TimesTen In-Memory Databases.
Structured Query Language (SQL92), Java Database Connectivity (JDBC), and Open Database Connectivity (ODBC) technologies are used to access Oracle TimesTen In-Memory Databases just as with standard relational databases. SQL provides a layer of abstraction between applications and Oracle TimesTen In-Memory Database internal code, enabling easy extension and alteration of application capabilities. Rapid integration of new services to an existing Oracle TimesTen In-Memory Database is accomplished by simply adding application modules, tables, and columns. Further simplifying development, a cost-based optimizer within the Oracle TimesTen In-Memory Database automatically determines the fastest way to process queries and transactions.

3.4 Easing Operations

Adding even more value, the structure of Oracle TimesTen In-Memory Databases minimizes the effort required for installation, setup, and administration. In fact, the simplicity of the disk structures and relatively small and consistent size of an Oracle TimesTen In-Memory Database obviates the need for most traditional database administration tasks. Oracle TimesTen In-Memory Database provides command-line utilities for backup, restore, database copy and migration, policy setting, interactive query, and monitoring functions. For installations which utilize the Oracle Data Replication and Cache Connect to Oracle options, utilities are also available for configuring and monitoring replication and caching functions.

3.5 Deployment Scenarios

Oracle TimesTen In-Memory Database supports a variety of configurations and scenarios, including transient look-up caches, operational data stores, and mission-critical transaction processing systems. While Oracle TimesTen In-Memory Databases are often embedded inside applications, traditional client-server access is supported. Many organizations find client-server access to an in-memory database advantageous for adjunct functions, such as reporting or speeding transactions when large numbers of application-tier platforms access a common database. Oracle also provides replication and caching options to enhance the applicability of the Oracle TimesTen In-Memory Database in load balancing and multi-tier application architectures.

3.5.1 Replication — TimesTen to TimesTen

Mission-critical application deployments often implement the Replication — TimesTen to TimesTen option. The Replication — TimesTen to TimesTen software enables the implementation of an active-standby or active-active configuration, providing the foundation for failover scenarios that improve database availability. In addition, the Replication — TimesTen to TimesTen software option supports multi-way and peer-to-peer replication. Replication can occur at the table or data-store level using asynchronous or synchronous transmissions. The Replication — TimesTen to TimesTen software also features conflict detection and resolution and automatic resynchronization after recovery of a failed instance. By utilizing the ability to sustain multiple synchronized data copies, architects can create solutions which distribute and load balance database transactions across many servers — further enhancing the performance of applications and services.

3.5.2 Cache Connect to Oracle

The Cache Connect to Oracle option enables tight integration between an Oracle Times-Ten In-Memory database and a standard Oracle Database 10g, Oracle Database 11g, Oracle Real Application Clusters (RAC) 10g, or Oracle RAC 11g instance. By utilizing Cache Connect to Oracle software, the Oracle TimesTen In-Memory Database can function as a real-time dynamic data cache for either type of disk-based Oracle database. Furthermore, the automatic data loading and update synchronization features of Cache Connect to Oracle vastly simplify the management of dataset refreshes between a standard Oracle Database and an Oracle TimesTen In-Memory Database.

Database architects can choose to cache a full set or subset of disk-based database tables in an Oracle TimesTen In-Memory Database. Administrators can specify the caches for read-only access, uni-directional updates, or bi-directional updates. Asynchronous and write-through updates are supported, and multiple cache instances can be created against a single Oracle database.

The Cache Connect to Oracle option is particularly useful when deploying the Oracle TimesTen In-Memory Database in an application tier caching role to improve Web service response times. For example, a read-only cache can store data on top sales items for an on-line store. Entries related to less frequently purchased products can be held on a disk-based data store. An updatable cache containing the most recently used data can improve performance for active sessions and users. For Web services with extremely large user volumes, instantiating multiple updatable caches can help improve throughput.
Using an Oracle Times-Ten In-Memory database with the Cache Connect to Oracle option can help enterprises capture data in real time for later transfer to disk storage. For example, an application that gathers sensor data in real time can initially store the data to an Oracle Times-Ten In-Memory database. The Cache Connect to Oracle option can transfer the captured data to a standard Oracle database for longer term storage at a later time.

4.0 SGI Altix Servers — Ideal Platforms for Oracle TimesTen In-Memory Database Processing

Robust, scalable, high-performance SGI Altix servers give enterprises an ideal platform for reaping the full benefits of an Oracle TimesTen In-Memory Database. SGI Altix servers combine the power of Intel Itanium 2 processors and industry-standard Linux environments with the SGI NUMAflex architecture and global shared memory to create a platform that is powerful, uniquely flexible, and highly reliable.

4.1 SGI Altix 4700 Servers and SGI Altix 450 Servers

Modular, high-performance SGI Altix 4700 and SGI Altix 450 servers help organizations maximize the performance and scalability of Oracle TimesTen In-Memory Database solutions (Figure 3). SGI Altix 4700 and SGI Altix 450 servers are comprised of blades — including interchangeable compute, memory, I/O, and special purpose blades — in compact packaging. SGI Altix servers provide excellent performance density and deliver as much as one teraflop per tall rack.
Based on the innovative SGI NUMAflex system architecture, these servers support the independent scaling of processors, memory capacity, memory bandwidth, interconnect bandwidth, I/O connectivity, and I/O bandwidth. As a result, enterprises can mix and match a variety of standardized blade choices, and take advantage of system configuration flexibility to meet current and changing requirements easily and cost-effectively.

Figure 3. The SGI Altix 4700 and SGI Altix 450 servers

SGI Altix 4700 and SGI Altix 450 servers offer exceptional configuration flexibility and scalability (Table 3). Two alternate compute blades, supporting dual-core Intel Itanium 2 Series 9000 processors are available. One blade is optimized for maximum performance with top memory bandwidth and the other blade is optimized for cost-effective compute density.

- The SGI Altix 4700 server supports up to 1,024 cores (512 sockets) with up to 12.3 TB of globally shared memory under one instance of the Linux environment and 128 TB of globally addressable memory per system.
- The SGI Altix 450 server supports up to 76 cores (38 sockets) with up to 912 GB of globally shared memory under one instance of the Linux environment.

Within the SGI Altix 4700 and SGI Altix 450 servers, the SGI NUMAlink4 interconnect provides a raw single link aggregate transfer rate of 6.4 GB/second. The NUMAlink 4 interconnect efficiently propagates coherent memory traffic across large system domains. As a result, NUMAlink interconnect technology delivers best-in-class low hardware latencies. In fact, SGI Altix servers achieve direct memory access to remotely located memory in a few hundreds of nanoseconds versus up to 30 microseconds for competing interconnects.

	SGI Altix 450 Server	SGI Altix 4700 Server
Processors	• Intel Itanium 2 dual-core processors • 1.4 GHz/12 MB cache • 1.6 GHz/8 MB, 18 MB, or 24 MB cache • 533 MHz Front Side Bus Support	
Number of cores (under one instance of Linux)	• 4 to 76 per system	• 6 to 1024 per system
Interconnection between nodes for global shared memory	• SGI NUMAlink 4 • Fat-tree topology • 6.4 GB/sec bidirectional bandwidth per link	
Memory	• Up to 48 GB DDR2 per blade • Up to 912 GB per system globally shared memory	• Up to 48 GB DDR2 per blade • Up to 12.3 TB per system globally shared memory • Up to 128 TB per system globally addressable memory
Memory-only modules available?	• Yes	
Base I/O blade	• Two low-profile PCI-X slots • One SAS port • Two Gigabit Ethernet ports • Four USB ports	
PCI blades available	• PCI-X and PCI-Express expansion blades	
System disk	• Up to two hard drives • Mix or match 300GB SAS or 500GB SATA2 hard drives	• Up to two hard drives • 300GB SAS hard drives
Disk expansion	• JBOD, RAID, Data Servers, Tape, and Libraries	
Rack	Short rack • 41.8" H x 25.8" W x 40.9" D • 20U internal usable space per rack • Up to 38 sockets per short rack • Tall rack also available	Tall Rack • 79.5"H x 25.8" W x 45.0" D • 42U internal usable space per rack • Up to 78 sockets per rack

Table 3. SGI Altix 4700 and SGI Altix 450 server configuration details

4.2 Memory Scalability and Speed — Key to In-Memory Database Performance

The memory architecture of a server greatly influences the peak scalability and throughput of an in-memory database solution. The unique global shared-memory scalability of SGI Altix 4700 and SGI Altix 450 servers provides the ability to efficiently manage and access large data sets and execute more instructions per cycle. Through the SGI NUMAflex memory architecture, SGI Altix servers can optimize the performance of Oracle TimesTen In-Memory Database solutions.

A system architecture that provides efficient direct processor load-store access to the entire address space is key to making effective use of global shared memory. The globally shared memory architecture of SGI Altix servers enables all system processors to directly access memory, regardless of memory location. The SGI NUMAflex architecture couples the large physical address space of the Intel Itanium processor with a system interconnect capable of distributing that address space seamlessly across hundreds or even thousands of nodes. As a result, processor issued load instructions can read any address in the entire global shared memory space of local and remote nodes. The SGI Altix 4700 and SGI Altix 450 servers also integrate SGI's Peer I/O technology, enabling high-speed access to large shared memory for all system components.

The SGI NUMAflex architecture demonstrates industry leading performance with memory bandwidth seven times faster than the nearest enterprise competitor. In fact, cost-effective SGI Altix servers utilize the high performance SGI NUMAflex architecture to manage larger data sets and move data faster than leading systems from IBM, HP, and Sun (Table 4).

System Category	Maximum Global Addressable Memory	Memory Bandwidth
Sun Fire E25K Server	1 TB	115 GB/second
HP 9000 Superdome Server	2 TB	533 GB/second
IBM System p595 Server	2 TB	811 GB/second
SGI Altix 4700 Server	128 TB	5478 GB/second

Table 4. Memory capacity and bandwidth comparisons

Furthermore, the ability of SGI Altix servers to configure processors and memory independently provides enterprises with economic advantages. SGI Altix server configurations support as much as 1 TB of memory with only 6 cores (3 processors), while competing systems require at least 128 cores (64 processors) to handle similar memory configurations (Table 5). By utilizing SGI

Altix servers, architects can build platforms to precisely match the need of a particular database and avoid purchasing unnecessary processors, lowering both system costs and software licensing fees.

System Category	Maximum Memory Configuration per Core	Number of Cores Required for 1 TB System
Sun Fire E25K Server	8 GB	128
HP 9000 Superdome Server	16 GB	64
IBM System p595 Server	32 GB	32
SGI Altix 4700 Server	208 GB	6

Table 5. Processor core to memory capacity ratio comparisons

4.3 Optimized for Uptime and Smooth Operations

With the execution of mission-critical database transactions at stake, organizations need system platform features that help minimize the need for planned downtime and mitigate the impact of unanticipated system faults. SGI Altix servers boast the reliability, availability, and serviceability (RAS) features to meet stringent operational demands for Oracle TimesTen In-Memory Database deployments. In addition, numerous and continued contributions by SGI help make the Linux environment ready for large memory systems. The advanced reliability, availability, and serviceability features of SGI Altix servers and features of the Linux operating system work together to help organizations maximize uptime and simplify the operation of an Oracle TimesTen In-Memory Database solution.

4.3.1 Reliability

The SGI Altix 4700 and SGI Altix 450 servers are comprised of reliable, modular blades built of meticulously selected, high quality components from top tier suppliers to ensure robust component reliability. SGI minimizes potential hardware faults and error conditions through careful design, and the selection of system components that automatically detect and correct errors. For example, Itanium 2 processors utilized by the SGI Altix 4700 and SGI Altix 450 servers include built-in capabilities which detect and correct most on-chip errors. In addition, SGI Altix systems with dual-core Itanium 2 processors leverage the Intel Cache Safe technology which disables L3 cache lines when an uncorrectable error occurs, enabling systems to continue operation in the face of errors that can otherwise cause disruption to operation.

Redundant hardware components further improve the reliability of SGI Altix platforms. Redundant, variable speed fans adjust automatically in order to keep servers running at optimal temperatures, and redundant power supplies help keep platforms running in the face of power faults.

SGI Altix systems are also designed to reliably detect all errors in system memory, directories, and data paths. SGI utilizes Error Correcting Code (ECC) on all system buses and memories to detect and correct single bit errors and detect double-bit errors, as well as cyclic redundancy check (CRC) on all NUMAlink 4 channels between SGI Altix blades. SGI Altix servers also contain an extensive environmental monitoring and control system to protect hardware operation.

Furthermore, SGI continuously makes error avoidance and management contributions to the Linux community. In fact, SGI has done more than almost any other server vendor to ensure the performance and reliability of the Linux environment for large-scale memory systems, including memory uncorrectable error recovery enhancements, hardware error reporting improvements, technology to reduce panics on double-bit errors, fault-containment for cross-partition jobs, and more.

4.3.2 Availability

The hardware partitioning capabilities of SGI Altix servers improve platform availability and ease maintenance, providing organizations with an ideal platform for hosting Oracle TimesTen In-Memory Databases.

- A single SGI Altix server can be subdivided into multiple, logical systems without the need to adjust physical cables.
- Memory protection is built into the SGI designed chipset within each CPU blade, ensuring memory faults are contained to individual partitions and unexpected writes from one partition to another are avoided.
- SGI Altix NUMAlink 4 router chips integrate a feature called reset fences which isolates partitions, ensuring that a reset in one partition does not impact other partitions in the platform.
- The block transfer engine (BTE) built into Altix memory hub chips provides a reliable means to transfer data between partitions without altering CPU memory protection, enabling partitions to share data via high speed copying as desired and ensuring disruption in a remote partition does not crash or hang a partition that is actively performing a block transfer.

SGI Altix hardware partitions also reboot independently without affecting operations in other partitions, providing a number of benefits that can improve overall availability of an Oracle TimesTen In-Memory Database. For example, hardware can 6
be removed or added in one partition without disrupting other partitions. Software updates can be applied to individual partitions without halting the entire SGI Altix platform. On SGI Altix servers, the impact of necessary reboots and ill-behaved software remain isolated within individual partitions, enabling software development, testing, and production environments to safely co-exist on a single platform.

4.3.3 Serviceability

The SGI Altix design includes important serviceability features that enhance system monitoring, enable online system management and maintenance, and simplify fault analysis.

- **Designed for serviceability**
  Blades are housed in a chassis referred to as an Individual Rack Unit (IRU) and are electronically isolated so that individual blades can be replaced without powering down the IRU. The advanced, modular design of SGI Altix blades enables easy access to individual system components for service, maintenance, or upgrades while system operations continue on remaining blades.

- **Continuous operation**
  In most cases, power supplies, fans, and individual PCI cards can be hot-swapped without interrupting the operation of the system or partition containing the component. Helping to minimize unplanned downtime, compute and memory blades can be disabled without affecting system operation, enabling repair to wait until a scheduled maintenance window. Furthermore, the advanced RAS capabilities of the Intel Itanium 2 processor such as Cache Safe, minimize the likelihood of a CPU failure. In addition, processors and memory are always subjected to self-test at boot time and are automatically de-allocated if failures occur. SGI Altix systems can boot without the decommissioned component, enabling operations to continue.

- **Simplified monitoring and maintenance**
  To ease monitoring and maintenance, SGI Altix 4700 and SGI Altix 450 IRU enclosures contain an embedded system controller that runs off standby power and is operational whenever the enclosure is connected to an active power source. The SGI Altix system controller network manages the hardware partitions within each system, providing pinpoint power control, system booting, and runtime system monitoring, along with support for configuration control, hot swapping, and diagnostics execution. The system controller can transparently extract all internal register states and actions from compute, memory, and router blades while the system is running, providing a wealth of input data that enables a fault analyzer to produce failure data reports down to the field replaceable unit (FRU) level. In addition, the system controller can read the complete hardware configuration down to the level of individual FRU serial numbers in real time, helping support rapid and accurate notification and transmittal of essential information for system service actions.

4.4 Standards-Based Platforms

The use of proprietary technology in computing platforms often creates vendor lock-in, increases costs, and limits the longevity, flexibility, and usefulness of a solution. Built on industry-standard components and running the open source Linux operating system, SGI Altix platforms deliver uncompromised performance and the flexibility to execute thousands of commercial off-the-shelf software applications that are available for the Novell SUSE Linux Enterprise Server and Red Hat Enterprise Linux Advanced Server operating systems. By choosing standards-based SGI Altix servers for Oracle TimesTen In-Memory Database solutions, organizations can protect technology investments, reduce costs, and simplify solution integration.

4.5 Total Cost of Ownership

IT managers are constantly seeking platforms which lower the total cost of ownership of solutions. With SGI Altix servers, organizations can build entry-level architectures that can scale both compute capacity and capital expenditures incrementally. A pay as you go architecture enables organizations to add individual processors, memory components, and I/O devices as needed. For example, unique memory-only blades enable architects to configure additional memory on SGI Altix 4700 servers and SGI Altix 450 servers without adding CPUs or I/O devices. By decoupling processor and memory configuration, SGI servers can help reduce overall hardware expense, minimize software licensing costs, reduce utility charges, and decrease server footprint and associated real estate fees. In fact, SGI Altix server configurations can cost 75 percent less, require 50 percent less datacenter floor space, and consume 25 percent less power than solutions from other vendors.

5.0 Summary

With increasing frequency, enterprises depend upon the execution of business transactions in real time. Built on industry-standards, SGI Altix platforms and the Oracle TimesTen In-Memory Database combine to form an easy to adopt solution for delivering the performance, scalability, and reliability required to meet the real-time demands of modern enterprise applications. SGI's long-standing leadership in designing systems to process and store extremely large data sets contributes to the unique ability of SGI Altix platforms to maximize the performance of an Oracle TimesTen In-Memory Database. The massive memory capacity, high-speed performance, and robust design of SGI Altix platforms enable the effective processing of large data sets in real time, helping organizations accelerate breakthrough business improvements, innovation, and information transformation.

Furthermore, by utilizing a standards-based design approach and leveraging advanced technology from innovative industry leaders, SGI Altix servers deliver fast, scalable performance at an affordable price point. SGI collaborates with Intel on system design requirements for large-scale computing for Intel Itanium and Intel Xeon processors, partners with Novell on Linux support for scalability, performance, and certification of SGI servers and storage systems. In addition, SGI works closely with Red Hat on Linux support with a special emphasis on security and adherence to standards, and contributes thousands of lines of code to the Linux community, including code that supports large-scale computing, reliability, and system stability.

5.1 For More Information

For additional information on SGI Altix servers and the Oracle TimesTen In-Memory Database, please contact your local SGI representative or consult the Web references listed in Table 6.

Web URL	Description
http://sgi.com/products/servers/altix/	SGI Altix Family
http://sgi.com/products/software/linux/	Linux Operating System
http://sgi.com/industries/imdb/	In-Memory Database Applications
http://sgi.com/support/	SGI Support Services
http://oracle.com/database/timesten.html	Oracle TimesTen In-Memory Database

Table 6. Related Web sites

Corporate Office
SGI
1140 East Arques Avenue
Sunnyvale, CA 94085-4602
650.960.1980

North America +1 800.800.7441
Latin America +55 11.5185.2860
Europe +44 118.912.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000

4038 [9.2007]

J15319