

THROUGHPUT COMPUTING

Changing the Economics and Ecology of the Data Center
with Innovative SPARC[®] Technology

White Paper
November 2005

Table of Contents

Executive Summary	1
Business and Technology Challenges in the Age of Participation	2
Business Requirements	2
The Diminishing Returns of Complex Processor Design	2
Throughput Computing	4
CoolThreads Processor Technology for a Thread-Rich World	5
Hardware Multithreading	5
Chip Multithreading (CMT)	6
Sun's Unique Ability to Deliver Throughput Computing	7
A Range of Servers Based on SPARC Processors with CMT Technology	7
The UltraSPARC T1 Processor with CoolThreads Technology	7
Systems Based on the Proven UltraSPARC, IV+ Processor	7
The Solaris Operating System	8
Multithreaded Environments for Scalable Applications	9
A Seamless Application Environment for Customers and Developers	10
Conclusion	11

Executive Summary

Building on the foundation of the Information Age that preceded it, the emerging Participation Age is driving a fundamentally new set of requirements for information technology (IT). Converging IP, voice, and data networks, greater availability of bandwidth, and the emergence of web services and service oriented architectures (SOA) are all changing basic assumptions about how technology is used and deployed. The notion of connectivity too is changing, and network access is an increasingly common assumption for a wide range of disparate devices — from phones and PDAs to cars, parking meters, and packages of razor blades. Perhaps most importantly, many organizations are coming to grips with the significant energy resources used to power IT, yielding a growing understanding that ecological responsibility can impact the planet as well as the bottom line.

As recently as a few years ago, most organizations thought they had all the server capacity they were likely to need. Now growing demands along with dynamic and unpredictable usage patterns are driving dramatic needs for increased utilization, performance, capacity, and scalability in the data center. In fact, IDC now predicts a 50-percent increase in the installed base of servers in the United States alone by 2009, further estimating the annual power bill for that installed base at five billion dollars¹ (excluding cooling). Unfortunately, the need for infrastructure to support new and growing services is happening at a time when many data centers are facing hard limits of space, power, and affordability. Gartner now estimates that more than 80 percent of all data centers are already constrained by electrical power, physical space, or cooling capacity². Simply adding more of the same kinds of systems is clearly no solution, only contributing to new levels of consumption, complexity, and waste.

Sun's Throughput Computing initiative is designed to specifically address these seemingly contradictory trends — providing radically better levels of application and computational throughput while fundamentally changing the equation on space and power consumption. Chip multiprocessing and Sun's radical chip multithreading (CMT) processors with CoolThreads™ technology allow individual systems to deliver significantly more throughput in footprints that are the same or considerably less than previous-generation systems — enabling truly massive levels of simplification and consolidation. Simpler, cooler, and more efficient chip and system designs also dramatically lower both power and cooling bills, prolonging the useful life of many data centers while lowering dependence on shared energy resources. Efficient power and cooling is critical, as a recent survey indicated that data center managers report a 10 percent or more increase in power requirements over the past year and fully 41 percent reported that they anticipated upgrading their power and cooling systems over the next three years.³

Best of all, Sun is bringing this innovative new technology to market with the familiar SPARC® architecture in the form of the UltraSPARC® T1 processor with CoolThreads technology. Thanks to Sun's Binary Compatibility Guarantee and the robust Solaris™ Operating System, the increased resources and increased utilization from CMT processors are immediately available to existing applications and environments. With its strong research and development investments and extensive intellectual property in processor and system design, operating systems, and software, Sun is uniquely positioned to rapidly bring the benefits of Throughput Computing to its customers, without requiring so much as a recompile.

1. "Server Power Consumption Reemerges as a Critical Cost Factor in Datacenters", Vernon Turner, IDC, August 2005
2. Informal poll from Gartner's 2004 Datacenter Conference, as reported in "Dealing with a Less Than Perfect Data Center Location", by Johanna Ambrosio, Data Center Futures, TechTarget mailing, January 21, 2005
3. Joint study published by AFCOM and InterUnity Group in April 2005, and reported on at http://searchdatacenter.techtarget.com/originalContent/0,289142,sid80_gci1081270,00.html

Chapter 1

Business and Technology Challenges in the Age of Participation

Businesses today are increasingly defined by their applications, and now more than ever, an organization's prospects for success are increasingly fixed to its ability to deploy technology in an agile and effective fashion. The risks are extreme. In today's competitive and highly-regulated business environment, the cost of technology failure can be rapid and severe. Even small lapses in IT competence can result in wide-spread damage and loss.

Business Requirements

Increasing the pressure, an endless variety of new networked devices and users are demanding ever-higher levels of performance, capacity, availability, and security from the applications and services that serve them. Real estate concerns along with very real and rising energy costs for both power and cooling are now significant factors that discourage merely adding endless racks of traditional servers. The cost and complexity of managing very large numbers of systems is another pressing concern, especially when coupled with the very low levels of utilization typically found in traditional infrastructure.

To respond to these myriad challenges, business must:

- Increase application throughput along with capacity and performance to address pressing business needs as well as capture new customers and opportunities
- Reduce power, cooling, and real estate costs both to save money and to enable necessary growth and scalability
- Maintain application compatibility and enhance security across the organization to preserve investments and limit risks to the firm and its clientele

Beyond mere packaging, these issues drive to the very technology used to design processors, systems, and applications. Processor design in particular can have enormous ramifications for business-level issues and solutions. Unfortunately, traditional high-frequency, single-threaded processors are increasingly yielding diminishing returns. Even with ever-higher clock rates, these processors are producing only small improvements in real-world application performance. At the same time, these high-frequency processors generate escalating costs in the form of higher levels of power consumption, and significantly higher levels of heat load that must be addressed by multiple large and expensive HVAC systems. With economic and competitive pressures at an all-time high, most understand that significant change is needed.

The Diminishing Returns of Complex Processor Design

While optimistic marketing statements constantly call attention to presumably impressive multiple-gigahertz frequencies and high levels of cache for new generations of processors, corresponding small gains in real-world system performance and productivity continue to frustrate IT professionals. Throughput Computing, along with Sun's focus on optimizing real workload performance is designed to help resolve these divergent trends. This approach provides higher levels of *delivered* performance and computational throughput while greatly simplifying the data center. Understanding the importance of throughput computing requires a look at how both processors and systems have been designed in the past, and the trends that are defining better ways forward.

The oft-quoted tenant of Moore's Law states that the number of transistors that will fit in a square inch of integrated circuitry will approximately double every two years. For over three decades the pace of Moore's law has held, driving processor performance to new heights. Processor manufacturers have long exploited these chip real estate gains to build increasingly complex processors, with instruction-level parallelism (ILP) as a goal. Today these traditional 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 these techniques have produced faster processors with impressive-sounding multiple-gigahertz frequencies, they have largely resulted in complex, hot, and power-hungry processors that don't serve many modern applications, or the constraints of today's data centers. In fact, many of today's data center workloads are simply unable to take advantage of the hard-won ILP provided in these processors. As shown in Table 1, applications with high shared memory and data requirements are typically more focused on processing a large number of simultaneous threads (thread-level parallelism) rather than running a single thread as quickly as possible (ILP).

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

Workload Attributes	Web-Centric	Application-Centric	Data-Centric			
Benchmark	Web (SPECweb99)	Application (SPECjAppServer2002)	SAP-SD 2Tier	Data (TPC-C)	SAP-SD 3Tier (DB)	DSS (TPC-H)
Application category	Web server	Server, Java™	ERP	OLTP	ERP	DSS
Instruction-level parallelism	Low	Low	Medium	Low	Low	High
Thread-level parallelism	High	High	High	High	High	High
Instruction/Data working set	Large	Large	Medium	Large	Large	Large
Data sharing	Low	Medium	Medium	High	High	Medium

(SPEC, SPECweb, and SPECjbb are registered trademarks of the Standard Performance Evaluation Corporation (SPEC). SAP is a registered trademark of SAP AG in Germany and other countries. TPC-C and TPC-H are trademarks of the Transaction Processing Performance Council (TPC). For more information see www.spec.org, www.sap.com/benchmark, and www.tpc.org.)

Complicating matters, the disparity between processor speeds and memory access speeds means that memory latency dominates application performance, erasing even very impressive gains in clock rates. While processor speeds continue to double every two years, memory speeds have typically doubled only every six years. This growing disconnect is the result of memory suppliers focusing on density and cost as their design center, rather than speed. Unfortunately, this relative gap between processor and memory speeds leaves ultra-fast processors idle as much as 85 percent of the time, waiting for memory to return required data. Ironically, as traditional processor execution pipelines get faster and more complex, the effect of memory latency grows—fast, expensive processors spend more cycles doing nothing. Worse still, idle processors continue to draw power and generate heat. Its easy to see that frequency (gigahertz) is truly a misleading indicator of real performance.

Figure 1 illustrates how even doubling processor performance (frequency) often provides 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 and dominant influence of memory latency.




Figure 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

While some vendors have seemingly awakened to the inherent limitations of traditional, frequency-based processor designs, they are now attempting to graft power-saving technologies and multiple cores onto old, once-discarded architectures. Unfortunately, these efforts represent stop-gap measures at best. Effective Throughput Computing can only be realized with fundamentally new processor designs that deliver truly compelling benefits to customers while leaving legacy approaches behind.

Throughput Computing

Put simply, Throughput Computing is concerned with the ability to get more done within the constraints of space, power, cooling, and time. Beyond any particular technology, the focus is necessarily one of providing users and applications with more available resources, while affording the business maximum flexibility. Specifically, Throughput Computing serves businesses and their customers by:

- Maximizing throughput to meet growing workload demands driven by increased participation
- Reducing operating costs through space and power consumption, while helping to enable growth within current data center thermal and power limits
- Facilitating effective consolidation to drive simplification and increase system utilization while further reducing requirements for space, power and cooling
- Improving agility and efficiency to align IT with the needs of the business
- Providing investment protection and reduced risk through binary compatibility and consistent architecture

Though technology for Throughput Computing necessarily resides in everything from system design to OS and application technology, getting more done requires a special focus on resolving the challenges posed by traditional processor design.

Chapter 2

CoolThreads Processor Technology for a Thread-Rich World

Dissatisfied with the performance characteristics of traditional single-threaded processors, Sun's extensive in-house design team—one of the largest microprocessor design engineering teams in the world—has taken a bold new approach to processor design. Sun understands that the network-computing environments found in most modern data centers are inherently multithreaded, where the execution speed of an individual thread is typically less important than overall application throughput. For this reason, Sun is focusing on processors and architectures that maximize throughput for commercial network-computing workloads. These efforts are resulting in new chip multithreaded (CMT) processor technology that leverages the additional gains delivered by Moore's Law to provide thread-level parallelism rather than instruction-level parallelism.

Recently, a number of processor vendors have started making multi-core processors available that place multiple processors on a single die (so-called chip multi-processing). Unfortunately, many of these efforts too have resulted in complex, hot, high-power processors that don't solve the real issues driving processing and productivity in the data center. The most significant potential benefits of multi-core technology derive from a fundamental simplification and rethinking of how processors are designed and built.

Hardware Multithreading

Unlike traditional single-threaded processors and even most current multi-core processors, hardware multithreaded processors allow rapid switching between active threads as other threads stall for memory. Figure 2 illustrates a hypothetical multithreaded processor core that is designed to switch between up to four threads as each thread stalls for memory. With this approach, the processor's execution pipeline remains active doing real useful work, even as memory operations for stalled threads continue in parallel.

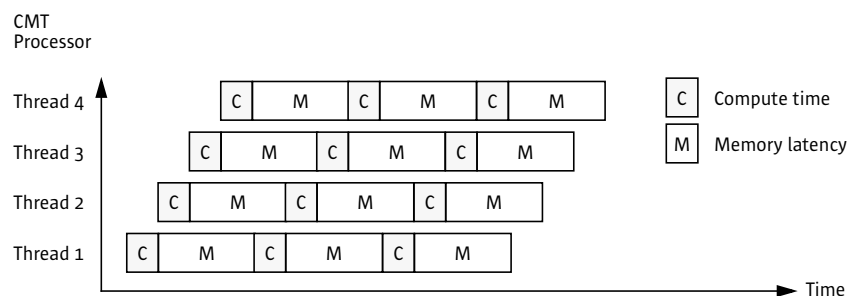


Figure 2. A hypothetical multithreaded processor core switches between a number of active threads, doing useful work even while threads stall to perform memory-related operations

Hardware multithreading provides real value since it increases the ability of the execution pipeline to do actual work. Utilization of the processor pipeline is greatly enhanced since a number of execution threads now share its resources. The negative effects of memory latency are masked since the processor and the memory subsystems remain active in parallel to the processor execution pipeline.

Chip Multithreading (CMT)

Like single-threaded processors, multi-processing (multi-core) technology can also be used to scale and multiply the benefits of hardware multithreading. Sun calls the result chip multithreading (CMT). Unlike complex single-threaded processors, CMT processors utilize the available transistor budget to implement multiple hardware multithreaded processor cores on a single silicon wafer or chip. Because these individual processor cores implement much simpler pipelines (emphasizing TLP over ILP), they are also substantially cooler and require significantly less electrical energy to operate. This innovative approach results in CoolThreads processor technology—multiple physical instruction execution pipelines (one for each core), with several active thread contexts per pipeline or core.

Application throughput is greatly improved with CoolThreads processors, as is utilization of pipeline resources. Thread-rich applications common in commercial workloads benefit greatly from this model, whether comprised of larger multithreaded applications, or of large numbers of smaller single-threaded applications. The number of simultaneous threads that can be accommodated is quite large, and a wide range of processor designs are possible. Figure 3 illustrates a hypothetical example where four hardware multithreaded cores are combined into a single CMT processor, supporting up to sixteen active execution threads.

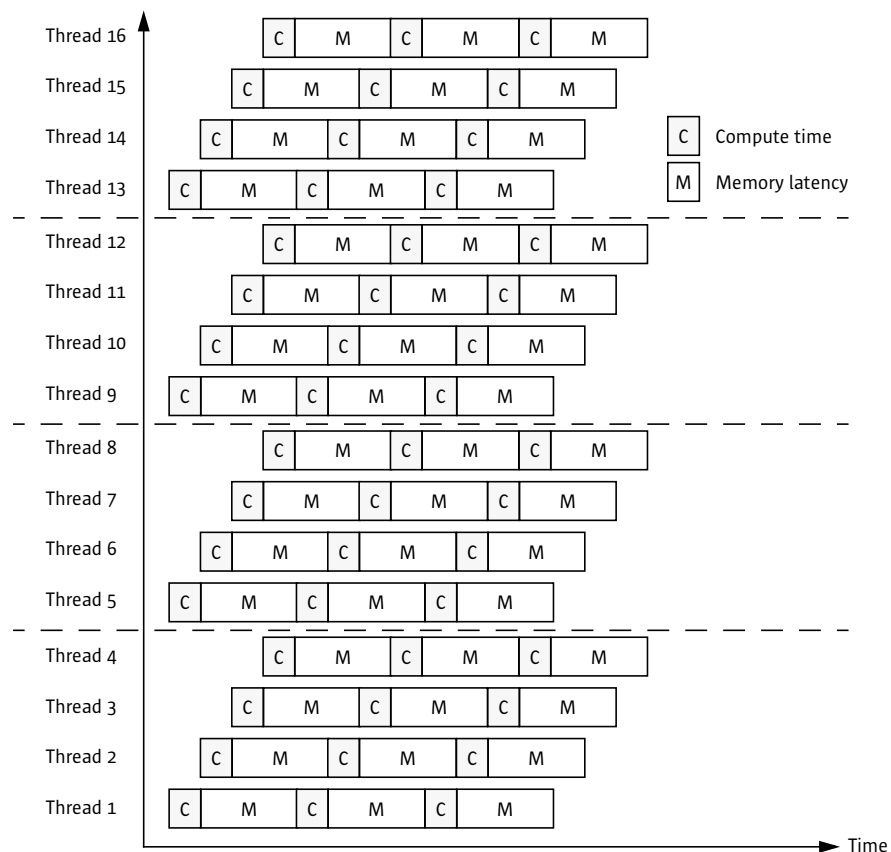


Figure 3. A hypothetical CMT design supporting four multithreaded cores and 16 active thread contexts

CMT technology essentially results in implementing a multithreaded Symmetric Multiprocessing (SMP) system on a single chip—an environment that is extremely familiar to Sun.

Chapter 3

Sun's Unique Ability to Deliver Throughput Computing

Deploying radical new technology in a seamless fashion requires a systematic approach, one facilitated by Sun's end-to-end technology ownership. In fact, Sun is delivering the considerable benefits of CMT processors in a form that can be leveraged directly and immediately, without requiring so much as a recompilation. Sun's broad set of intellectual property and continuous multi-billion dollar focus on research and development make it uniquely qualified to deliver the benefits of Throughput Computing. Without control of significant processor or operating-system technology, this approach would be impossible for other system vendors to duplicate.

A Range of Servers Based on SPARC Processors with CMT Technology

Sun has long been an innovator with SMP technology—both in terms of systems and operating systems—building a wide range of robust multiprocessor SPARC systems that deliver considerable scalability. Sun is keenly focused on building systems that specifically address real-world workloads, a focus that has led to multiple implementations of the SPARC processor architecture to directly address the needs of different workloads and customers.

The UltraSPARC T1 Processor with CoolThreads Technology

Sun's new UltraSPARC T1 processor with CoolThreads technology combines chip multi-processing (up to eight cores per processor) and hardware multithreading (up to four threads per core) with simple and efficient, fully-standard SPARC V9 instruction pipelines. In particular, a single UltraSPARC T1 processor provides support for up to 32 active threads on two square inches of chip real estate—requiring only the power of a light bulb to operate. These processors set a new standard for extreme throughput, performance, and breakthrough power savings.

As a full SMP system on a chip, a single UltraSPARC T1 processor offers dramatic simplification and increased reliability. Systems based on these processors are expected to deliver not just industry-leading performance on a per-system and per-rack basis, but also industry-leading performance per watt as well as industry-leading performance per rack unit. The vastly superior energy and thermal characteristics of UltraSPARC T1 processors result in extremely low power requirements and produce relatively little heat.

Systems based on the UltraSPARC T1 processor are expected to be ideal for commercial applications, especially those powering web and transaction services. In addition to providing ideal platforms for Java application servers and enterprise application servers (e.g. ERP and CRM) these systems will enable considerable web-tier consolidation in the data center, dramatically reducing the number of servers required to deliver a given workload while requiring less power and space. At the same time, these reductions will help make room for the expansion of services and content in the data center as required by the business.

Systems Based on the Proven UltraSPARC® IV+ Processor

Extending the successful UltraSPARC® III and UltraSPARC® IV processors, the new UltraSPARC IV+ processor enables existing customers of Sun's mid-range and high-end Sun Fire™ servers to seamlessly increase performance by up to five-fold. In concert with advanced features such as Dynamic Reconfiguration, these systems can be upgraded “on the fly” to new UltraSPARC IV+ processors, even as they continue to operate and deliver application and database services.

The UltraSPARC IV+ processor offers a completely new design, with enhanced processor cores, a new three-level cache hierarchy, and 90nm technology. In spite of these advances, the UltraSPARC IV+ processor remains pin-compatible with earlier UltraSPARC III and UltraSPARC IV processors. Depending on the application, Sun Fire V490-E25K servers with UltraSPARC IV+ processors can deliver up to five times the performance of systems equipped with UltraSPARC III processors and up to twice that of systems equipped with UltraSPARC IV processors. Perhaps most importantly for asset protection, UltraSPARC IV+ processors can operate in existing Sun Fire V490-E25K systems, side-by-side with UltraSPARC IV and UltraSPARC III processors¹—with each running at their rated speed.

Sun Fire servers based on UltraSPARC IV+ processors are optimized for large data and application-centric workloads such as those required by large applications and databases, and these systems have already set a number of world records in industry benchmarks (http://www.sun.com/servers/sparc_benchmarks/). These systems represent Sun's most scalable servers, and with the UltraSPARC IV+ processor they deliver breakthrough performance and leading investment protection with no increase in footprint or power consumption, and no increase in price.

The Solaris Operating System

It's one thing to have a fast multithreaded processor or a well-designed system, but those abilities are wasted if the operating system can't deliver available processor resources to the application. Nowhere is Sun's commitment to scalability and multithreading more apparent than in the Solaris Operating System. This proven scalability for SMP systems lets multithreaded applications and their data grow to almost arbitrary size, limited only by the platforms that host them. The Solaris 10 OS also provides innovative features that help application performance through better utilization, virtualization, tuning, and management of the resources of CMT processors as well as applications and data. The Solaris OS also provides comprehensive security that is built in, not bolted on.

- **Extensive threading support**—Because of Sun's long lineage in SMP systems, the Solaris OS was designed from the beginning to support considerable multithreading abilities. This threading support serves both to help ensure that the operating system scales to very large numbers of processors and threads, and that individual multithreaded applications can scale effectively as well. For example, the Solaris 10 OS provides threaded network drivers to provide the scalability needed to support the considerable network throughput required by multiple applications running over multiple high-speed network connections.
- **Solaris Containers**—In addition to the extensive threading support in the Solaris 10 OS, Solaris Containers technology allows the partitioning and virtualization of the considerable resources of CMT processors. Sun's Dynamic System Domains already provide the ability to dynamically partition a single Sun Fire mid-range server into multiple independent operating environments. Now Solaris Containers take the server utilization concept to a new level by allowing *all* servers running the Solaris 10 OS (including individual Dynamic System Domains) to be partitioned to sub-CPU granularity, without replicating the Solaris OS image.

Each Solaris Container provides a virtualized Solaris OS image, including a unique root file system, a shared set of system executables and libraries, and whatever resources the root administrator assigns to the container when it is created. Each Solaris Container even has its own IP address, and individual containers can be booted and shut down—rebooting in seconds if the need arises. Solaris Containers can also help enhance security and isolation to help enable more effective consolidation efforts. Because of the low overhead of Solaris Containers, individual applications can be given their own private container and environment, virtually eliminating error

¹.Sun Fire V490 and V890 servers support mixing and matching of UltraSPARC IV+ and UltraSPARC IV processors

propagation, intentional and unintentional intrusions, and unauthorized access—both to other containers and to the system at large.

- **Dynamic tracing with DTrace**—When production systems exhibit nonfatal errors or sub-par performance, the sheer complexity of modern distributed software environments can make accurate root-cause diagnosis extremely difficult. At the same time, the inability to meet anticipated user load or the failure to meet service-level requirements can be very costly to the business. Unfortunately, most traditional approaches to solving this problem have proved time-consuming and inadequate, leaving many applications languishing far from their potential performance levels.

The Solaris DTrace facility provides dynamic instrumentation and tracing for both application and kernel activities—even allowing tracing of application components running in a Java Virtual Machine (JVM™)¹. DTrace enables developers and administrators to explore the entire system to understand how it works, track down performance problems across many layers of software, or locate the cause of aberrant behavior. Tracing is accomplished by dynamically modifying the operating system kernel to record additional data at locations of interest. Best of all, although DTrace is always available and ready to use, it has no impact on system performance when not in use, making it particularly effective for monitoring and analyzing production systems.

Perhaps most importantly, the Solaris OS supports all Sun SPARC systems with the Solaris Binary Compatibility Guarantee, helping to preserve investments in existing applications and training while enabling organizations to seamlessly take advantage of innovative CMT technology.

Multithreaded Environments for Scalable Applications

CMT processors and Throughput Computing are designed to deliver more resources to applications that are throughput oriented, namely those that seek to accomplish a large amount of work in aggregate. A number of classes of applications benefit directly from the ability to scale throughput with CMT processors:

- **Multithreaded native applications** — Multithreaded applications are characterized by having a small number of highly-threaded processes. Examples of threaded applications include Lotus Domino or Siebel CRM.
- **Multi-process applications** — Multi-process applications are characterized by the presence of many single-threaded processes. Examples of multi-process applications include the Oracle database, SAP, and PeopleSoft.
- **Java applications** — Java applications embrace threading in a fundamental way. Not only does the Java language greatly facilitate multithreaded applications, but the Java Virtual Machine is a multithreaded process that provides scheduling and memory management for Java applications. Java applications that can benefit directly from CMT resources include application servers such as Sun's Java Application Server, BEA's Weblogic, IBM's Websphere, or the open-source Tomcat application server. All applications that use a Java™ 2 Platform, Enterprise Edition (J2EE™ platform) application server can immediately benefit from CMT technology.
- **Multi-instance applications** — Even if an individual application does not scale to take advantage of a large number of threads, it is still possible to gain from CMT architecture by running multiple instances of the application in parallel. If multiple application instances require some degree of isolation, Solaris Containers technology can be used to provide each of them its own separate and secure environment.

1. "The terms "Java Virtual Machine" and "JVM" mean a Virtual Machine for the Java™ platform.

A Seamless Application Environment for Customers and Developers

Far from typical burdensome migrations and ports required by other vendors' technology, Throughput Computing with CMT technology is remarkable in its simplicity. With Sun's binary compatibility guarantee, applications written for the SPARC/Solaris platform simply run with no modification whatsoever. Customers not only don't need to replace their existing applications, they also avoid costly re-deployment and the associated training costs. In addition, the extensive catalog of existing Solaris/SPARC applications is immediately available to systems based on CMT technology.

Throughput Computing reflects a massive simplification for developers as well. Because applications move to CMT technology without so much as a recompilation, developers avoid expending time and money better reserved for adding functionality to their products. Developers need only qualify and test for a specific version of the Solaris OS in order to encompass all of the platforms supported by that OS version.

Chapter 4

Conclusion

In concert with the demands of the today's complex data centers and the needs of the planet, Sun's Throughput Computing strategy offers a compelling approach to designing, building, and deploying server platforms and infrastructure that succeeds in surmounting many of the limitations of traditional processor and system design. Consistent with Sun's eco-responsibility stance, these innovations promise to deliver new levels of performance for throughput-oriented workloads typical of commercial applications, while fundamentally changing data center economics in terms of space, power, and heat production.

An effective Throughput Computing strategy could only come from an organization with considerable investments and experience in a broad range of technologies, including processor and system design, operating system leadership, and threaded application environments such as those enabled by Java technology. Sun's investments in processor technology in particular have yielded UltraSPARC T1 processors with CoolThreads technology that—in spite of their radical new design—preserve customer's investments in applications, software and training, even as they facilitate new levels of flexibility, agility, and environmental responsibility for the business.

Because Sun stands behind its promise of binary compatibility for the proven and secure Solaris OS, both native applications as well as Java applications simply run on Sun systems equipped with CMT processors, with no modification whatsoever. As a result, these processors and systems can deliver on their potential immediately—providing vastly improved throughput in existing or considerably smaller physical footprints with drastically improved power and cooling characteristics. Together Sun's Throughput Computing innovations help ensure that—along with the Information Age—Sun will help enable the Participation Age.

Copyright © 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A.

All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Sun, Sun Microsystems, the Sun logo, CoolThreads, Solaris, Java, Sun Fire, JVM, and J2EE are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a). DOCUMENTATION IS PROVIDED AS IS AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS HELD TO BE LEGALLY INVALID.