THE CLIPPER GROUP
Navigator™

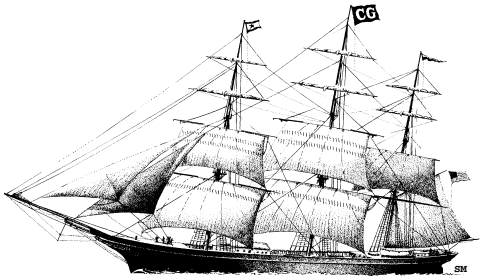

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2008040 August 11, 2008

Sun and Fujitsu Ignite Enterprise Systems with a New SPARC

Analyst: David Reine

Management Summary

No matter what business you are in, keeping your existing customer base happy and loyal is your most important mission. Nowhere is that more evident than in the world of consumer electronics where vendors provide upgrades as a matter of course for everything from your iPod to your Personal Computer. Software compatibility between releases ensures that the latest version of any commodity processor is compatible with its predecessor and when you need additional memory, usually you can either add additional modules to the motherboard or replace existing ones with higher capacity replacements. In order to keep customers up-to-date with the latest updates, software companies are vigilant to keep their customers current with their latest revisions. Data Protection companies, like Symantec and McAfee, are always on the alert to keep their clients protected from the latest viruses and to keep firewalls insulated against any attempt to invade your privacy. As long as you keep your subscription up-to-date, these companies will continue to provide you with the latest updates, even unsolicited.

The same requirements exist in the enterprise data center, except on a much grander scale. The IT staff is responsible to provide an ever-expanding network of employees, customers, and partners with high performance, highly reliable, mission-critical servers to perform an ever-increasing set of mission- and business-critical functions. In most large enterprises, this usually includes an environment consisting of high-performance computing (HPC) applications, along with a set of transactional processing activities involving relational databases that have grown to multi-terabyte capacities. In the largest enterprises, these environments usually operate in a legacy UNIX environment where the IT staff can take advantage of the economies of scale and reliability available in a scale-up architecture with a symmetrical multiprocessing (SMP) environment. Three companies have traditionally provided these high-end UNIX platforms and compete with each other for market share: HP, with their *Itanium*-based *Superdome* line running under *HP-UX*; IBM with their *System p POWER* platforms with *AIX*; and Sun with their *SPARC* line of *Solaris* servers. Over the past two years, HP and IBM have done a good job of providing the scalability required to meet a changing enterprise environment. Sun had been eclipsed on this innovation front, until now.

With their recent announcement, Sun and Fujitsu put a new spark into SPARC-based platforms. They have increased performance, lowered energy consumption, lowering the total cost of ownership (TCO) in SPARC, and preserved the investment that legacy data centers have made in Sun products. To learn more about *SPARC64 VII* and *SPARC Enterprise*, please read on.

IN THIS ISSUE

Legacy Data Center Requirements

An epidemic of under-utilized servers is running rampant, like an uncontrolled virus, throughout the data centers of almost every enterprise. If an evildoer plants a single virus on just one mission-critical enterprise server, disabling 85% of its compute cycles, the data center staff will employ a strike team to cut that virus out and restore the server to full health. How do we rationalize, then, the conditions existing today where many enterprise servers are operating at only 15% efficiency, while the IT staff ignores the other 85% of the CPU's capability? The cost to power these under-utilized servers, and cool the data center, increases the TCO of the data center and has a direct *negative* impact on the bottom-line of any enterprise. In addition, wasting natural resources limits the availability of energy to run new applications to meet the needs of a growing enterprise. *Consolidating* multiple servers onto a single scale-up server, while *virtualizing* the environment to enable multiple applications to operate in parallel is one method of removing much of the complexity out of the IT architecture and simplifying the infrastructure to enable the staff to restore order to the data center. **While it is painful to think about redesigning the IT environment, it is irresponsible to ignore the necessity**.

Two of the most urgent imperatives every CIO must face in changing the compute paradigm of the data center are (1) the need to make the architecture more **energy-efficient**, reducing enterprise demand on electricity, thereby improving performance per watt, and (2) **maintaining 24x7x365 availability** to ensure continuous access to mission- and business-critical applications. Improving the reliability, availability, and serviceability (RAS) of the application platform is not an option; it is a necessity! Downtime no longer is measured in hours but in hundreds of thousands of dollars per hour or more, an expense that no enterprise can afford. While availability is a priority, TCO and operational familiarity are significant secondary objectives.

Many of the servers installed throughout the enterprise were acquired 3-5 years ago, in the value-packed era immediately following the tech-bubble bursting in the early part of this decade. These servers were designed with single-core processors from Sun (*SPARC*) and companies such as Intel and AMD, using the x86 architecture for scale-out environments running *Windows*- and *Linux*-based applications. They were not designed for the scalability implicit in SMP architectures used in UNIX platforms and designed to grow with the enterprise. While lacking many of the RAS characteristics found in UNIX and mainframes systems, these x86 scale-out platforms were characterized as "good-enough" for the infrastructure environments and applications that they ran. This may be true for a single print server or file server, but when a data center installs hundreds of platforms in a mission-critical environment, good enough usually is *not* good enough. **Consolidating these disparate systems requires the highest levels of consistent and reliable performance and I/O throughput, not just good enough.** They also require the flexibility to access libraries of both UNIX and Linux applications, to avoid reinventing the wheel, with the RAS characteristics necessary to ensure system availability.

Over the past few years, there has been a pronounced emphasis among server providers to deliver systems geared to satisfy the requirements and desires of the most common denominator, the small and medium business, the SMB. With a typical user count of between 100 and 1000, there are literally thousands of SMBs with a requirement for anywhere from 10 to 50 servers to deploy in a scale-out, rack-mounted architecture, or packaged in blade chassis, to satisfy their mission- and business-critical application needs. Now, there is a striking need to apply that same innovation to the enterprise data center where legacy applications are outgrowing the scalability of their mission-critical platforms. In order to accomplish this, enterprise data centers are now beginning to deploy high-end SMP platforms implemented with multi-core processors to achieve the desired scalability and virtualization to improve server utilization and reduce the TCO.

In fact, multi-core processors have now taken over the mission-critical server domain, not just Intel's *Xeon* or AMD's *Opteron*, but from legacy vendors who have been building UNIX servers for the past two decades: Sun and Fujitsu with their *SPARC* architecture, IBM with *POWER*, and HP with *HP-UX*, now on *Itanium*. Data centers in this domain often are heavily involved in high-performance computing and back office transaction processing. In addition, there is also a significant presence of servers designed to manage the web infrastructure. *HPC applications* have a requirement for significant scalability, a large shared memory, and a high volume of floating point computations. Back office applications also require high scalability, and are characterized by high availability, improved response times, and improved utilization, as well as a requirement to protect business- and mission-critical data. The IT

staff needs to protect the investment in their legacy platforms, both from an application standpoint and a staffing position. Web servers are characterized by high throughput, low cost, and the need for high utilization of servers. They are also noted for low power consumption. Many of these enterprise data centers have deployed multiple SUN servers, including many with SPARC inside.

In all cases, these data centers must reduce the IT infrastructure in order to reduce the TCO and improve profitability. They can accomplish this through consolidation and virtualization, reducing the number of platforms, improving the utilization of each platform, lowering energy consumption; and reducing floor space. They cannot, however, lose any of the RAS characteristics for reliability, availability, and serviceability. These must remain, or improve, in order to satisfy enterprise service level agreements (SLAs)

By upgrading the existing architecture with denser, more highly scalable platforms, the data center can reduce the TCO without radically changing the compute paradigm. By replacing the existing environment with servers that have binary compatibility with the prior platforms, the IT staff can reduce application porting costs and staff retraining. Ease of deployment and ease of use is essential. **In order to meet all of these requirements, Sun and Fujitsu have continued their 20-year partnership to deliver the latest in SPARC innovation, with the new quad-core *SPARC64 VII*.**

Fujitsu SPARC64 VII Technology

The **S**calable **P**rocessor **ARC**hitecture, or **SPARC**, technology dates back over 20 years, to its origination at SUN in 1985. *SPARC64 VII* is the latest version of the *SPARC V9* specifications. Today, there are two roots for SPARC innovation: the SPARC64 architecture developed by Fujitsu and found in *SPARC Enterprise* mid-range and high-end servers and *UltraSPARC T1/T2/T2 Plus* processors developed by Sun and also found in the SPARC Enterprise low-end servers. Both variations have a common architecture that is administered by SPARC International. Because of this common architecture, any application running on Fujitsu's *PRIMEPOWER* or Sun's *Sun Fire* servers will run on SPARC Enterprise servers, thus preserving the investment made in SPARC technology and lowering TCO for the data center.

The first SPARC64 processor was developed in 1998, with improvements in reliability being integrated in *SPARC64 V* in 2002. In 2007, Fujitsu incorporated a dual-core technology into *SPARC V9* for all SPARC Enterprise servers, integrating all of the high reliability and availability features of SPARC64 V into *SPARC64 VI*. Now, Fujitsu and Sun are adding dual-thread, quad-core SPARC64 VII technology into all SPARC Enterprise mid-range and high-end models, with the SPARC64 VII being socket-compatible with SPARC64 VI.

SPARC64 VII is a multi-core, multi-threaded architecture with a 512KB L1 cache and a 6MB L2 cache to reduce memory latency. It uses a 65nm manufacturing technology, as compared to the 90nm technology of SPARC64 VI, enabling SPARC64 VII to provide four cores in about the same space as existing dual-core CPUs. The result is superior performance and improved cost efficiency. While SPARC64 VI has a nominal power consumption of 120 watts for a dual-core, dual-thread environment, SPARC64 VII has a nominal consumption of 135 watts for a quad-core, dual-thread architecture, running at 2.5GHz, reducing the power per thread ratio from 30 watts/thread to less than 17 watts/thread.

SPARC64 VII utilizes a simultaneous multi-threading architecture (SMT), compared to a vertical multi-threading architecture (VMT) that was used in SPARC64 VI. Multi-threading minimizes CPU core wait times and increases core utilization. While VMT improved thread execution, SMT also enables two threads to run in parallel. (See Exhibit 1 on the next page.) All of these improvements in core utilization reduce the data center's TCO.

SPARC64 VII takes advantage of high performance technology first introduced with SPARC64 V, such as *Out of Order Execution* to initiate instructions when they have executable data, and *Branch Prediction Mechanism*, to predict and prepare for the most likely instruction sequence. A more complete listing of SPARC64 high-performance technology features can be found in Exhibit 2, also on the next page.

SPARC Enterprise Solutions

SPARC Enterprise is the new SPARC/*Solaris* standard, combining the best technologies of its partners, Sun and Fujitsu, to deliver a versatile UNIX engine for a wide range of enterprise capabilities. These are mission-critical platforms that deliver the highest performance, reliability, and virtualization to enable efficient application processing in a scale-up engine. With the SPARC64 VII processor to drive them, along with high-performance bus technology, SPARC Enterprise servers provide scalability and high throughput in a compact energy efficient package. Sun and Fujitsu

Exhibit 1 – Multi-Threaded Architecture

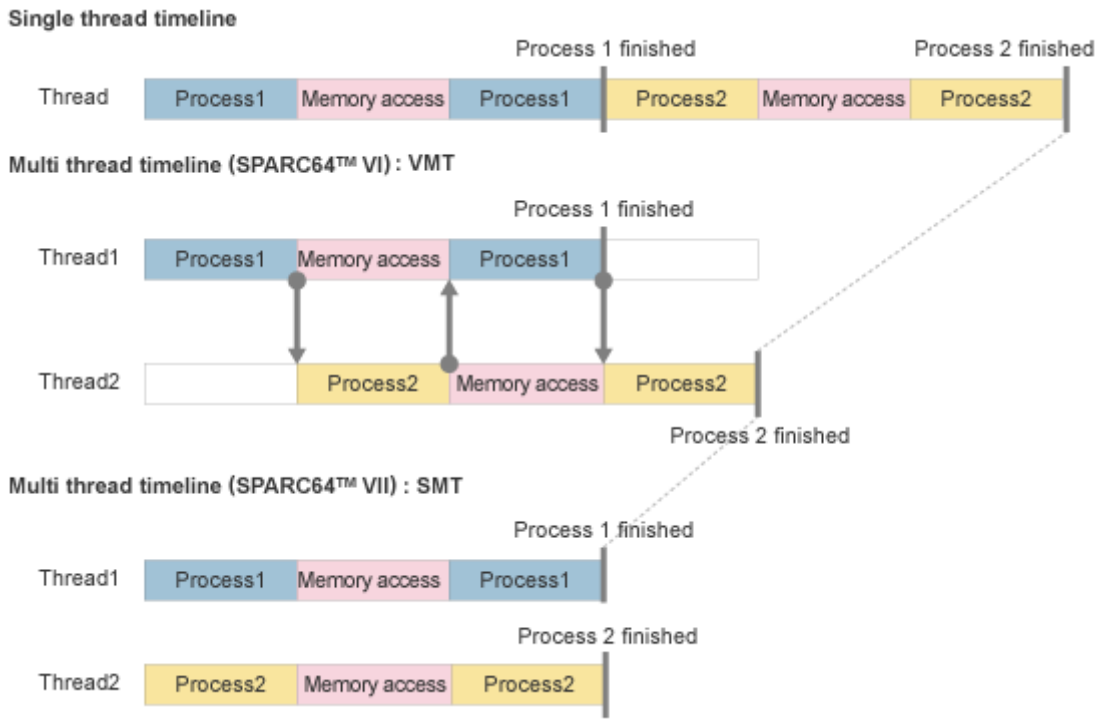

Source: Fujitsu

have now added SPARC64 VII to their complete family of highly scalable mid-range to high-end SPARC Enterprise servers: *M4000*, *M5000*, *M8000*, and *M9000*. These platforms join the *SPARC Enterprise T Series of UltraSPARC* servers to support HPC and back office applications, as well as Web Infrastructure.

These servers all support both the dual-core SPARC64 VI and the quad-core SPARC64 VII, running at speeds of up to 2.5GHz, delivering 1.8 times the performance of the previous generation. They all utilize DDR2 SDRAM memory with ECC/Extended ECC protection. Combined with the outstanding RAS characteristics found in Fujitsu mainframe technology for error detection and recovery, the M Series continues to provide business continuity in a mission-critical environment. All models also have the internal management necessary to reduce any possibility of failure by constantly monitoring all circuits, processors, and memory. Instructions and data are retried and rerouted automatically, locking out and identifying any potential point of failure. With the higher energy efficiency of quad-core technology, the M Series consumes less power, by as much as 44% per server, providing a lower overall TCO.

Sun and Fujitsu protect the investment made by the enterprise data center in the previous generation of the M Series by allowing upgrades from SPARC64 VI to SPARC64 VII CPU modules and enabling the IT staff to mix SPARC64 VI and VII modules in the same server, even in the same partition. Also, they configured the SPARC Enterprise family to take advantage of all of the performance features of the *Solaris 10 Operating System*, including the ability to deploy thousands of containers (virtual partitions) within a single hardware partition.

SPARC Enterprise M4000

The SPARC Enterprise M4000 is a versatile 6U rack-mounted server with industry-leading mainframe-class service levels. With up to four SPARC64 processors, supporting 16 cores and up to 32 threads, the M4000 provides the mid-sized data center with the scalability required in a compact server. Moreover, with up to 32GB/s of

Exhibit 2 – SPARC64 Technologies

- Pipeline
- Superscalar
- Out of Order Execution
- Branch Prediction
- Non-blocking cache
- Hardware fetch
- Advanced Data Protection
- Dynamic Degradation

Source: Fujitsu

system bandwidth, the M4000 has the throughput required to utilize fully the server compute capability. It is ideal for smaller business intelligence and data warehouse applications.

The M4000 scales to 128GB of memory with 292GB of internal disk and up to 25 PCI expansion slots. The data center can divide the M4000 into two hardware partitions with more than 8,000 Solaris Containers.

SPARC Enterprise M5000

Like the M4000, the SPARC Enterprise M5000 is rack-mounted, consuming 10U of rack space. With up to eight SPARC64 processors, supporting 32 cores and up to 64 threads, the M5000 provides a mid-sized data center with twice the scalability of the M4000 with the same throughput (I/O) capability. The M4000 is ideal for business tier workloads, including databases and ERP and CRM deployments.

The M5000 scales to 256GB of memory with 584GB of internal disk and up to 50 PCI expansion slots. The data center can divide the M5000 into four hardware partitions with more than 8,000 Solaris Containers.

SPARC Enterprise M8000

The SPARC Enterprise M8000 is a freestanding server with up to 16 SPARC64 processors, with 64 cores and up to 128 threads, making it an ideal engine for a data center with large mission-critical applications, such as ERP and CRM, enterprise database consolidation/virtualization, and the need for secure, on-demand performance. With up to 184GB/s of system bandwidth, the M8000 has the enhanced throughput required to utilize fully the compute capability of 64 cores. Moreover, with the Fujitsu mainframe heritage for reliability and availability, it is the ideal platform for a 24x7 environment: mainframe features without the mainframe cost.

The M8000 supports up to 512GB of memory with up to 2.34TB of disk capacity. It has 32 PCI Express slots, internal, and up to 112 slots via external I/O expansion units. The M8000 supports up to 16 Dynamic Domains and offers maximum granularity for virtualization and consolidation down to the CPU level.

SPARC Enterprise M9000

The SPARC Enterprise M9000 is a larger freestanding server, with up to 64 SPARC64 processors, 256 cores and up to 512 threads. It has the performance, scalability, and security for the most demanding enterprise environments. It combines the highest levels of performance and reliability for the M series, through extensive virtualization and automation, capitalizing on the full functionality of the Solaris OS. It is the most powerful server in the Enterprise line. It is ideal for the massive vertical scaling required for mission-critical enterprise consolidation.

With up to 737GB/s of system bandwidth, the M9000 has the throughput required to utilize fully the compute capability of all 256 cores. Moreover, with its mainframe heritage for reliability and availability, it is the ideal platform for a 24x7 environment.

The M9000, with the same service levels as the M8000, has significantly more scalability, supporting up to 2TB of memory and up to 9.34TB of disk capacity. It has up to 128 PCI Express slots in the base and expansion cabinets, and up to 288 additional slots via external I/O expansion units. The M9000 supports up to 24 Dynamic Domains with more than 8,000 Solaris Containers.

Conclusion

Over the past two years, Sun and Fujitsu have trailed the market in delivering the kinds of performance boost that their customers have been seeking. **Now, however, with the introduction of SPARC64 VII, they have provided an upgrade for the SPARC Enterprise M Series that will not only enhance the scalability and performance of each model of the family, but will also reduce the TCO of the data center.** These savings come through reductions in energy consumption, reduced space requirements via consolidation and virtualization, and the investment protection of existing hardware and software. They have positioned the M Series as a cost-effective, long-term solution for the business- and mission-critical data center.

The M Series has the scalability to provide the right amount of performance and throughput where you need it, when you need it. With scalability from four to 64 CPUS, the M Series provides your enterprise with the flexibility required to respond to an ever-changing business environment. If you are looking to improve performance and lower costs, look to the M Series platforms; they may have the right impact for your bottom line.

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**, **Clipper Notes**, 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.