

Sun Fire™ Enterprise Servers and the UltraSPARC® IV Processor

Scaling Application Performance with Throughput Computing

February 2004



Table of Contents

Delivering on Sun’s Throughput Computing Vision1
The Diminishing Returns of Complex Processor Design	2
Chip Multi-Threading (CMT): A Better Approach for Network Computing	4
Sun’s Balanced Architecture for Throughput Computing	5
Sun Fire Enterprise Servers7
A Scalable Range of High-Performance Servers	7
Unprecedented Investment Protection	9
Key Technology for Scalability and Performance	9
Real-World Performance Results13
SPEC CPU2000 Int_Rate (Sun Fire E6900 server)	13
SAP R/3 Enterprise 4.70 2-Tier SD Benchmark (Sun Fire E20K and E4900 Servers)	15
Manugistics Fulfillment v7.1 “Plan” Benchmark (Sun Fire 6800, E4900, and E6900 Servers)	16
SPEC JBB2000 (Sun Fire E6900 server)	18
SPEC OMPL2001 (Sun Fire E25K Server)	19
Conclusion21
References23

Chapter 1

Delivering on Sun's Throughput Computing Vision

The need for computational throughput in the modern data center has never been greater. Truly immense numbers of people and devices now depend on network computing for their day-to-day interactions with everything from wireless phones, to PDAs, to packages of razor blades increasingly assuming a network presence. This enormous increase in connectivity brings new opportunities and new demands for applications and services that can redefine the ways that companies do business, both internally and with their customers.

Unfortunately, IT budgets remain strictly limited and organizations are trying to do more with less — a goal often frustrated by complex and sprawling data-center environments. In fact, Sun estimates that as much as 70 percent of IT budgets are currently spent just maintaining data-center operations, with only 30 percent left over to add new services. Simply deploying more servers is clearly not an option for most, with additional servers often contributing to greater complexity and higher administrative costs while requiring more time to add new network services.

In the face of these challenges, Sun understands that increased performance and utilization can help reduce the cost of operations. Getting more work done with fewer servers lowers the cost and complexity of managing the IT environment, ultimately making more funding available for desperately-needed services that actually improve the business. Sun's *throughput computing* vision is dedicated to doing just that, providing greater real-world performance and throughput that lets organizations be more productive with fewer servers while taking costs out of operations.

More than just a vision, Sun is now delivering on throughput computing with the new Sun Fire™ E2900, E4900, E6900, E20K and E25K servers running the proven and robust Solaris™ Operating System. Along with the vertical scalability needed for many classes of applications, these systems deliver up to twice the computational throughput and compute density of Sun's previous-generation data-center servers in the same physical footprint. At the same time, these servers offer full binary compatibility and unprecedented levels of investment protection for owners of Sun's previous-generation systems.

This document describes Sun's throughput computing vision, the chip multi-threaded UltraSPARC® IV processor, and the new Sun Fire Enterprise servers along with the real-world performance and scalability improvements they can provide.

The Diminishing Returns of Complex Processor Design

While optimistic marketing statements constantly call attention to presumably impressive multiple-gigahertz frequencies for new generations of processors, corresponding small gains in system performance and productivity continue to frustrate IT professionals. Throughput computing is designed to help resolve these divergent trends by providing higher amounts of *delivered* performance and computational throughput while simplifying the data center. Understanding the importance of throughput computing requires a look at how both processors and systems have been designed.

The oft-quoted Moore's Law states that the number of transistors that will fit in a square inch integrated circuitry will double approximately every two years. For over three decades the pace of Moore's law has held, driving processor performance to new heights. Processor manufacturers have exploited the chip real estate gained with Moore's law to build increasingly complex processors with instruction-level parallelism as a goal. Today these processors employ very high frequencies along with a variety of sophisticated tactics to accelerate a single instruction pipeline, including:

- Large caches
- Superscalar designs
- Out-of-order execution
- Very high clock rates
- Deep pipelines
- Speculative pre-fetches

While this approach has produced faster processors with impressive-sounding multiple-gigahertz frequencies, real-world application performance improvement has not kept pace. (Table 1-1).

Processor	Celeron	Pentium 4	% Difference
Frequency	1.2 GHz	2.2 GHz	83%
Performance (Business Winstone 2001 ¹)	44.7	54.8	20%
Power	29.9W	55.1W	84%
Cost ²	\$103	\$562	446%

1. Source: PC Magazine (February 2002)

2. Cost source: CNET (January 4, 2002), Yahoo Finance News (January 8, 2002)

This example illustrates the familiar 80/20 rule, but in an undesirable direction. In this case, an 80-percent increase in frequency and power consumption yielded only a 20-percent increase in


Table 1-1: Business Winstone 2001 performance gains lag processor frequency gains

real-world system performance. And regrettably, the small 20-percent performance improvement came with more than a 400-percent increase in cost over the slower processor. Intelligent system design is principally a matter of removing bottlenecks, and this data clearly shows that processor frequency was not the most critical performance bottleneck for this benchmark.

These diminishing returns in real-world application performance are principally caused by much slower gains in memory speeds over time as compared to those for processors. While processor speeds continue to double every two years, memory speeds have typically doubled only every six years. This relative gap between processor and memory speeds leaves ultra-fast processors spending as much as 75 percent of the time waiting for memory to return required data, rather than performing useful work. Ironically, as processors get faster and more complex, the effect of memory latency grows — fast, expensive processors spend more cycles doing nothing. Its easy to see that frequency (megahertz) is truly a misleading indicator of real performance.

Figure 1-1 illustrates how even doubling processor performance can provide only a small relative increase in application performance. In this example, though the compute time is reduced by half, only a small overall improvement in execution time results due to the constant presence of memory latency.

Figure 1-1: Increasing single-threaded processor performance by 100 percent (a 50-percent reduction in compute-time) provides only a small relative gain in application performance due to memory latency



Further complicating performance expectations, many of today's data-center workloads are unable to take advantage of the instruction-level parallelism designed into modern processors. Applications with high shared memory and data requirements are typically more focused on processing a large number of simultaneous threads rather than running a single thread as quickly as possible (Table 1-2). In practice, many of these applications respond better to thread-level parallelism than to the instruction-level parallelism provided by complex modern processors.

Table 1-2: Attributes of common commercial workloads favor thread-level parallelism over instruction-level parallelism

Attribute	Web Services			Client Server		Data Warehouse
	Web (Web99)	Application server (JBB)	Data (TPC-C)	SAP 2T	SAP 3T (DB)	DSS (TPC-H)
Application category	Web server	Server, Java™	OLTP	ERP	ERP	DSS
Instruction-level parallelism	Low	Low	Low	Medium	Low	High
Thread-level parallelism	High	High	High	High	High	High
Instruction/Data working set	Large	Large	Large	Medium	Large	Large
Data sharing	Low	Medium	High	Medium	High	Medium

Chip Multi-Threading (CMT): A Better Approach for Network Computing

Unsatisfied with these diminishing returns, Sun's in-house design team of 1400 engineers — the second largest microprocessor design engineering team in the world — is taking a revolutionary approach to processor design. Sun understands that the network-computing environments found in most data centers are inherently multi-threaded, where the execution speed of an individual thread is typically less important than overall throughput. While continuing to design UltraSPARC processors that are optimized to process a single thread, Sun is also focusing on processors that are designed to maximize throughput for network-computing workloads. These chip multi-threading (CMT) processors take advantage of the additional gains delivered by Moore's Law to provide thread-level parallelism rather than instruction-level parallelism. Starting with the ability to run two concurrent threads in the UltraSPARC IV processor, support for multi-threading at the chip level will eventually enable single processors to process tens of threads simultaneously.

Figure 1-2 illustrates a hypothetical CMT processor that is designed to switch between up to four threads as each thread blocks for memory. With this approach, the processor remains active doing real useful work, even as memory operations for other threads continue in parallel.

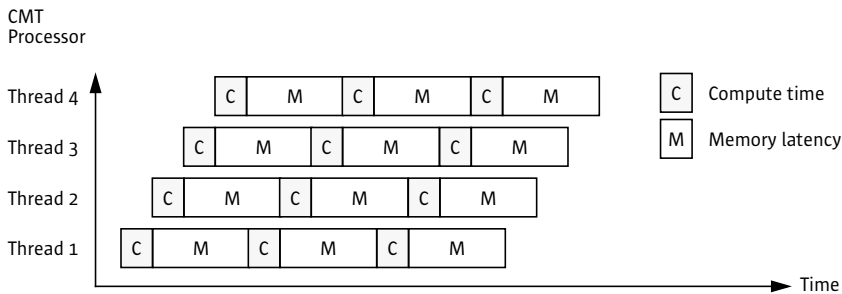
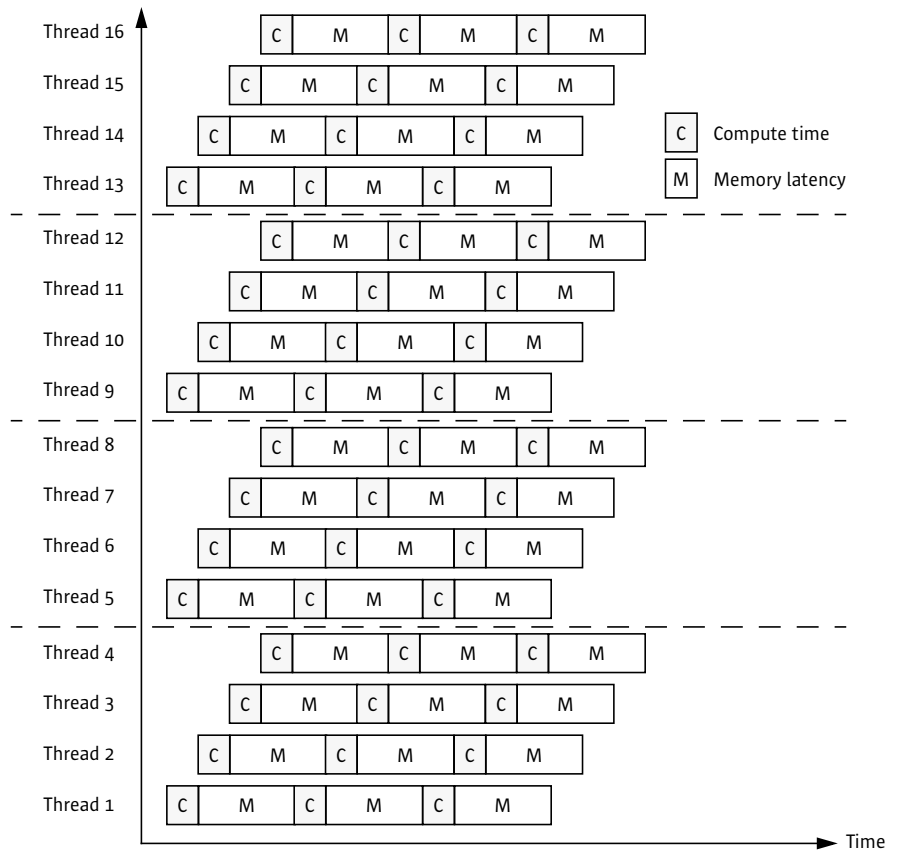


Figure 1-2: A hypothetical chip multi-threaded processor could process a number of threads concurrently by doing useful work while some threads perform memory operations

Application throughput is greatly improved as is utilization of processor resources. Thread-rich applications and environments benefit greatly from this model, whether through larger multi-threaded applications, or large numbers of smaller single-threaded applications. The number of simultaneous threads that can be accommodated is actually quite large. Figure 1-3 illustrates another hypothetical example where several CMT processors are combined into a single processor, supporting up to sixteen simultaneous threads.

Figure 1-3: More advanced chip multithreaded processors can process even larger numbers of concurrent threads



Sun's Balanced Architecture for Throughput Computing

Even with the promising performance gains of chip multi-threading, Sun understands that delivering real application performance and throughput requires a balanced approach. System-level architecture must allow for the considerable scalability that CMT can provide, without introducing other bottlenecks that can constrain performance. The operating system too is critical, and it must fully utilize processor resources while providing the ability to scale along with the most demanding thread-rich environments, large databases, and datasets.

For Sun, these design principals are simply second nature. As an early proponent of threaded applications and environments, Sun's architectures and operating systems have long assumed a multi-threaded workload. As a result, Sun is ideally positioned to deliver CMT to its greatest advantage with a minimum of disruption to ongoing operations. The new Sun Fire Enterprise E2900, E4900, E6900, E20K, and E25K servers continue this long tradition of designing for vertical scalability, while providing the necessary I/O bandwidth and headroom to support the performance gains of CMT processors.

- Each Sun Fire Enterprise server supports multiple chip multi-threaded UltraSPARC IV processors with each processor capable of running up to two concurrent threads, providing up to twice the throughput of previous-generation UltraSPARC® III processors. Up to 72 UltraSPARC IV processors are supported in the high-end Sun Fire E25K server for support of up to 144 concurrent threads.

- 64-bit memory addressing is key to providing high-end scalability and wider data paths. 64-bit UltraSPARC processors along with 64-bit support in the Solaris Operating System allow massive data sets and applications to be moved directly into main memory. UltraSPARC processors and the Solaris Operating System include support for 64-bit computing today while preserving 32-bit compatibility for existing applications.
- All Sun Fire Enterprise servers provide very large memory support (up to 576 GB in the Sun Fire E25K server) and they all feature the high-throughput Sun Fireplane crossbar system interconnect. These systems provides the memory capacity and considerable I/O bandwidth needed to help enable large multi-threaded environments.
- The Solaris Operating System has long provided symmetrical multiprocessing (SMP) and multi-threading support. In addition, the Java™ language automatically provides a thread-based programming model so Java technology applications (including the Sun™ Java™ Enterprise System) are automatically accelerated by UltraSPARC IV performance gains as well as those of future CMT designs.

Chapter 2

Sun Fire Enterprise Servers

By delivering up to twice the computational throughput and compute density of previous-generation systems, Sun Fire Enterprise servers demonstrate that throughput computing is not only real, but that its adoption is virtually transparent. After all, even the highest levels of performance are difficult to justify if the technology involved requires large re-engineering efforts and disruptions to critical on-going operations. This transparent performance improvement across a full line of mid-range and high-end servers is the result of Sun's careful design and application of key technologies that preserve customers investments while delivering higher levels of realized application performance.

A Scalable Range of High-Performance Servers

With their high levels of performance and density, Sun Fire Enterprise servers (Figure 2-1) offer a simplification for the data center. Fewer processors are needed and fewer servers are required to provide a given level of service — yielding lower administrative, resource, and facilities costs. Real estate costs are minimized since these servers occupy exactly the same physical footprint as their UltraSPARC III based predecessors. With proven technologies such as fifth-generation Dynamic System Domains and Dynamic reconfiguration, these systems can also achieve very high levels of utilization through flexible server consolidation.



Figure 2-1: Sun Fire Enterprise servers, including the Sun Fire E2900, E4900, E6900, E20K, and E25K servers (The Sun Fire E20K and E25K servers share the same cabinet design)

In addition to chip multi-threaded UltraSPARC IV processors, large memory capacities and large sustained levels of system bandwidth are provided across the Sun Fire Enterprise server product line to support demanding applications, large databases, and growing data sets. A scalable range of servers makes this innovative technology available for a wide range of applications and requirements.

- ***The Sun Fire™ E2900 server***

The Sun Fire™ E2900 server supports up to 12 UltraSPARC IV processors (up to 24 concurrent threads) along with up to 96 GB of RAM in a rugged and compact rack-optimized package. Support for up to 6 PCI cards and two internal disks is provided and the system offers a sustained bandwidth of 9.6 GB per second. The Sun Fire E2900 server supports a single instance of the Solaris Operating System.

- ***The Sun Fire™ E4900 server***

The Sun Fire™ E4900 server supports up to 12 UltraSPARC IV processors (up to 24 concurrent threads) along with up to 96 GB of RAM and up to 16 PCI slots in a desk-side or rack-mount system (one per rack). The system's Sun Fireplane interconnect provides 9.6 GB per second of sustained bandwidth. Up to two Dynamic System Domains are supported providing up to two fault-isolated instances of the Solaris Operating System.

- ***The Sun Fire™ E6900 server***

The Sun Fire™ E6900 server supports up to 24 UltraSPARC IV processors (48 concurrent threads) along with up to 192 GB of RAM and up to 32 PCI slots in a data-center rack configuration. The Sun Fireplane interconnect supplies 9.6 GB per second of sustained system bandwidth. Up to four fault-isolated Dynamic System Domains can be configured, providing support for up to four fault-isolated instances of the Solaris Operating System.

- ***The Sun Fire™ E20K server***

Packaged in a data-center rack, the Sun Fire™ E20K server supports up to 36 chip-multithreaded UltraSPARC IV processors (72 concurrent threads) and up to 288 GB of RAM. Up to 36 hot-swap PCI+ slots are supported with the new hsPCI+ I/O assembly that also delivers a 60-percent

performance improvement. Each server provides 29 GB per second of sustained bandwidth (86.4 GB/second peak) and up to nine fault-isolated Solaris Operating system domains are supported.

- ***The Sun Fire™ E25K server***

The flagship Sun Fire™ E25K server utilizes the same rack as the Sun Fire E20K server and supports up to 72 chip multi-threaded UltraSPARC IV processors (up to 144 concurrent threads), up to 576 GB RAM and up to 72 hot-swap PCI+ slots using the new high-performance hsPCI+ I/O assembly. Each Sun Fire E25K server provides 57 GB per second of sustained system bandwidth (172.8 GB per second peak). Up to 18 fault-isolated Solaris Operating System domains are supported in a single Sun Fire E25K server chassis.

Unprecedented Investment Protection

Organizations make considerable investments in their IT infrastructure and Sun strives to protect those investments for as long as possible. Sun Fire Enterprise servers accomplish this goal through a variety of means:

- ***100-percent binary compatibility***

Consistent with Sun's binary compatibility guarantee, these servers offer 100-percent binary compatibility with earlier generations of Sun systems. Applications simply run without modification or recompilation and customers can retain their existing Solaris™ 8 (2/04) or Solaris™ 9 (4/04 or later) Operating Systems.

- ***Mixed-generation processor support***

Sun Fire Enterprise servers can run both UltraSPARC III and UltraSPARC IV processors (of multiple speeds) in the same or separate domains. With multi-speed processor support, individual processors continue to operate at their rated speeds. Even more impressive, previous-generation Sun Fire™ V1280-15K servers can also run both UltraSPARC III and UltraSPARC IV processors, enjoying the same mixed-speed processor support.

- ***Uniboard resource provisioning and upgrades***

With common components throughout the Sun Fire Enterprise server product line, customers can choose either UltraSPARC IV or UltraSPARC III Uniboards (and CPU/memory boards for the Sun Fire E2900 server) without being forced into expensive and disruptive box-swap upgrades. Customers retain their investments in existing Sun Fire V1280-15K server platforms as well as investments in existing UltraSPARC III Uniboards and CPU/memory boards.

- ***Preserved application infrastructure investments***

With the robust Solaris Operating System, the Sun Fire Enterprise server product line preserves investments in applications, software development, training and infrastructure.

Key Technology for Scalability and Performance

Scalability is critical for organizations that need to make the most cost-effective use of their computing resources. Systems must gracefully handle peak workloads and allow growth of the computing environment in response to business conditions. The scalable Sun Fire Enterprise server product line provides a range of binary-compatible systems and performance points to accommodate diverse applications from entry-level and mid-range to high-end data-center servers. To meet changing demands, and to adapt to dynamic workloads, all of the Sun Fire Enterprise servers utilize key technology to provide additional scalability.

Chip Multi-threading (CMT) UltraSPARC IV Processors

Rather than merely ramping frequencies, the fourth-generation UltraSPARC IV processor (Figure 2-2) provides Sun's first chip multi-threading implementation supporting up to two independent and concurrent threads per processor. Even more impressive, these processors can deliver a near-linear scaling over the single-threaded UltraSPARC III processor in real-world applications (chapter 3).



Figure 2-2: The chip multi-threaded UltraSPARC IV processor is both binary and pin-compatible with the previous-generation UltraSPARC III processor

UltraSPARC IV processors are built using 130 nm technology and utilize an ultra-dense manufacturing process that approximately doubles the number of transistors on the die. Each UltraSPARC IV processor is comprised of approximately 66 million transistors compared to 29 million for each UltraSPARC III processor. Even though the UltraSPARC IV processor doubles the computational capabilities, it is both binary-compatible, and pin-compatible with the earlier UltraSPARC III processor.

The Sun™ Fireplane System Interconnect

With the Sun™ Fireplane Interconnect, Sun continues the tradition of providing superior memory and I/O bandwidth along with predictable performance and scalability under demanding loads. Features of the Sun Fireplane interconnect include:

- A fast operating frequency for greatly increased performance over previous designs
- High-performance crossbar switch design allowing simultaneous transfers between processors, memory, and the I/O subsystem
- Low latency memory access
- Completely separate address/control and data paths for flexible implementation
- Out-of-order transaction processing enabling multiple simultaneous transactions on the bus
- High-throughput paths to memory via 576-bit wide paths (including ECC)
- Economical implementation through distributed control (no central memory controller)
- Integrated support for multiprocessor configurations

The Sun Fireplane interconnect can be built with various levels of interconnect switching to implement different kinds of systems. Sun Fire E2900, E4900, and E6900 servers all utilize a three-level interconnect with the first two levels of switching residing on the Uniboard (Sun Fire E4900 and E6900 servers) or CPU/memory board (Sun Fire E2900 server). With all of these systems, the third level of the Sun Fireplane interconnect is implemented by a 10-way Address Repeater and a 10 x 10 Data Switch. For redundancy and serviceability, the third level of switch logic is mounted

on removable and hot-swappable Sun Fireplane Switch boards that attach to a passive system centerplane.

A four-level interconnect is used on the high-end Sun Fire E20K and Sun Fire E25K servers that support up to 36 and 72 processors respectively. In these servers, the third level of the interconnect is located on an Expander Board that also serves as a mounting bracket to hold the Uniboard and I/O assemblies. The fourth level of switching logic is located on the redundant active centerplane of the system consisting of three 18 x 18 crossbars — one each for addresses, responses, and data packets. It is important to remember that the Uniboard used in these systems is exactly the same component that is used in the mid-range Sun Fire E4900 and E6900 servers.

Fast PCI I/O Assemblies

All of the Sun Fire Enterprise servers provide fast I/O capabilities with PCI I/O assemblies that connect directly to the Sun Fireplane interconnect. To support high-performance I/O in significant volume, Sun Fire E20K and E25K servers utilize an innovative hot-swap PCI+ (hsPCI+) I/O assembly that delivers up to a 60-percent performance improvement over the previous-generation I/O assembly. The hsPCI+ I/O assembly also provides for more 66 MHz PCI slots than the previous-generation I/O assembly with up to 54 66MHz PCI slots supported in a single Sun Fire E25K server. Each hsPCI+ I/O assembly holds up to four PCI cards in hot-swap cassettes. This approach enables hot-swap and Dynamic Reconfiguration capabilities for individual standard PCI cards.

The Sun Fire™ Link Cluster Interconnect

Many organizations today require the ability to scale applications beyond the computational resources of a single server. Coupled with the Solaris Operating System and appropriate middleware, the Sun Fire™ Link Interconnect can bring the power of multiple Sun Fire Enterprise servers to bear on everything from large-scale commercial computational activities to today's most important "grand challenge" scientific enquiries. The Sun Fire Link interconnect provides high-throughput, low-latency clustering technology while middleware software exploits remote shared memory (RSM) techniques, bypassing the TCP/IP stack for high performance and reduced latency. Appropriate middleware includes:

- Oracle 10G and 9iRAC for large databases
- Sun™ Cluster 3.1 software for high-availability solutions
- Sun HPC Clustertools™ software for high performance and technical computing (HPTC) environments

The Sun Fire Link assembly provides an interface between the Sun Fireplane system interconnect and two fiber-optic connections to the Sun Fire Link network. Sun Fire E6900, E20K, and E25K servers can all participate in Sun Fire Link clusters as well as previous-generation Sun Fire 6800, 12K and 15K servers. Unlike PCI adapters, the Sun Fire Link assembly taps directly into the Sun Fireplane interconnect and runs at system interconnect speeds. Each Sun Fire Link assembly replaces a standard I/O assembly for the Sun Fire™ 6800, E6900, 12K/15K, and E20K/E25K servers. To match system form factors, there is one Sun Fire Link assembly for the Sun Fire 6800 and E6900 server, and a different assembly for the Sun Fire™ 12K, 15K, E20K, and E25K servers although the assemblies share logic.

Sun Fire Link configurations can either be directly connected (2-3 nodes) or switched (2-8 nodes). For switched configurations, two Sun Fire Link switches route signals from each clustered node to any of up to seven other clustered nodes with minimal latency (3-4 microseconds). Up to

eight Sun Fire Link optical modules can be installed in each rack-mountable Sun Fire Link switch, supporting clustered configurations of up to eight nodes with redundant switch connections.

Solaris™ Operating System

Nowhere is Sun's commitment to scalability more apparent than in the Solaris Operating System's support for massive systems and applications, including:

- Up to 128 processors in a single system and 848 processors in a clustered environment
- Configurations with up to 576 GB of main memory
- A 64-bit process address space (on 64-bit processors), supporting extremely large applications with enormous amounts of data
- One million processes
- One terabyte native file systems, and up to 252 TB file systems using Sun StorEdge™ QFS software
- Four billion open sockets, providing practically limitless scalability for network applications
- IPv6 support, enabling a 128-bit IP address space and lifting the limits faced by large organizations and global service providers
- 64-bit Java™ virtual machine, enabling very large Java technology applications and seamless integration with IPv6 environments

This scalability lets applications and their data grow to almost arbitrary size. The Solaris Operating System also provides features that help applications run faster by optimizing their memory utilization.

- ***Multiple page size support***

Many applications such as database management systems use a small amount of code to process a large amount of data, resulting in a small locality of reference for instructions and a large locality of reference for data. In order to improve processor cache behavior for these situations, the Solaris Operating System enables applications to have multiple page sizes for different memory segments, ranging from eight kilobytes to four megabytes. Using multiple page sizes can increase performance by reducing Translation-Lookaside Buffer (TLB) misses. Sun's implementation allows multiple page sizes to be specified while leaving the binary untouched. As a result, this approach allows page size to change dynamically as the application executes.

- ***Memory placement optimization (MPO)***

As both processor and system technologies evolve, an increasing number of systems have multiple layers of memory. For example, as chip fabrication techniques improve, more cache memory can be co-located with the processor, helping to reduce instruction and data fetch times. The Solaris Operating System is designed to be memory hierarchy-aware, enabling it to optimize memory management to best match applications' memory-access patterns. Rather than forcing recompilation for different memory hierarchies, Sun's philosophy of binary compatibility across the product line dictates a mechanism whereby the application can give hints as to its data access patterns. The Solaris Operating System then manages the memory hierarchy to provide the optimum locality of data for the application.

Chapter 3

Real-World Performance Results

Previous-generation Sun Fire servers have demonstrated near-linear scalability when adding additional processors, a testament to both balanced system architecture as well as the robust symmetric-multiprocessing capabilities of the Solaris Operating System. With Sun Fire Enterprise servers, these benefits now apply to the additional computational resources provided by chip multi-threaded UltraSPARC IV processors. The benchmark results provided in this section demonstrate the industry-leading result and world-record performance of the Sun Fire Enterprise server product line and the ability of these systems to meet the demands of very challenging and diverse workloads.

Ultimately, system benchmarks are only one indicator of performance that be used to evaluate a system's suitability for a given set of applications. To help design and size appropriate configurations, Sun also offers an extensive program of reference architectures that provide tried and tested configurations for specific application scenarios. Customers can even evaluate their specific applications directly through Sun's iForce™ Solution Centers. By using iForce Solution centers, customers can build proof-of-concept configurations using actual Sun equipment, run benchmarks using their own applications and data, and evaluate systems based on actual, rather than theoretical performance — all at no cost.

SPEC CPU2000 Int_Rate (Sun Fire E6900 server)

The SPEC CPU2000 benchmark is the world's leading measure of processor performance. SPEC CPU2000 is available in two distinct applications suites comprised of codes that test integer and floating point performance respectively.

The integer component of the SPEC CPU2000 benchmark consists of twelve benchmarks written in both the C and C++ languages, including:

- 164.gzip — File compression
- 175.vpr — FPGA circuit placement and routing
- 176.gcc — C program language compiler
- 181.mcf — Combinatorial optimization
- 186.crafty — Game playing: chess
- 197.parser — word processing
- 252.eon — Computer visualization
- 253.perlbmk — PERL programming language
- 254.gap — Group theory, interpreter
- 255.vortex — Object-oriented database
- 256.bzip2 — Compression
- 300.twolf — Place and route simulator

The two suites (integer and floating point) can be run as either *speed* or *rate*. Speed refers to running a single copy of the benchmark to see how fast individual components respond. The rate metric runs multiple copies of the kernels concurrently, giving a measure of how the system works under load. Each suite is further broken down into peak and base metrics. The base metric assumes that the same compiler flags are used to compile the entire suite, while the peak metric allows for individual kernel tuning by the compiler. Rate, speed, peak, and base metrics should not be compared against each other, as they test things in different ways and fair comparison by SPEC rules requires that only like metrics be compared.

SPEC CPU2000 Int_Rate Benchmark Results

The Sun Fire E6900 server delivers outstanding throughput performance as measured by the SPEC CPU2000 rate benchmark suite. The mid-range Sun Fire E6900 server equipped with 24 1.2 GHz UltraSPARC IV processors performs comparably to other vendors' high-end systems on this benchmark. These results were accomplished with eight fewer, lower-frequency processors than the competition and demonstrate the considerable cost savings available with the Sun Fire E6900 server and CMT technology (Table 3-1).

Processor	Number and Type of Processors	Jobs	Base	Peak
HP Integrity Superdome	32 x 1.5 GHz Itanium2	32	453	453
IBM p690	32 x 1.7 GHz Power4+	32	322	339
Sun Fire E6900 server	24 x 1.2 GHz UltraSPARC IV	48	295	337
Fujitsu HPC1500	32 x 1.35 GHz SPARCV	32	246	288
Sun Fire 6800 server	24 x 1.2 GHz UltraSPARC III	24	158	180

Table 3-1: SPEC CPU2000 Benchmark results for the mid-range Sun Fire E6900 server and several competitors high-end systems

The SPECint_rate2000 peak result of 337 for the Sun Fire E6900 server equipped with 24 UltraSPARC IV processors demonstrates:

- System performance scaling of 1.87 times as compared with the Sun Fire 6800 server with the same number and speed of UltraSPARC III processors
- 33 percent better per-processor performance and comparable overall performance to the 1.7 GHz IBM p690 (within 1 percent)
- 56 percent better per-processor performance than the 1.35 GHz Fujitsu HPC1500
- Better system performance than the 1.35 GHz 32-way Fujitsu HPC1500
- Comparable per-processor performance to the 1.5 GHz HP Integrity Superdome (within 1 percent)

Disclosure

SPEC and the benchmark name SPECcpu2000 are registered trademarks of the Standard Performance Evaluation Corporation. Competitive benchmark results stated above reflect results published on www.spec.org as of February 11, 2004. The comparison presented above is based on the best peak performance for all systems. Sun Fire E6900 server results were submitted to SPEC on 2/11/04 and are currently under review as of this writing.

SAP R/3 Enterprise 4.70 2-Tier SD Benchmark (Sun Fire E20K and E4900 Servers)

The SAP Standard Application Benchmarks test and prove the scalability of mySAP Business Suite. The benchmark results provide basic sizing recommendations by testing new hardware, system configurations, system software components, and Relational Database Management Systems (RDBMS). The SAP 2-Tier SD benchmark represents an ERP/OLTP system.

The benchmarking procedure is standardized, well-defined, and monitored by the SAP Benchmark Council that is made up of representatives of SAP and technology partners involved in benchmarking. Originally introduced to strengthen quality assurance, the SAP Standard Application Benchmarks can also be used to test and verify scalability, concurrency and multi-user behavior of system software components, RDBMS, and business applications. All performance data relevant to system, user, and business applications are monitored during a benchmark run and can be used to compare platforms and as basic input for sizing recommendations.

The Sun Fire Enterprise servers again demonstrated superior performance setting both world-record performance (Sun Fire E20K server) as well as excellent SAPs/processor performance (Sun Fire E4900 server). In these benchmarks, both the database server and the application servers were all run on the same Sun Fire Enterprise server. This approach demonstrates the effectiveness of running multiple applications in a single instance of the Solaris Operating System and reinforces Sun's position as a leader in server consolidation.

Sun Fire E20K Server, World-Record Performance

The SAP-SD 2-Tier benchmark was run on a 36-way Sun Fire E20K server with 288 GB of memory using 11 Sun StorEdge T3 Arrays for mass storage. Software used included the Solaris 9 Operating System (Update 6) and Oracle 9i (9.2.0) with SAP R/3 Enterprise 4.70. The Sun Fire E20K server recorded 5,050 benchmark users with a response time of 1.72 seconds, processing 517,330 order line items per hour.

The Sun Fire E20K achieved a world-record on the SAP R/3 Enterprise 4.70 SD 2-tier benchmark beating the previous best number of benchmark users by 25 percent. These results also demonstrate superior performance per processor versus a 32-way, 1.5 GHz Itanium 2-based NEC 5800, beating NEC by 11 percent on a per-processor basis.

Sun Fire E4900 Server, Excellent SAPs/Processor Performance

The SAP-SD 2-Tier benchmark was also run on a two-way Sun Fire E4900 server with 32 GB of memory using 4 Sun StorEdge T3 arrays for mass storage. Software used included the Solaris 9 Operating System (Update 6) and Oracle 9i database (9.2.0) with SAP R/3 Enterprise 4.70. The Sun Fire E4900 server supported 320 benchmark users with a response time of 1.83 seconds processing 32,330 order line items per hour.

For small processor counts, the Sun Fire E4900 server demonstrated excellent per-processor performance, achieving 810 SAPs per processor. This result is better than over 80 percent of all per-processor results, better than most of the Xeon based systems running at 2.4 to 3.0 GHz, and on par with most of the 1.5 GHz Itanium 2-based systems.

Disclosure

Directly comparable results using SAP Enterprise 4.70 from the official SAP benchmark Website as of 12/10/04. R/3, mySAP, and other SAP products and services mentioned herein are trademarks or registered trademarks of SAP AG in Germany and several other countries. For more information on this benchmark, please see the SAP benchmark Website (<http://www.sap.com/benchmark/>)

Manugistics Fulfillment v7.1 “Plan” Benchmark (Sun Fire 6800, E4900, and E6900 Servers)

The Manugistics Fulfillment v7.1 “Plan” benchmark simulates the processing of stock keeping units (SKUs) in a supply chain management (SCM) system. This retail batch-forecasting benchmark offers loads at varying levels. Vendors can run either a 10-million SKU model or a 20-million SKU model. In running the Manugistics Fulfillment V7.1 “Plan” benchmark, Sun ran single-system configurations for the 10-million SKU model and also used multi-system configurations to demonstrate scalability to handle the 20-million SKU model. As of this writing, Sun is the first vendor to run the 20-million SKU model in a benchmark, and the first vendor to achieve more than 30 million SKUs per hour.

Demonstrating Leadership and Single-system Scalability for the 10-Million SKU Model

For the 10-million SKU model, the Sun and Manugistics performance team completed a set of fulfillment benchmarks on a previous-generation Sun Fire 6800 server, a Sun Fire E4900 server, and a Sun Fire E6900 server. All of these results beat the nearest competing system, an IBM p690 that required 32 higher-frequency processors. The Sun Fire E4900 server beat this result with just 12 UltraSPARC IV processors, less than half the processors required by the IBM system (Table 3-2). In addition, the Sun Fire E6900 server processed almost 30 million SKUs/hour, the highest ever single-server result for this benchmark (as of this writing).

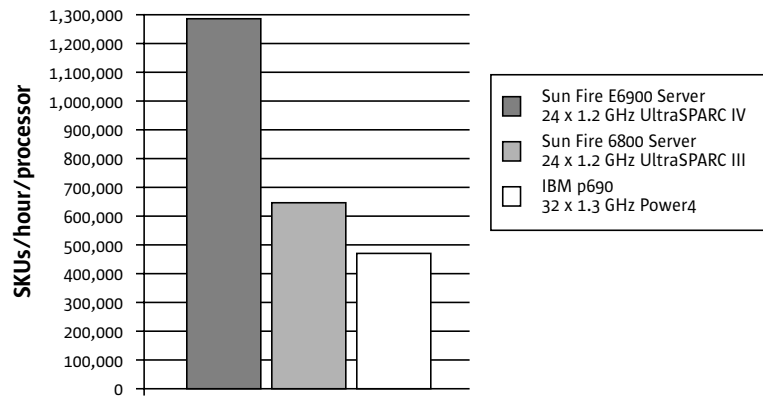
Table 3-2: Single-system Manugistics results for Sun Fire Enterprise servers demonstrate almost twice the throughput of previous-generation systems, and considerably better performance than IBM

System	SKUs/hour	SKUs/hour/processor	Number of processors	Processor type
Sun Fire E6900 server	29,751,044	1,239,627	24	UltraSPARC IV, 1.2 GHz
Sun Fire 6800 server	15,501,285	645,886	24	UltraSPARC III+, 1.2 GHz
Sun Fire E4900 server	15,428,165	1,285,680	12	UltraSPARC IV, 1.2 GHz
IBM p690	15,075,000	471,093	32	Power 4, 1.3 GHz

These results also demonstrate how the new Sun Fire servers nearly double the throughput of previous-generation Sun Fire systems. In the tests, a 24-processor Sun Fire E6900 server achieved 1.92 times the throughput of the previous-generation Sun Fire 6800 server. The Sun Fire E4900 server also nearly matched the results of the UltraSPARC III based Sun Fire 6800 server with only half the number of UltraSPARC IV processors. As shown in Figure 3-1, these results demonstrate:

- 2.6 times better per-processor performance than an IBM p690 server
- The new Sun Fire Enterprise servers equipped with UltraSPARC IV processors deliver almost twice (1.92 times) the real-world processing performance of previous-generation Sun Fire servers
- Continued performance leadership with supply chain management solutions
- The ability to scale resources along with SKU size to meet customer’s processing window

Figure 3-1: UltraSPARC IV equipped Sun Fire Enterprise servers deliver almost three times the per-processor performance of the IBM p690 and almost twice the per-processor performance of previous-generation Sun systems based on the UltraSPARC III processor



Demonstrating Multiple-System Scalability for the 20-Million SKU Model

In addition to single-system benchmarks, the Sun and Manugistics team completed a set of fulfillment benchmarks that demonstrate how Sun Fire servers can scale processing power and problem size. In these tests, the 10-million SKU model was run on both single and dual-system configurations (using the Sun Fire 6800 server). Additionally, the model was doubled to 20 million SKUs to test scalability and performance with a larger processing load. As shown in Table 3-3, these results demonstrate that Sun and Manugistics can scale in a near-linear fashion.

System	SKUs	SKUs/hour	SKUs/hour/ processor	Number and type of processors	Threads	Time(s)	Version
Dual Sun Fire 6800 servers	20 million	30,578,093	647,043	48 UltraSPARC III, 1.2 GHz	144	2366.4	v7.1
	10 million	30,495,617	635,325	48 UltraSPARC III, 1.2 GHz	144	1186.4	v7.1
Sun Fire 6800 server	10 million	15,529,230	647,051	24 UltraSPARC III, 1.2 GHz	96	2329.8	v7.1
IBM p690 server	10 million	15,075,000	471,093	32 Power 4, 1.3 GHz	128	2400	v7.1

Table 3-3: Single and multiple system benchmarks with both 10-million SKU and 20-million SKU benchmarks demonstrate near-linear scalability and capacity for very large workloads

Even though these tests were run on previous-generation 24-processor Sun Fire 6800 servers, the single-system results beat IBM's 32-processor results with fewer processors. In addition, these tests show that Sun can scale the problem size while allowing the runtime to remain constant with the use of multiple servers. As of this writing, no other vendors have attempted a problem on this scale and Sun is the only vendor with a Manugistics-certified success on the 20-million SKU database. Sun not only ran the larger model but demonstrated predictability, linear scaling, and great efficiency with a result that shows less than a 2-percent difference between the 20-million/48-processor result, and the 10-million/24-processor result. Further, these results demonstrate:

- 37.4 percent better per-processor performance than the IBM p690
- Leadership in being the first vendor to run the 20-million SKU fulfillment benchmark and the first vendor to break the 30-million SKUs/hour barrier
- Leadership in both overall performance and per-processor performance
- Ability to deal with the real-world complexity of retail supply chain management
- Ability to scale system resources with SKU size to meet defined processing windows
- Java technology scalability

SPEC JBB2000 (Sun Fire E6900 server)

The industry-standard SPECjbb2000 (Java business benchmark) is designed to evaluate the performance of server-side Java applications. The SPECjbb2000 benchmark evaluates a 3-tier system, the most common type of server-side Java application today where middle-tier activities predominate, including processing of business logic and object manipulation. In the SPECjbb2000 benchmark, clients are replaced by driver threads and database storage replaced by binary trees of objects. The benchmark consists of applying increasing amounts of workload, thus providing a graphical view of the scalability of the server platform under test.

SPEC JBB2000 Benchmark Results

The SPEC JBB2000 benchmark results (Table 3-4) demonstrate the considerable improvements in performance of the UltraSPARC IV based Sun Fire E6900 server relative to the UltraSPARC III based Sun Fire 6800 server and competitive systems — a direct result of chip multi-threaded UltraSPARC IV processors.

Table 3-4: SPEC jbb2000 Benchmark results for the mid-range Sun Fire E6900 server, the previous-generation Sun Fire 6800 server, and several competing high-end systems

System	Date	JVM	Number of processors	GHz	JBB ops/second	JBB ops/second/processor
Sun Fire E6900 server	1/04	32-bit	24	1.2	421,733	17,574
Sun Fire 6800 server	02/03	32-bit	24	1.2	231,121	9,630
IBM p690	12/01	64-bit	32	1.3	339,484	10,609
IBM p690	06/03	32-bit	32	1.7	553,480	17,296
Fujitsu PP250	4/03	64-bit	64	1.35	835,479	13,054

The results of this benchmark demonstrate:

- The Sun Fire E6900 server with UltraSPARC IV processors delivered a 1.83 times performance improvement over a Sun Fire 6800 server with the same number and speed UltraSPARC III processors.
- The Sun Fire E6900 server equipped with 1.2 GHz UltraSPARC IV processors outperformed the 32-processor 1.3 GHz IBM p690 by 24 percent.
- The 24-processor Sun Fire E6900 server also had better per-processor performance than the 32-processor 1.7 GHz IBM p690 and the 32 processor 1.35 GHz Fujitsu 2500 server.

Disclosure

Sun Fire E6900 server results were submitted to SPEC on 2/10/04 and are currently under review as of this writing. SPEC and the benchmark name SPECjbb are registered trademarks of the Standard Performance Evaluation Corporation. Competitive benchmark results stated herein reflect results published on www.spec.org as of February 10, 2004. For the latest SPECjbb2000 benchmark results, visit: <http://www.spec.org/osg/jbb2000/>.

SPEC OMPL2001 (Sun Fire E25K Server)

Ideal for HPC and scientific applications, the SPEC OMPL2001 benchmark provides a tool for evaluating performance based on large technical and scientific applications written to the parallel OpenMP application programming interface (API). The OpenMP API supports multi-platform shared-memory parallel programming in C, C++, and FORTRAN. The SPEC OMPL2001 benchmark applications are selected to reflect the needs of engineers and researchers to model large and complex tasks, and place heavy demands on systems and memory.

The SPEC OMPL2001 benchmark is designed to test large systems of from 32 to 256 processors and is comprised of the following codes:

- 311.wupwise_l — quantum chromodynamics
- 313.swim_l — shallow water modeling
- 315.mgrid_l — multi-grid solver in 3D potential field
- 317.applu_l — parabolic/elliptic partial differential equations
- 331.art_l — neural network simulation of adaptive resonance theory
- 321.equake_l — finite element simulation of earthquake modeling
- 329.fma3d_l — finite-element crash simulation
- 325.apsi_l — solves problems regarding temperature, wind, and distribution of pollutants
- 327.gafort_l — genetic algorithm code

The numbers in the benchmarks' names serve as identifiers to distinguish programs from one another and from similar codes in SPEC CPU2000.

Sun Fire E25K Delivers World-Record Performance for SPEC OMPL2001

The Sun Fire E25K server delivered world-record performance on the SPEC OMPL2001 benchmark suite. This performance was produced on a Sun Fire E25K server equipped with 72 1.05 GHz UltraSPARC IV processors. The result of 316,182 bested the previous world record holder (HP) by over four percent even though the UltraSPARC IV processors tested run at a 43-percent lower clock frequency than those of the HP system. These results demonstrate the power of Sun's chip multi-threading technology (Table 3-5).

Table 3-5: World-record SPEC OMPL2001 Benchmark results for the high-end Sun Fire E25K server and several competing high-end systems

System (Processor)	Active Threads	Base	Peak
Sun Fire E25K server (1.05 GHz UltraSPARC IV)	143	240,622	316,182
HP Integrity Superdome (1.5 MHz Itanium 2)	64	289,967	303,161
Fujitsu HPC2500 (1.3 GHz SPARC V)	124	262,140	—
SGI Origin 3900 (600 MHz R14K)	128	193,992	197,838

The upcoming Sun Studio 9 compiler was used to help achieve these outstanding results. As of this writing, the EA2 BETA release of the compiler is available for download at <http://www.sun.com>. The compiler also includes improved performance analysis tools that allow parallel programmers to see all of their threads at the same time. This feature allows programmers to see what part of their code is working well, and where they need to spend more time optimizing.

Disclosure

SPEC and the benchmark name SPECcomp2001 are registered trademarks of the Standard Performance Evaluation Corporation. Competitive benchmark results stated above reflect results published on www.spec.org as of February 4, 2004. The comparison presented above is based on the SPEC OMPL2001 metric. Sun Fire E25K server results were submitted to SPEC on 2/9/04 and are currently under review as of this writing.

Chapter 4

Conclusion

The Sun Fire Enterprise servers demonstrate that Sun is ideally-qualified to deliver the benefits of chip multi-threading through its throughput computing strategy. With throughput computing, Sun Fire Enterprise servers, and the UltraSPARC IV processor, Sun provides up to twice the *delivered* throughput and compute density of previous-generation systems while maintaining a high level of investment protection for current and future customers of Sun technology.

Sun's considerable experience with large symmetric multiprocessing architectures and the robust Solaris Operating System make the transition to throughput computing virtually transparent. Throughput computing is a natural extension of Sun's network computing philosophy of balanced system architecture, 64-bit processor and operating system support, as well as full thread support in both the operating system and Java technology. Through important innovations such as the Sun Java Enterprise System, Sun's customers can immediately leverage chip multi-threading performance gains without negatively-impacting applications or accepting gratuitous operating system migrations that disrupt operations.

For customers who want to experience the benefits of throughput computing for themselves, Sun offers an extensive range of Sun Infrastructure Solutions and Reference Architectures that can help speed design and deployment of complex services. Customers can also take advantage of Sun's iForce Solution Centers and iForce partners to benchmark and evaluate their own applications and data on a range of Sun Fire Enterprise Servers and related equipment. Coupled with these opportunities, Sun Enterprise Servers help take the cost and risk out of deploying scalable vertical computing solutions.

References

Sun Microsystems posts product information in the form of data sheets, specifications, and white papers on its Internet World Wide Web Home page at: <http://www.sun.com>. Look for the this and other Sun technology white papers:

Throughput Computing: Driving Down the Cost of Network Computing, White Paper, Sun Microsystems

SUN™ Copyright 2004 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054, U.S.A. All rights reserved.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

TRADEMARKS

Sun, Sun Microsystems, the Sun logo, Sun Fire, Solaris, Java, Sun HPC Clustertools, Sun StorEdge, and iForce are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

Postscript and Display Postscript are trademarks of Adobe Systems, Inc., which may be registered in certain jurisdictions.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. SUN MICROSYSTEMS, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS PUBLICATION AT ANY TIME.

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA Phone 1-800-555-9786 or 1-800-555-9SUN Web sun.com



Sun Worldwide Sales Offices: Africa (North, West and Central) +33-13-067-4680, Argentina +5411-4317-5600, Australia +61-2-9844-5000, Austria +43-1-60563-0, Belgium +32-2-704-8000, Brazil +55-11-5187-2100, Canada +905-477-6745, Chile +56-2-3724500, Colombia +571-629-2323, Commonwealth of Independent States +7-502-935-8411, Czech Republic +420-2-3300-9311, Denmark +45 4556 5000, Egypt +202-570-9442, Estonia +372-6-308-900, Finland +358-9-525-561, France +33-134-03-00-00, Germany +49-89-46008-0, Greece +30-1-618-8111, Hungary +36-1-489-8900, Iceland +354-563-3010, India-Bangalore +91-80-2298989/2295454; New Delhi +91-11-6106000; Mumbai +91-22-697-8111, Ireland +353-1-8055-666, Israel +972-9-9710500, Italy +39-02-641511, Japan +81-3-5717-5000, Kazakhstan +7-3272-466774, Korea +822-2193-5114, Latvia +371-750-3700, Lithuania +370-729-8468, Luxembourg +352-49 11 33 1, Malaysia +603-21161888, Mexico +52-5-258-6100, The Netherlands +00-31-33-45-15-000, New Zealand-Auckland +64-9-976-6800; Wellington +64-4-462-0780, Norway +47 23 36 96 00, People's Republic of China-Beijing +86-10-6803-5588; Chengdu +86-28-619-9333; Guangzhou +86-20-8755-5900; Shanghai +86-21-6466-1228; Hong Kong +852-2202-6688, Poland +48-22-8747800, Portugal +351-21-4134000, Russia +7-502-935-8411, Singapore +65-6438-1888, Slovak Republic +421-2-4342-9485, South Africa +27 11 256-6300, Spain +34-91-596-9900, Sweden +46-8-631-10-00, Switzerland-German 41-1-908-90-00; French 41-22-999-0444, Taiwan +886-2-8732-9933, Thailand +662-344-6888, Turkey +90-212-335-22-00, United Arab Emirates +9714-3366333, United Kingdom +44-1-276-20444, United States +1-800-555-9SUN or +1-650-960-1300, Venezuela +58-2-905-3800