THE CLIPPER GROUP
Navigator™

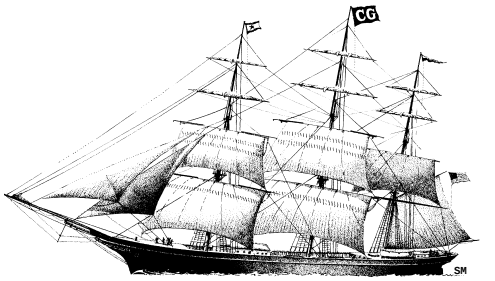

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2006025 April 18, 2006

Sun Improves Performance and Lowers TCO with UltraSPARC T1

Analyst: David Reine

Management Summary

The entertainment industry is littered with "one-shot wonders", performers who make an instant success, achieve their 15-minutes of fame, and then disappear from view as quickly as they arrived. Then, there are the stars - men and women who repeat their successes year after year, performers whose name alone on a record label or movie theater marquis can guarantee a hit. In the movie industry, for example, there are Clint Eastwood, Paul Newman, and Robert Redford, men who have spanned the past five decades with no acknowledgement to age, using their talent to evolve into roles that remain current to entertain succeeding generations of moviegoers. The same can be said of the leading lights in the music industry, performers such as Barbara Streisand and Wayne Newton, Elton John and Cher. These singers continue to bask in the gleam of the stage, entertaining young and old for decades. All of these stars have gone through periods when it may have appeared that time had passed them by, but talent and a will to succeed always return them to the spotlight.

The same is true for those on the Information Technology (IT) stage. In the 1950's and 60's, the computer industry was led by a group of companies that became known as IBM and the Bunch – with the Bunch an acronym for Burroughs, Univac, NCR, Cray, and Honeywell. How many of these names do we recognize today as leaders in the IT world. Through sale and merger, all of this bunch has lost their brand significance with CIOs around the world. One other company, however, has retained its preeminence through the decades, Sun Microsystems. Founded in 1982 with only four employees, **Sun has survived the passage of time (unlike Data General and Digital Equipment), not by following, but by leading the industry with new technology and sharing new ideas with all. Sun has flourished because of their ability to recognize both the revolutions in technology and the evolutionary trends required to reduce the total cost of ownership (TCO) of the IT infrastructure.**

First, Sun responded to industry demands for a better x86 environment by introducing AMD's *Opteron* processor in its *Sun Fire* product set, giving the data center compatibility for 32- *and* 64-bit processing within a single platform, then implementing dual-core x64 technology with Opteron in order to improve performance and reduce power consumption. **Now, Sun has extended the power and performance criteria in their own *SPARC* technology**, introducing the *Sun Fire T1000* and *T2000* with the multi-core *UltraSPARC T1* architecture. To see how the UltraSPARC T1 CPU can improve the performance of your data center while reducing your TCO, please read on.

IN THIS ISSUE

The Performance – Power Dilemma

The data center is under attack. Not by competitors, seeking out your corporate secrets. Nor by terrorists, trying to disrupt the capitalist economy. It is under attack from out-of-control server proliferation that has resulted in over-provisioned server farms running at 25% utilization factors. Under-utilized platforms acquired to execute one specific application are taking up too much valuable floor space and wasting too much of the energy required to both run the servers and, at the same time, the air conditioning necessary to cool them. They are affecting your bottom line.

Increased traffic, from an ever-increasing user base across the Web, is instigating the need to communicate, in parallel, with multiple servers in an ever-increasing server network. The CxO has learned to cope with the acquisition costs associated with this increased work load, but he is definitely having trouble coping with the recurring costs, the greatest constraints to growth, associated with those acquisitions, i.e., annual maintenance charges, floor space, and power, to name a few. These budget-busters are seriously affecting the total cost of ownership (TCO) of the IT infrastructure and reducing the profitability of the enterprise. Limitations in power are even influencing the ability of the enterprise to respond to mission-critical demands, especially when the local public utility tells you that they cannot even supply your data center with the power that you need to be competitive, regardless of the cost per kilowatt.

Obsolete servers, i.e. platforms more than six minutes old using single-threaded x86 architecture from Intel, traditionally measure performance as a function of CPU clock speed. A 2GHz *Xeon* server is better than a 1GHz Xeon, and a 3GHz Xeon is better than a 2GHz model. Unfortunately, the higher the clock speed, the more energy is required to run the server, often upwards of 120 to 150 watts. **With the cost of energy rising, the cost of powering even a small data center can become prohibitive**. The existing performance-centric chip design philosophy no longer matches the needs of an evolving data center requirement for reduced costs.

Within the past year, the x86-microprocessor industry has seen an evolution toward multi-core processors, following the lead established by Sun Microsystems, with *SPARC*, and IBM, with *POWER*. **Multi-core designs changed the CPU paradigm, enabling each core to run slower and cooler, drawing less power, needing less air conditioning, while increasing the overall performance of the processor.** AMD led this paradigm shift in the x86 arena with their open systems *Opteron* CPU, with Intel following their lead, evolving *Xeon* from single- to dual-core. Multi-core architecture has reinforced the availability of multi-threaded environments, promoting infrastructure transformation, with the consolidation of heterogeneous application sets onto a single platform. This enables the data center to take advantage of the idle time, the other 75%, to reduce the wasted resources in the infrastructure. As a result, fewer servers can accomplish more with the same amount, or less, expended energy. This is similar to one way that the federal government addressed the energy crisis thirty years ago – enabling right-turn-on-red at traffic lights to reduce wasted gasoline from idling. As a by-product, improved gasoline utilization reduced automobile emissions, improving the air quality. There is a by-product in the data center, also. A reduction in *power expended* means less heat created, therefore less air conditioning required for cooling.

IBM took the multi-core concept one small step forward, packaging two dual-core POWER5 CPU microprocessors into a quad-core module, creating a four-core processing module that uses one memory and one I/O controller. Now, Sun has taken that one giant leap further, combining eight UltraSPARC cores, with four threads in each, into a single *UltraSPARC T1 CoolThreads* module, capable of executing 32 separate streams at the same time. Combining this technology with the open architecture of Solaris and SPARC v9 binary compatibility provides an ideal path for data center consolidation and reduced TCO.

The Sun SPARC Evolution

Sun introduced platforms based upon SPARC technology 20 years ago and used them to lead the IT industry, and the rest of the world, into the Internet age. UltraSPARC T1 marks the ninth generation in the evolution of SPARC innovation. It represents a 1000x performance improvement over past two decades. In addition, UltraSPARC T1 reflects the latest advancements that Sun has implemented from ten years experience with multi-threaded processors. It also reflects the concern that Sun has with protecting the investment that customers have made in prior generations of SPARC servers. The T1 provides binary compatibility for applications written to the Solaris 10 operating system for all supported UltraSPARC systems via a SPARC v9 implementation.

By 1993, Sun had sold one million systems and by 1996 had evolved the UltraSPARC into a 64-bit architecture. Sun reached its peak in 2001 as an $18B leader in network computing, but then was one of the first companies to suffer from the backlash of the Internet economic downturn. Sun hit a lull in both sales and processor development in the years immediately following the bursting of the dot-com bubble, but they did not go to sleep, as is evident in

its position at the forefront of the paradigm shift to multi-core architecture. Now, however, it seems that it is returning to the spotlight with what must be considered a major innovation in high-performance, energy-efficient microprocessing. The UltraSPARC T1 technology is designed to double or triple a data center's processing capability without increasing its size. It addresses the urgent need for energy conservation while creating higher performance and faster response. It enables the improvement in processor utilization, following up, as phase 2, from Sun's announcement last year of their *Galaxy* line of Sun Fire servers based upon AMD'S Opteron microprocessor.

UltraSPARC T1 is available with 4, 6, or 8 cores, with four threads per core for a total of 32 threads, with the expectation to double the thread capability by 2008. UltraSPARC T1 takes advantage of the multi-threading capability of Solaris 10 with all cores able to see memory at the same time through a 134GB/sec crossbar switch. This enables all threads to work in parallel for faster performance. Originally codenamed *Niagara*, the T1 consumes only 72 watts of power, about the same as the average light bulb, leading Sun to stress their eco-responsibility leadership to the environment. (Will the new systems have a *green* bezel?) Sun's *CoolThreads* technology uses about 2 watts per thread, significantly lower than the x86 design from Intel, or the POWER architecture from IBM. T1 reduces the generation of heat, thus improving reliability through cooler execution. With a single floating-point unit, however, UltraSPARC T1 is not ideal for workloads heavily weighted toward math processing as in scientific computing environments.

Now, Sun has moved UltraSPARC T1 to open source, making multi-threaded processor technology available to all companies, in order to lower the barriers to innovation and thus help to contribute to Solaris application development. This should reduce development time for new solutions making Sun's Sun Fire platforms that much more attractive. Sun's expectations are that open sourcing will achieve the following.

- Seed a new SPARC ecosystem;
- Speed innovation into the data center; and
- Enhance the growth of a multi-thread application environment.

Sun has enumerated a number of benchmarks in which T1 has displayed a price/performance leadership position, including several from SPEC measuring Java and web performance and from *Lotus Notes*[1], reflecting how they are positioning these servers for database, web, and application tier workloads. Sun has formulated a new metric, however, for customers to use in order to evaluate platform energy efficiency and TCO. It measures space, wattage, and performance – **SWaP**[2] – and can be a useful tool in comparing the recurring costs that can plague rack-dense environments in your enterprise.

Exhibit 1 –
Sun Fire T1000 Entry Configuration

- 1x6 Core UltraSPARC T1 processor running at 1.0GHZ with 3MB L2 cache;
- 2GB of DDR-2 memory;
- 4x1GbE ports;
- 1 PCIe slot;
- 1 power supply;
- Solaris 10 Operating System

Source: Sun Microsystems

Sun Fire UltraSPARC T1 Solutions

Sun has now delivered two platforms designed for the enterprise looking for massive scale in their scale-out environment. Based upon the UltraSPARC T1 architecture, the *Sun Fire T1000* and *Sun Fire T2000*, complement the Sun Fire Opteron-based platforms that support the x86 requirement. They are small, rack-optimized servers with the Sun Fire T1000 starting at $2,995 when purchased with a *System Ready*[3] plan. It is a 1U server with 6 or 8 cores, up to 16GB of memory, and one 80GB SATA disk. It is ideally suited for compute, data, and transaction-intensive applications, and has been designed for web, portal, and network infrastructure, which have been dominated by small x86 systems. See Exhibit 1,above, for Sun Fire T1000 configurability.

The Sun Fire T2000 is a 2U-server designed with extensive internal redundancy and hot-plug capability for maximum uptime for multi-threaded application services and web-tier consolidation projects. It delivers a high level of reliability due to the reduction in the number of components enabled by the T1 architecture, measured in the thousands. It is ideal for OLTP, CRM, ERP, database, and collaboration environments. It has an entry price of $7,795 when purchased with a System Ready plan. Like the T1000, it runs under control of the newest iteration of the Solaris operating system, Solaris 10. See Exhibit 2, in the next column, for complete Sun Fire T2000 configurability.

[1] As always, the data center needs to be aware of the configuration of the tested server and how it compares to the proposed system.

[2] SWaP measures performance / (space (RU) x power).

[3] Sun offers a full protection plan that includes the essential, recommended services and one-to-three-year pricing so that customers can predict and manage their IT budgets and control the TCO.

One of the advantages that Sun can bring to your data center with these new Sun Fire servers is their 25-year relationship with Oracle. The T1000 and T2000 become more than just eco-responsible servers when you need to add database access to the basic platform. First, with Solaris 10 as one of the two development environments at Oracle, your data center can be sure of having the latest and greatest Oracle versions available upon announcement. You will not have to wait for porting to your environment. Second, as a trusted partner, Oracle has made remarkable (for them) concessions to Sun in order to put the T1 servers on *at least* an equal footing with competitive platforms. With the advent of multi-core processors, licensing of application software has become a minefield, with some ISVs charging by the socket while others are charging by the core. Sun and Oracle have struck a deal that appears guaranteed to benefit the user with lower software licensing costs, for a change.

Oracle is pricing their database products for the T1000 and T2000 at .25x per core. This means that an eight-core server will have a software charge of two Oracle licenses. This pricing gives the Sun Fire T1000 a favorable position when compared with x86 servers at .50x per core for Oracle and low-end IBM *pSeries*, with Oracle pricing at .75x per core. This gives Sun an edge from a sales perspective and gives the enterprise CIO a new view for TCO analysis.

Sun appears so convinced that the data center will become enamored with T1 servers, they have introduced a "Try and Buy" program that enables both resellers and customers to install and test the platform for 60 days, free of charge, with an option to purchase. Sun is sure that, when the data center sees the improved performance and reduced energy bills, the enterprise will insist that the server remains. Sun has reported that, so far, over 60% of these systems have gone to new customers, indicating strong interest outside of the Sun installed base.

The Sun Fire's tight integration with Solaris 10 enables the data center to take advantage of functionality built into the operating system to manage the environment. This includes the ability to assign specific threads to specific cores, or to disable a core while the system is operational. In addition, Sun's N1 System Manager (Rev 1.2) is now available to support both servers. This integrated offering allows the data center to perform operating system provisioning and system management of the new UltraSPARC T1 Sun Fire servers along with the Opteron family. The new Sun Fire servers can also use the Advanced Lights Out Management (ALOM) card for remote management to monitor CPU temperature, fan speed, thermal conditions, and drive status.

Exhibit 2 – Sun Fire T2000 Entry Configuration

- 1x4 Core UltraSPARC T1 processor running at 1.0GHZ with 3MB L2 cache;
- 8GB of DDR-2 memory;
- 4x1GbE ports;
- 4 USB ports;
- 3 PCIe/2 PCI-X slots;
- 2 73GB SFF SAS disks;
- 2 Hot-swappable, redundant power supplies
- 3 Hot-swappable, redundant fans;
- Hardware RAID, standard;
- Solaris 10 Operating System

Source: Sun Microsystems

Conclusion

After a few years in the shadows, Sun has returned to the spotlight of center stage. UltraSPARC T1 (along with dual-core Opteron) enables Sun to re-establish itself as a key player in the server market. Through innovation, and the formation of key partnerships (such as AMD and Oracle), Sun has become fashionable, again! **UltraSPARC T1 really does set a new standard for performance, space reduction, and energy efficiency!**

By releasing the specifications for T1 to the open source community, Sun stands to gain from the development of any new technology designed for the T1 architecture. In turn, if Sun gains, any UltraSPARC T1-based data center will also gain.

The success of the Sun Fire 1000 and 2000 will depend upon the availability of enterprise solutions designed for these multi-threaded platforms. With compatibility across the heterogeneous Solaris platforms, the T1000 and T2000 have a flying start with application availability in the thousands.

Does your enterprise data center need to double (or more) the performance capability of the server network? Has your CIO run out of ways to reduce the energy consumption in the data center? Have you run out floor space in your existing facility? Look at the new Sun Fire servers from Sun. With low acquisition costs and an even lower TCO, it may be the solution for which your enterprise has been searching.

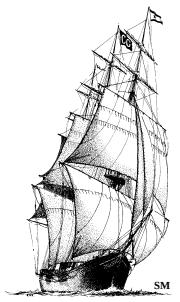

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

David Reine **is Director, Enterprise Systems for The Clipper Group.** Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial "123" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group *Captain's Log***, **The Clipper Group Voyager**, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *"Navigating Information Technology Horizons"*, and *"teraproductivity"* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.