

WHITE PAPER

Sun's Throughput Computing Strategy to Create a Quantum Change in Server Performance

Sponsored by: Sun Microsystems, Inc.

Vernon Turner
February 2004

Christopher G. Willard, Ph.D.

EXECUTIVE SUMMARY

Sun Microsystems (Sun) has announced that a quantum change in server performance is on its horizon, and the impact on enterprise IT strategy will be significant. Sun's *throughput computing* initiative is based on new thinking about processor design and a sharp focus on improving the actual workloads that processors are able to support. Between 2005 and 2007, Sun plans to introduce processors that offer 15–30 times the throughput of current processors and that cross the entire breadth of its server product line.

Sun is pursuing a solution to the industrywide problem that plagues server design: Processor speeds are outrunning memory access by nearly four orders of magnitude. As a result, increases in processor clock rates do not lead to servers with greater throughput because the processors are stalled waiting for data to be read and written to memory. Server components are out of balance and memory performance is throttling overall system performance.

Sun's innovation is the use of chip multithreading (CMT) processors that manage two or more threads concurrently and can quickly switch between threads (i.e., independent software processes) when the active thread stalls. In addition to switching among threads, CMT processors will have multiple cores — that is, two or more separate arithmetic logic and control units that are each capable of executing instructions. Throughput improvements are expected to be proportional to the product of the number of threads, the number of cores, and the performance of each core.

Sun intends to leverage CMT at the processor level with its extensive capabilities for processing threads in the operating and software environments. Sun's support for threads goes back to 1988, when it shipped SunOS 4.1 with the first lightweight process (LWP) library, which allowed applications to use threads. The Solaris operating system has supported threading since its release in 1992, and Java 1.0 provided support for threads when it was released in 1995. The Sun Java Enterprise System (formerly known as SunONE) also supports threading.

IDC believes that Sun's CMT strategy and focus on throughput computing have the potential to support a quantum leap in server performance. Sun has the opportunity to exploit tight integration throughout the system — among its CMT processors, the servers that incorporate CMT processors, and its operating environment and software infrastructure.


IDC encourages IT planners to focus on throughput as the key indicator of server performance and to monitor Sun's progress in delivering on its CMT strategy. Furthermore, IT planners should examine the queue of IT service ideas that do not currently meet ROI forecasts based on today's cost of processing. Successful companies will be those that apply tomorrow's more powerful servers to their competitive advantage.

TRENDS IN INFORMATION TECHNOLOGY

Change is afoot in the information technology (IT) marketplace in general and in the server marketplace in particular. IDC estimates that the installed base of networked computing devices is approximately 1 billion and growing. From a broader perspective, the appearance of new categories of devices will expand market boundaries and enable entirely new IT solutions. IDC forecasts that the worldwide installed base of networked *computers* in 2012 will grow by a factor of 1,000 to be 1.08 billion units, as Figure 1 shows. This number will grow to 6 billion when handheld, entertainment, and industrial/automotive devices are included and to more than 16 billion units when appliances and toy devices are added to the mix. Finally, when radio frequency identification (RFID) *sensors* and *tags* are included, the total number of networked devices will exceed 1 trillion.

FIGURE 1

Worldwide Installed Networked Devices in 2012



Source: IDC, 2004

The ability to support billions of networked devices has significant implications for infrastructure requirements in the next decade. All IT systems and technologies will be challenged to manage the expanding network. Our special interest is in whether servers will be capable of taking on significantly greater workloads in the future.

IDC believes that server technology is approaching a workload performance plateau. Processors in today's servers have evolved from personal workstation processors. These processors were designed initially for a single user and now support multiple workloads as IT departments consolidate applications and partition and provision server resources.

Servers supporting new Web services architectures are expected to support thousands of users. The increasing demand for simultaneous support of multiple processes creates an opportunity for new technology to enter the market — technology that provides an alternative to incremental improvements in, or scaling of, current technology.

Quantum Change

In the IT industry, steady, incremental improvements in products and services are punctuated from time to time by quantum change. The introduction of RISC processors is one example; the move from single-processor servers to SMP servers is another. Performance curves reflect a discontinuity as engineers take new design approaches. Such quantum changes in computer performance and price/performance have affected the market, first, by allowing existing users to address next-generation applications, and second, by making computer technology available to new classes of users.

While processors continue to gain speed exponentially, the overall throughput of servers is not improving proportionally. Concern about this problem is industrywide. Attempts to work around the problem have spawned a collection of higher-level architectures using multiple processors; NUMA, clusters, blade servers, and grids all attempt to improve performance with multiple independent or semi-independent memory systems.

As improvements in semiconductor process technology provide engineers with increasing transistor budgets, IDC expects to see different strategies emerge to solve the problem, which, simply stated, is that processors supporting typical server workloads are spending far too many instruction cycles in a wait state. Processor architectures will be improved with a mixture of insights based on historical improvements in processor performance and new thinking about processor design.

Quantum Change: Implications for Business

When innovation leads to quantum change and reshapes the competitive landscape, it also provides extraordinary opportunities for businesses. Namely, an order of magnitude change in product performance enables new opportunities within the enterprise where IT can be put to work.

- ☒ The ability to reach out to individuals and personalize products for the mass market was initially enabled by another quantum change: the ubiquitous Internet. In the near future, the ability to provide personalized products and services again will advance dramatically. Supported by digital cellular devices, wireless networks, and high-capacity servers, innovative companies will reach out to customers wherever they are providing specialized information on demand.
- ☒ RFID further extends the reach of IT systems by identifying objects when they are in the proximity of a sensor. Early adopters are using RFID to automate the tracking of materials and products both within the enterprise as well as in the supply chains that link trading partners and customers. The use of RFID extends computer networks beyond intelligent devices to include devices that have no inherent information processing capabilities. More sophisticated RFID devices will carry additional information such as how an object is to be handled (e.g., refrigerated, kept upright, routed with highest priority).

As competition increasingly requires personalization, and computer networks extend into new classes of devices, the volume of data generated will increase exponentially. Enterprise IT organizations will be challenged by new requirements to manage this data and new opportunities to deploy IT systems that create competitive advantage.

Balancing Processing Speeds and Memory Access

Moore's law continues to drive significant improvements in computer system component technologies, but not at the same pace for all components. As a result, computing system architectures continue to flex as they reflect new *balance points* among key technologies: processing, memory, storage, display, power, network, and software.


The issue of shifting balance points is increasingly important for microprocessors and memory both in workstations and especially in servers. As described by Moore's law, the density of semiconductor devices continues to increase. For microprocessors, smaller transistor geometries have led to higher switching speeds, increased transistor budgets, and the shift to 64-bit processing. For memory, Moore's law has improved capacity. However, the relationship between processor speed and memory latency (i.e., the time it takes for data from memory to reach the processor) has not followed the same growth path and is out of balance. Namely, processors increasingly outrun the ability of memory to supply the processor with data.

The Memory Wall

The imbalance between processor and memory performance was noted by W.A. Wulf and S.A. McKee in their 1995 article, "Hitting the Memory Wall: Implications of the Obvious." Wulf and McKee observed that while processor speed and memory latencies were both improving exponentially, the exponential functions were not the same, and the difference was itself growing as an exponential function. Since 1980, according to a paper authored by M.V. Wilkes in 2001, memory speeds have improved by approximately 10% each year while processor speeds have improved by about 60%. As a result, following Wilkes' model, processors will soon outrun memory by nearly four orders of magnitude, as Figure 2 illustrates. These investigators conclude that faster processor technologies will have little or no effect on overall system throughput.

FIGURE 2

RELATIVE PERFORMANCE OF PROCESSORS AND MEMORY



Source: IDC, 2004 (based on M.V. Wilkes, "The Memory Gap and the Future of High Performance Memories," 2001)


Processing and Memory Access

The consequence of the imbalance between processor and memory is illustrated in Figure 3. The time-series plot of processing and memory access shows that Processor A divides its time between processing steps and pipeline stalls while data is written and retrieved from memory. Processor B is a faster processor, and memory latency remains the same. Figure 3 illustrates two important lessons:

- Increases in processor speed do not appreciably improve throughput when memory latency remains fixed and large compared with processing time.
- Fast processors outrun memory and remain in a stalled state for much of the time.

FIGURE 3

PROCESSOR SPEED AND MEMORY LATENCY



Source: IDC, 2004

Workloads and Processors

Different types of computer workloads generate different sets of requirements for computer processors and associated systems. Desktop systems can be viewed as single task computers in which the processor must be powerful enough to effectively address the user's most complex task (e.g., picture editing, encryption of Internet messages). However, underutilization of a desktop system is not an issue; it is acceptable to turn the system off and waste the cycles. Thus, for the desktop, processor cycles need to be free and plentiful on an as-needed basis.

In contrast, servers are designed to support large and varied workloads and to operate on a 24 x 7 basis. For servers, single task performance is not as important as the total task completion in a given amount of time (i.e., throughput). In addition, system efficiency, the server's price/performance ratio, becomes an issue. The more work a given server can perform, the fewer servers the enterprise will need to purchase, house, maintain, and manage.

The combination of improvements in processor speed and system scalability has allowed desktop processors to address many server-oriented workloads on the basis of acceptable task-level performance and good price/performance (i.e., if cycles are cheap enough, efficiency doesn't matter). However, as the memory-wall analysis shows, efficiency losses are growing rapidly as processor and memory systems become more out of balance.

Sun's Processor Strategy for Throughput Computing

Sun has developed a technology strategy for addressing future server throughput computing requirements. A major component of this strategy is the use of new approaches to processor architectures. The major technology that Sun is implementing at the processor level is CMT.

A major driver of Sun's strategy is the concept of thread-level parallelism (TLP). A *thread* is an independent software process consisting of a sequence of instructions to be executed by a processor. Threads occur within workloads as multiple independent applications (i.e., two separate user jobs) as well as within a single application (i.e., parallelism inherent to distinct tasks). TLP provides an opportunity for organizing simultaneous processing so long as there are multiple threads in the workload environment.

Thread-rich workloads are most common to servers, where the computer is a shared device supporting multiple applications and multiple workloads. Applications delivered as Web services are particularly rich in threads as multiple users launch independent processes with each browser click. Applications such as database queries, once multiforked in design, are now multithreaded.

All thread-rich workloads are not created equal. Sun divides the network computing space into network-facing and data-facing workloads. Network-facing workloads, typical of Web servers, are characterized by asynchronous, stateless threads that are I/O driven. Data-facing workloads, such as database servers, are characterized by stateful, often synchronized threads and are memory driven or compute driven. CMT processors for network-facing workloads will emphasize the ability to handle many threads in parallel at the expense of single-thread performance, while CMT processors for data-facing workloads will require more robust single-thread performance.

Freeing Up Processor Real Estate

When processor designers relax the requirement for the fastest possible single-thread performance, a significant amount of a processor's silicon "real estate" can be opened up for other uses. Based on early laboratory results, Sun told IDC, "A core [i.e., the basic hardware compute engine] providing half the single-thread performance of today's fastest and most complicated processor cores can be built using about 10% of the real estate." Doing so frees up 90% of the silicon for other uses, such as multithread support and integrated network and security functionality, depending on the workload.

Efficient Use of Clock Cycles

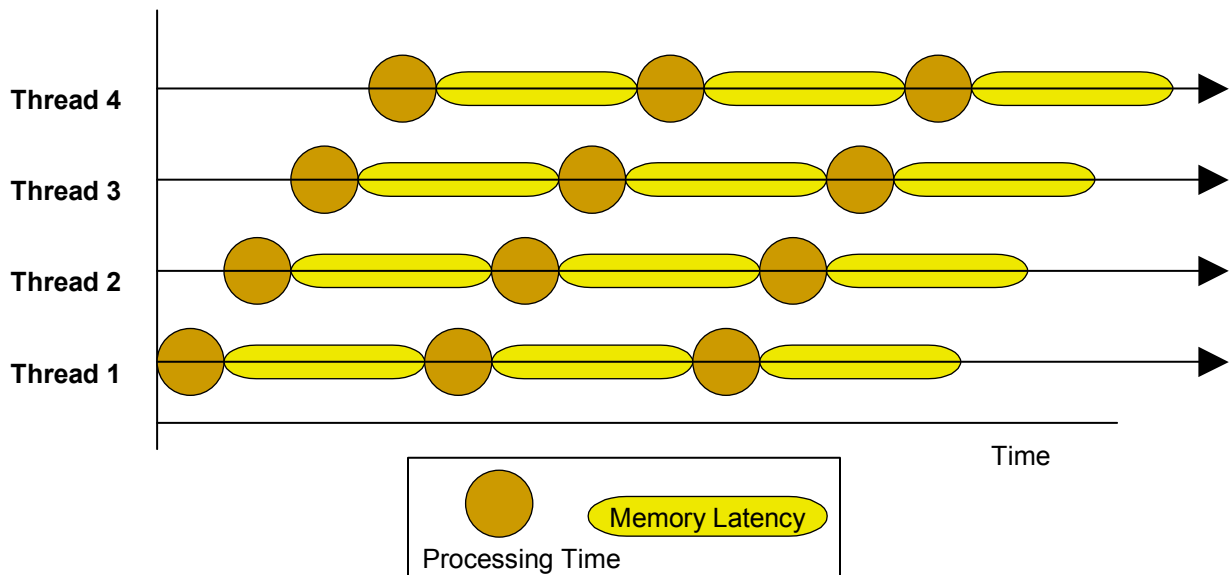
CMT processors can use some of the available real estate to implement mechanisms that can reduce the amount of clock cycles wasted in wait states. Specifically, CMT processors that manage two or more threads concurrently can quickly switch between threads when the active thread stalls. The goal of multithreading strategies is "latency hiding." Instead of attempting to reduce memory latency times, multithreading simply allows the processor to jump to a ready-to-run thread as soon as the current or active thread stalls for lack of data. The technical challenge is to

provide a *zero cycle context switch*, a processor that can move among threads with the absolute minimum of overhead. Figure 4 illustrates how multithreading works:

- ☒ A processor simultaneously holds states for multiple threads with the ability to switch among threads within a single clock cycle.
- ☒ When a currently executing thread stalls on a memory reference, the processor switches to the next runnable thread and begins executing it.
- ☒ When the second thread stalls on a memory reference, the processor switches to a third runnable thread and then to a fourth runnable thread when the third thread stalls.
- ☒ By the time the fourth thread stalls, the first thread has likely completed its memory access and is ready to resume running.

FIGURE 4

MULTITHREADING TO INCREASE THROUGHPUT



Source: IDC, 2004

While at any given time several threads on the processor will be in a wait state, one or more other threads will likely have the data needed from memory and execute. Ideally, the processor is always doing useful work by switching between threads whenever a running thread is stalled by a memory reference. Thus, although the performance of any given thread is not changed, the total number of threads completed in a given time period is increased.

Multithreaded processor performance is bounded by such factors as:

- ☒ The number of threads that a processor can physically manage
- ☒ The availability of threads (i.e., how well the applications and systems software is able to identify and manage threads)
- ☒ Memory performance (i.e., multithreading acts to hide the effects of memory latency but in doing so places greater demand on processor to memory bandwidth)
- ☒ Processor speed (i.e., multithreading allows the processor to use as much of its available resources as possible; however, the performance of those resources is ultimately bounded by processor speed)

The multithreaded approach addresses the imbalance between processing speed and memory latency. Together, a server with processors supporting multithreading and an operating environment capable of managing threads could provide significantly greater system throughput. The potential effect of multiple execution of multiple threads is shown in Figure 5. To improve throughput, system designers must meet the following criteria:

- ☒ Processor designers will need to allocate semiconductor resources (i.e., the transistor budget) to include support for the queuing and execution of a multitude of threads.
- ☒ Support for threads will be needed in the operating system, which will need to be multithreaded and tightly coupled to the processor's multithreading functionality.
- ☒ Support for threads will also be needed in the software infrastructure used to develop the applications that carry enterprise workloads. That is, support for multiple threads must reach from the processor hardware to the top of the application stack and deliver performance improvements to the enterprise.

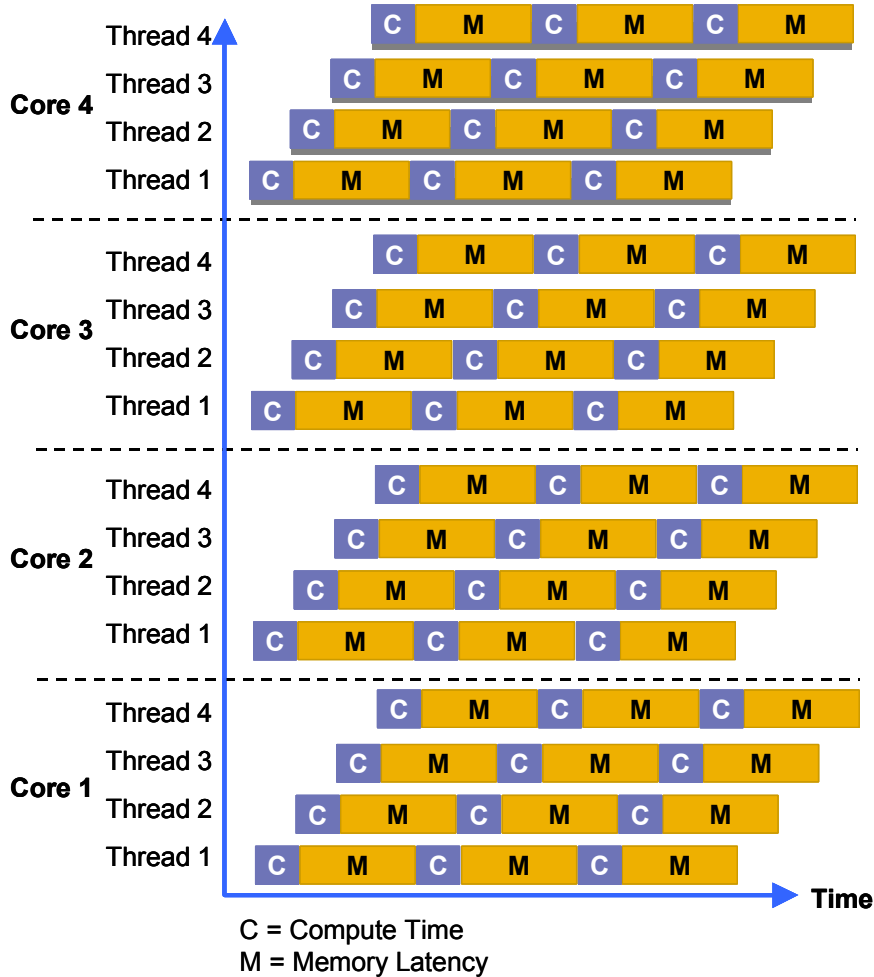
Solaris and Java Support Multithreading

Multithreading cannot be put to use without an operating environment capable of managing threads. Starting with Sun OS 4.1 in 1988, Sun shipped its first library that provided operating system support for threads. Through the 1990s, Sun continued to enhance its thread functionality in Solaris and to participate in standardizing the Sun approach as part of the POSIX standards. By 2000, with the release of Solaris 9, Sun's optional threading model became a standard model for Solaris users.

Not surprisingly, when Java 1.0 was released in 1995, the initial definition of the language included support for threading. Improvements to the Sun Java Enterprise System platform and the new HotSpot Java Virtual Machine demonstrate Sun's continued enhancement of Java's threading functionality.

FIGURE 5

CMT PROCESSOR



Source: Sun Microsystems, 2004

The Numbers Get Interesting

One important aspect of Sun's approach to throughput architectures is that it gives processor designers three factors to work with when designing future multithreaded processor configurations:

- ☒ The number of cores per processor (m)
- ☒ The number of threads managed by each core (n)
- ☒ The single-thread performance of each core (p)

These three factors combine multiplicatively to yield throughput performance or overall processor capacity (i.e., capacity increase = $m \cdot n \cdot p$). As transistor budgets increase over time, Sun claims that its processor architects can increase all three factors, thus significantly increasing performance. For example, a processor with four

cores, four threads per core, and each core running at twice the performance of current SPARC chips could provide 32 times the throughput capacity of current processor (i.e., $m \cdot n \cdot p = 4 \cdot 4 \cdot 2 = 32$).

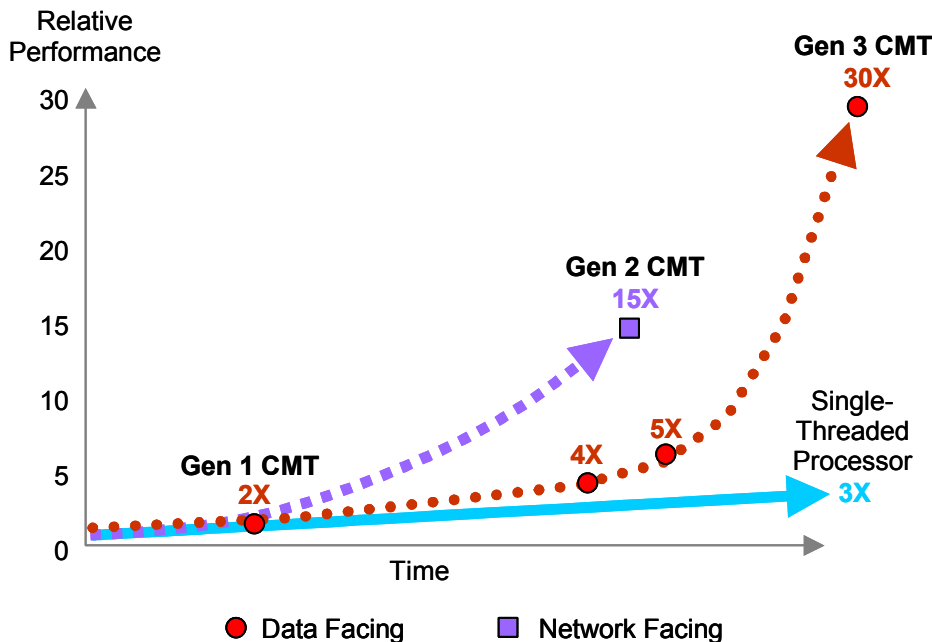
Sun's CMT Processor Road Map

Sun has announced systems using the UltraSPARC IV, the first milestone in its CMT road map, which is shown in Figure 6. UltraSPARC IV is a dual-thread processor offering customers up to twice the application throughput of the UltraSPARC III.

In its second and third generations of CMT processors, Sun will extend multithreading capabilities and integrated functionality. As shown in Figure 6, for network-facing applications, Sun expects to deliver CMT-enabled blade servers in 2005/2006 that provide 15 times more throughput than servers using the current blade processor, UltraSPARC III, which is used as a baseline for relative comparison. By 2007, for data-facing applications, Sun plans to deliver a CMT system processor with 30 times more throughput than today's UltraSPARC III. Over the same period, without CMT, assuming Wilkes' (2001) annual growth rate of 60% per year, Sun points out that the single-threaded performance of its processors would improve only by a factor of about three.

FIGURE 6

Relative Performance Improvement with CMT



Source: Sun Microsystems, 2004

IDC ANALYSIS

Today's servers can be idle up to 75% of the time while their processors stall waiting for data from memory, leaving considerable room for improvement in processor design. By focusing on increased application workload throughput instead of clock speed, Sun's CMT processors could deliver significant increases in application performance. Sun is developing processors that support both network-intensive and data-intensive server workloads and is positioning itself to address future growth in these workloads.

Alignment with Industry Trends

Sun's focus on significant improvements in price/performance ratios for its servers is well aligned with renewed demand by businesses that IT investments yield tangible benefits with manageable costs. While price points haven't been discussed, Sun is well positioned to offer high-performance servers that provide high availability and maintainability at a lower life-cycle cost.

Industrywide, server offerings are being configured for the next-generation computing architecture, which IDC calls utility computing. In the utility computing model, IT services are provisioned as needed from datacenters populated by powerful servers, high-capacity storage systems, and robust networks. As the number of networked devices continues to grow over the coming decade, the demand for high-performance, highly available, and cost-effective utility computing servers will grow proportionally.

The IT industry is generally interested in solutions to the performance imbalance among system components. Sun's focus on improving throughput is in alignment with that of its competitors, and its CMT approach is commonly viewed as viable and practical. With that said, servers using CMT processors are only beginning to enter the market, and proof points based on enterprise workloads are not yet available.

Sun's Challenges

Sun will face strong competition in a consolidated market of large and powerful surviving players. Both component and systems vendors are developing a variety of solutions to the system balance and throughput challenges. Sun's ability to win market share with sheer innovation will be tested. As Sun realizes, proven innovation must be woven into trusted sets of high-quality product and service offerings that allow customers to assimilate these new technologies successfully.

Sun will soon need to provide benchmarks to certify that its CMT strategy is on course. Sun's CMT processor architecture may provide improved throughput if Sun is successful in designing and fabricating a multithreaded processor with a zero-cycle thread switch and in finding enough parallel threads to occupy the processors. While IDC believes that the CMT approach is theoretically capable of providing improvements in throughput, the magnitude of those improvements cannot be accurately assessed until multithread processors have been fabricated and tested.

Scalable clustered and blade architectures also present a challenge to new architectural approaches to throughput computing. To the extent that Moore's law can provide reduced price for a constant level of performance, and the industry can continue to increase the scalability of clusters and blades, users may opt to expand their current systems with multiples of conventional processors rather than move to new architectural strategies. In addition, grid technology may provide a new dimension for scaling by allowing users to view collections of independent scalable servers as a single virtual system.

IDC believes that the IT community is currently recovering slowly from setbacks due to the global economic downturn. The IT consumer's appetite for new technology is limited and may remain so over the time frame when Sun intends to bring new products to market. Shifts as radical as Sun is forecasting may be difficult for consumers to absorb.

Sun's Opportunities

Sun's ownership of all its server technologies — including processor, operating system, and software infrastructure — provides the company with the opportunity to integrate its components in a balanced way over its multiyear planning horizon. Moreover, Sun's ability to integrate its software and processor technologies allows the company the opportunity to use innovation as a differentiator at every level — SPARC processors, Solaris operating system, Java development environments, and even its more encompassing offerings such as N1 and the Sun Java Enterprise System platform.

Furthermore, from a customer's point of view, moving to new Sun CMT-enabled platforms will not require a migration of software, compilers, and tools. No software disruptions are expected, and customers' investments will be protected. Sun is extending the SPARC architecture to further underpin support for thread-level processing, a capability that has been available for years in Sun's Solaris operating system and its attendant software infrastructure.

CONCLUSION

Sun has described a strategy for improving the throughput of servers with multithreading processors. CMT is one of several attacks on the industrywide problem of processors that operate more quickly than memory and thus often stall waiting for data. Sun views CMT as the next logical next step in 15 years of work on threaded operating environments and programming languages.

IT planners should value improvements in throughput because users will need more powerful servers to support new workloads in the decade to come. IDC's forecast growth in the installed base of networked devices will lead to a demand for high-performance servers. IT planners are encouraged to begin planning for the impact of a quantum leap in the number of devices on the next-generation network and to monitor proof points as Sun progresses toward delivering CMT processor-based products.

REFERENCES

- ☒ M.V. Wilkes. "The Memory Gap and the Future of High Performance Memories." AT&T Research Laboratories, Cambridge, U.K., 2001.
- ☒ W.A. Wulf and S.A. McKee. "Hitting the Memory Wall: Implications of the Obvious," *Computer Architecture News*, Vol. 23, No. 1: 20–24, March 1995.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2004 IDC. Reproduction without written permission is completely forbidden.