

WHITE PAPER

Green IT of SPARC Enterprise Servers with Quad-Core SPARC64 VII Processor

Sponsored by: Fujitsu Ltd.

Vernon Turner

Jean S. Bozman

Kenneth Cayton

October 2008

EXECUTIVE SUMMARY

Green IT is about far more than the environmental impacts of using technology. For example, IDC server research has found that reducing power is also an important cost issue for customers — and a top priority for datacenter managers. In addition, space has become an issue for many businesses as the number of server footprints has grown significantly over the past 10 years. The SPARC Enterprise M3000, based on a single quad-core SPARC64 VII processor designed and manufactured by Fujitsu Ltd., is targeted squarely at these IT challenges.

The M3000 has a footprint that uses half of the space, or less, and provides nearly twice the performance of an earlier generation of servers based on SPARC architecture and Solaris. In addition, reliability, availability, and serviceability (RAS) features and functionality have been significantly increased, over time, by leveraging technologies that were originally developed for mainframe-class systems. These changes have been accomplished while maintaining full binary compatibility with older SPARC/Solaris systems for investment protection and the ability to move customer applications, without recompilation, to a new, more powerful platform.

The M3000 server is a one-socket server that is positioned in the volume server space (IDC defines volume servers as those that are priced less than \$25,000). It fills the need for cost-effective Unix systems to cover the wide array of applications that run on volume systems, such as IT infrastructure, Web infrastructure, collaborative workloads, and software development. The M3000 is the newest member of the M-Series line, which is based on Fujitsu's SPARC64 processors, and is being jointly brought to market worldwide by Fujitsu and by Sun Microsystems.

Importantly, the M-Series models M4000, M5000, M8000, and M9000 have on-board virtualization technologies, easily taking advantage of the multiple cores and threads provided by the SPARC64 VII processors to increase server utilization — and to improve performance for customers' Solaris applications. The more scalable servers in the M-Series product range have partitioning; with the M3000, each individual server is, in effect, a partition — and multiple M3000s can be clustered together in scale-out configurations.

SITUATION OVERVIEW

Today's datacenter may be located in the same building as it was five or 10 years ago — but the demands being made on it are very, very different. Ten years ago, the focus was buying the servers with the most capacity and the best price/performance. Today, performance and price/performance are still important — but new criteria have emerged as challenges that must be met by datacenter managers:

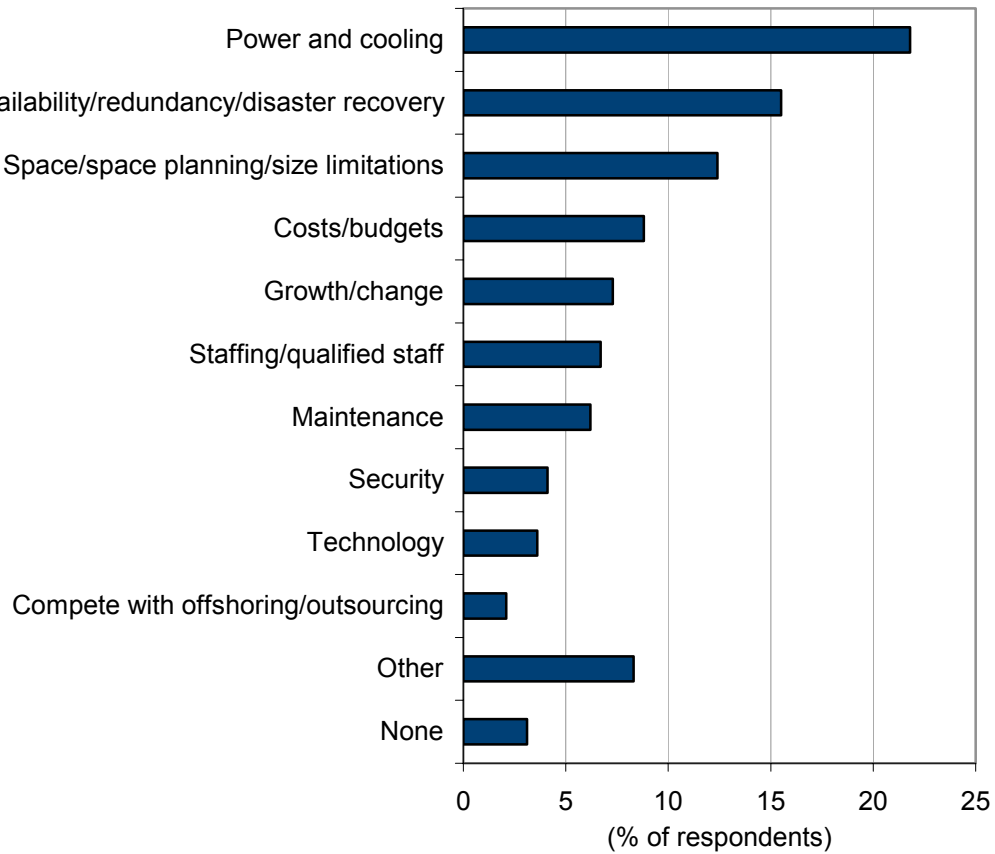
- ☒ Today's datacenter has many more servers inside it — and a limited supply of power and cooling to keep all of the servers working smoothly without overheating. IDC estimates that 30 million server systems were installed in 2007, and we expect the number of systems to increase at a 6.6% CAGR to over 41 million in 2012.
- ☒ Today's datacenter is being accessed across the Internet by end users and by customers, so there is no time for downtime — planned and unplanned downtime impact both a company's revenue and its profitability.
- ☒ Today's datacenter must find ways to deploy more processing power within smaller form factors, making virtualization, server utilization, and space considerations more important than ever before.

These concerns are known by IT managers and datacenter managers all over the world. As shown in Figure 1, the top-of-mind challenges for datacenter managers are:

- ☒ **Power/cooling.** Reducing power/cooling costs through improved energy efficiency has emerged as the top priority for many datacenter managers. The reasons are clear: limited energy supply in city areas, leading to "power capping." This means that all datacenter servers, storage, and networking devices must be powered within a certain "power envelope." In any case, use of closely packed servers within the datacenter leads to "hot spots" that must be controlled, meaning that datacenter servers must also operate within a "thermal envelope" — or operational costs will rise even further.
- ☒ **Availability of applications and data.** Availability is of critical importance in a computing environment that must operate on a 24 x 7 x 365 basis. End users within the organization — and end customers outside the organization — must be able to access data and applications on a 24-hour basis, from any time zone around the world. Therefore, there is "no time for downtime" — and all applications must be working all the time. If unplanned downtime occurs, workloads must be transferred to other servers for continued processing, using high-availability software to restart the processing.
- ☒ **Space utilization within the datacenter.** Today's datacenters are packed full of many small servers, and they also house scalable servers (midrange and high-end servers) that handle the largest and most complex workloads. This crowding in the datacenter contributes to power/cooling problems — and customers are looking for new ways to deploy servers in order to use the existing datacenter space more efficiently while supporting more energy-efficient operations.

FIGURE 1

Top Challenges in the Datacenter



n = 197

Source: IDC, 2008

Businesses benefit from lower operational costs. Those costs can be related to energy use, to IT staff time related to maintenance and upkeep for existing systems, and to downtime from unreliable systems, power outages, or network outages. Technologies and products that address this important need to reduce operational costs will contribute to overall company revenue and overall profitability. Another benefit of improved uptime for computer systems is business continuity, which provides very real benefits to an organization's end users — and to its customers.

Green IT and Power and Cooling

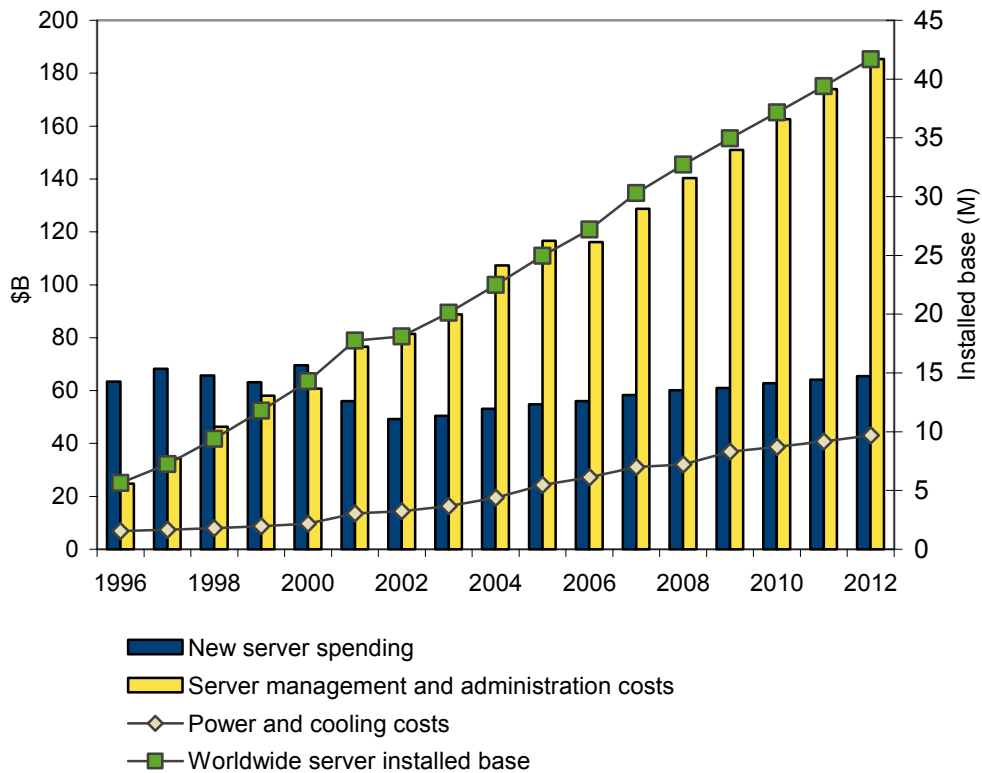
Green IT is an important initiative for today's businesses, which face limited power provisioning from utility companies, especially in large cities. The "power cap" experienced by companies is a limit by which IT organizations must operate — reducing the energy "envelope" for all of the servers under management. Therefore, customers must find ways to operate within this envelope, finding energy efficiencies wherever possible.

As shown in Figure 2, power/cooling costs have been rising rapidly in recent years, as the total number of servers deployed in customer sites has grown dramatically. The more power that is required to operate a system, the more heat energy that is dissipated during its operation. This phenomenon leads to "hot spots" throughout the datacenter — most often affecting the rows of rack-optimized servers, where many small servers are packed closely together. In fact, power/cooling costs have grown by a factor of four, compared with the cost of acquiring new servers. That is why server vendors are working to reduce the power/cooling requirements through the following approaches:

- ☒ Use of multicore processors, reducing the energy for each processor core
- ☒ New server designs that require less energy for operation
- ☒ Efficient cooling mechanisms for densely packed servers
- ☒ Detection of hot spots within the datacenter
- ☒ Use of hot aisle/cool aisle designs for the overall datacenter

FIGURE 2

Worldwide IT Spending on Servers, Power and Cooling, and Management/Administration, 1996–2012



Source IDC, 2008

High Availability and Reliability for Enterprise Applications

Today's enterprise is frequently competing in a global 24 x 7 marketplace and has little tolerance for any interruption to its ongoing operations. Because IT systems are the foundation for modern business processes, downtime in IT infrastructure means downtime for the overall business, or for part of it. Lack of business continuity leads to fewer sales, and it can cause customers to seek other providers of products and services that offer more reliable business services.

IT has a variety of technologies to improve availability of applications and data, such as high-availability failover software, workload balancing software, data replication, and enterprise system management "orchestration" that moves workloads to available computing resources. However, the interruption that may take place as workloads are shifted to alternate servers can, in itself, be a source of frustration to end users who find their work, or their purchase of goods and services, interrupted.

This means that each server system must provide the highest availability possible, through the support of on-board reliability features, such as error-correcting memory, memory protection features, and redundant hardware components that can be quickly repaired or replaced, if needed. Beyond that, the groupings of servers must be efficient and useful. In this way, clusters of servers can be grouped together for high-availability purposes, and software that speeds a restart of processing is very useful.

Price/Performance of Entry Unix Servers

Entry Unix servers have already become an important element of core computing for today's enterprises. They support network-intensive workloads, as requests come in from the Internet to access applications and data running in the datacenter. They support telecommunications infrastructure, supporting wired and wireless communications applications with high availability, where they typically are deployed in scale-out configurations with dozens, or even hundreds, of entry servers. They support government workloads, with high degrees of availability and security, given support for user-defined "roles" and identity authentication software.

Ideally, customers would like to acquire a system that would have midrange server capabilities, with high levels of availability, reliability, and security, but be able to buy it at an entry server price. Economic pressures have caused IT managers to search for lower price points for entry servers to complement the scalable servers that bring more capacity and more advanced management features to the largest and more demanding mission-critical workloads. There is a role for both types of servers — entry servers and scalable servers — to run the enterprise datacenter. Importantly, customers are looking for entry servers that incorporate innovations such as multicore/multithreaded technology and virtualization to support consolidation of workloads — and to save space in the datacenter — and to support energy efficiency that allows IT organizations to add capacity while reducing power/cooling costs.

Benefits of Multicore, Multithreaded Servers

Multicore processors have brought a new approach to system design. Multicore processing allows each processor core to run at lower power than would otherwise be the case for a single-core processor. These multicore systems, shipping in dual-core and quad-core versions, have impacted datacenter operations already, allowing each system to do more work in less overall datacenter space, contributing greatly to energy-efficiency plans, and reducing operational costs.

Grouping multiple cores together increases total throughput for computing workloads. At the same time, server power/cooling requirements are kept under control.

In addition, server virtualization improves the utilization of each processor. The combination of multicore and virtualization technologies has greatly improved server utilization so that fewer server "footprints" support datacenter IT workloads without requiring expansion of the entire datacenter.

Unix Server Workloads

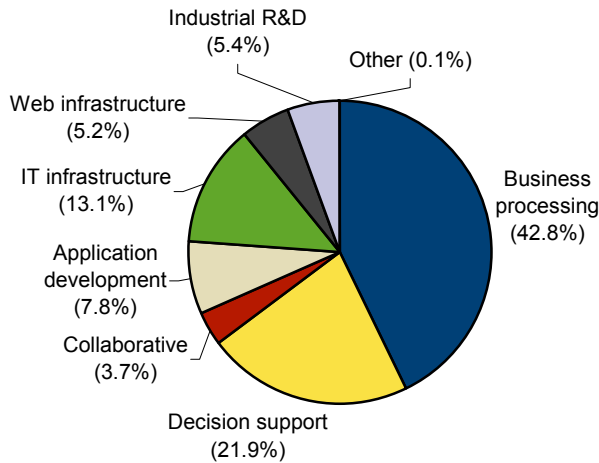
Unix servers have long played an important role in the datacenter by hosting mission-critical workloads and Web-enabled workloads (workloads that have an Internet component but that link to enterprise applications). Unix servers are widely deployed worldwide, with important roles in government, telecommunications, financial services, manufacturing, retail, and other vertical market segments.

The RAS features of Unix servers, which have been built into the platform over a period of many years, provide an important foundation for important applications and databases. Such reliability and availability are extremely important in environments where downtime would stop business operations, causing revenue loss and impact on profitability. Unix servers have a strong reputation for supporting important workloads, as shown in Figure 3. Typical workloads for Unix servers include business processing (e.g., online transaction processing [OLTP] and line-of-business [LOB] workloads such as ERP and CRM).

In many cases, large businesses view Unix servers as providing the type of reliability that otherwise could be provided only by scalable mainframe servers. IDC believes this is why the worldwide Unix server market remains strong, even as prices for individual point products reach lower price points due to intense competition in the marketplace. In 2007, the worldwide Unix server market generated more than \$17 billion in factory revenue and nearly \$19 billion in customer revenue (revenue based on installation at customer sites). IT spend on Unix servers for 2007 amounted to 32% of all IT spend for all types of servers worldwide, as measured by IDC each quarter.

FIGURE 3

Unix Workloads Worldwide Customer Revenue Share



Source: IDC, 2008

Entry-Level SPARC Enterprise Server

The M3000 system, based on the Fujitsu SPARC64 VII processor, provides a new entry point into the full range of products in the SPARC Enterprise server line. This reduces the initial price point in the M-Series product range for purposes of:

- ☒ Wide range of applications and databases for business units within the enterprise
- ☒ Web-enabled workloads for the enterprise
- ☒ Web 2.0 workloads for the Internet and Web serving
- ☒ Application development for SPARC64 Enterprise servers

Full Range of M-Series Servers

The M3000 will serve as the entry point to the M-Series server line, as sold by Fujitsu and by Sun worldwide. Customers can deploy multiple M3000s to host a portfolio of business applications, including customer-specific applications (often called "custom" applications) and ISV applications. They can also use the M3000 to host Web-based workloads, including Web-enabled applications that work with databases that are running on the larger M-Series models. Or they can use the M3000 as a test/development system for programming and certifying new applications to be deployed across the full family of M-Series servers.

What follows is a brief summary of all of the other server models in the M-Series product line that the M3000 server now joins:

- ☒ **M4000.** The M4000 is a four-socket system with up to 16 processor cores per system. It supports up to two Dynamic Domains and up to 128GB of memory. There are up to 5 I/O slots, supporting PCI-X or PCI Express (PCIe). Additional PCIe or PCI-X slots can be added with the External I/O Expansion Unit. System bus bandwidth of up to 32GBps is supported.
- ☒ **M5000.** The M5000 is an eight-socket system with up to 32 processor cores per system. It supports up to four Dynamic Domains and up to 256GB of memory. There are up to 10 I/O slots, supporting PCI-X and PCIe. Additional PCIe or PCI-X slots can be added with the External I/O Expansion Unit. System bus bandwidth of up to 64GBps is supported.
- ☒ **M8000.** The M8000 is a 16-socket system with up to 64 processor cores per system. It supports up to 16 Dynamic Domains and up to 512GB of memory. Four I/O units, each with up to 8 PCIe I/O slots, are supported. In total, the server supports 32 I/O slots. Additional PCIe or PCI-X slots can be added with the External I/O Expansion Unit. System bus bandwidth of up to 184GBps is supported.
- ☒ **M9000.** The M9000 can be deployed as a 32-socket system or as a 64-socket system — with one system per cabinet. Up to 256 processor cores per system are supported in the 64-socket model. The large system can support up to 24 Dynamic Domains. Up to 64 PCIe I/O slots are supported per 32-socket cabinet, while up to 128 PCIe I/O slots are supported per 64-socket system. Additional PCIe or PCI-X slots can be added with the External I/O Expansion Unit. System bus bandwidth of up to 737GBps per 64-socket system is supported.

The M-Series Product Line

Ease of Use

Ease of use is important for longtime Solaris customers — and it is also important for attracting new customers to the M-Series platforms. Many Solaris applications within the Solaris installed base are dated at Solaris 8 or Solaris 9 — and therefore many of those workloads need to move to new, more powerful hardware. However, due to concerns about costs associated with moving the older workloads, many have remained on systems based on SPARC hardware that is five years old, or older. When this happens, customer applications cannot benefit from improvements in SPARC processor technology, as implemented in SPARC64 multicore systems. Customers want a "seamless" transition to the new hardware, which the M-Series provides. They want to be able to take their older Solaris applications, and databases running under the Solaris operating system, and move them to new, faster hardware that has lower power/cooling costs for ongoing operations. The M-Series systems, including the M3000 server, make this possible through support of binary compatibility for older SPARC deployments. It will be possible to take Solaris 8 and Solaris 9 applications and to redeploy them on the M3000. This means that, for most applications, there is little or no need to modify existing Solaris applications in order to deploy them on the M3000. One way to do this is to run Solaris 10 "containers" on the M3000 hardware so that individual containers can host images of Solaris 8, Solaris 9, or Solaris 10, depending on customer preference. Dozens of containers can be hosted within a single Solaris 10 operating system domain.

Investment Protection

This kind of straightforward move to newer systems will make platform updating easier, and faster, than before. It will preserve customers' previous investments in the SPARC/Solaris platforms offered, over many years, by Fujitsu and by Sun Microsystems. This means that customers will avoid the time and money costs that are often associated with platform upgrades and replacements — and that there will be no disruption of IT operations when M-Series systems are introduced.

SPARC Enterprise servers provide investment protection for Solaris/SPARC customers, who can now move applications seamlessly to new SPARC64 platforms via binary compatibility. It is important for Solaris customers on SPARC systems to move existing Unix applications to new hardware — to gain more performance and improved energy efficiency — and to put them into production immediately. The result of these moves into the new platforms will be more server capacity, at less cost, and the ability to "scale up" workloads onto a range of M-Series systems, if needed, as demand for computing grows.

Upgrading in place is another option that allows customers to preserve previous investments in the SPARC hardware platform. SPARC Enterprise M-Series systems based on dual-core SPARC64 processors can be upgraded by adding the new generation of quad-core SPARC64 VII processors to the existing hardware chassis. This upgrade-in-place approach speeds system performance and improves performance/watt metrics without the need to replace other components of the system.

Support for Green IT

Green IT initiatives have become high-profile efforts worldwide. This is important across all geographic regions, including the Americas (North America and South America), EMEA (Europe, the Middle East, and Africa), Asia/Pacific (including Australia and New Zealand), and Japan. Consciousness about the importance of green IT has gone up dramatically in recent years, prompted by environmental concerns, the study of global climate change, and economic issues associated with energy efficiency and power and cooling inside datacenters.

Many factors are included in green IT initiatives, all along the life cycle of the server products. These factors include attention to the materials used in the servers, energy use of the servers while operating in the datacenter, and recycling of components after the servers are retired from use within the datacenter.

Energy efficiency for each server, as measured in performance/watt and other metrics, gives IT managers a way to evaluate the environmental impact of using servers. It also provides a reason to replace older, or aging, server products that are not as energy efficient as modern server designs. Indeed, power/cooling concerns have changed greatly over the past 10 years — and the new design points for servers address energy efficiency through multicore processing, use of new materials, and even the direction of cooling air as it passes through the server cabinet or server rack.

As a result of green IT initiatives, the SPARC Enterprise M-Series product range, as sold by both Sun and Fujitsu worldwide, is designed to minimize power consumption and to keep the server cool. All M-Series models use a 2.5-inch hard disk drive (HDD), supporting overall system compactness and power savings and a high efficiency AC-DC conversion unit to minimize power requirements. In addition, these models are equipped with efficient air-control mechanisms for cooling hot spots inside the server cabinet, including both processors and memory.

Introducing the M3000

The M3000 server is designed to be the entry point to the line of M-Series server offerings, including the M4000 and M5000 midrange enterprise servers and the M8000 and M9000 high-end enterprise servers.

Fujitsu and Sun have been shipping the M-Series server products worldwide since 2007, and both companies intend to expand this line of enterprise server offerings, which supports mission-critical workloads by providing RAS hardware features and support for the Solaris 10 Unix operating system, as well as many Fujitsu system management software products that provide efficient and reliable operation of M-Series server systems.

The M3000 is a 2U one-socket system using dual-core or quad-core versions of the SPARC64 VII RISC processor, which is a 65nm process semiconductor that is designed and manufactured by Fujitsu Ltd. The SPARC64 VII processor's quad-core capabilities will give it computing density — and yet the server is designed to be energy efficient. Each processor consumes 135 watts of power — and the entire unit operates within a thermal envelope of 470 watts at a rated voltage of 100 to 120 VAC. (with dual zones of cooling within each server).

This combination of compute density and energy efficiency means that the system provides powerful computing capabilities in a small form factor, supporting green IT. It supports international environmental initiatives, including the RoHS standard for materials (which began in Europe and was later adopted more widely) and the energy-saving programs in Japan. It uses fan-rotation control to minimize power usage and embedded air ducts to address hot spots within the system, cooling the entire unit while reducing ambient noise levels of 47 decibels (dB) at 25 degrees centigrade (or 77 degrees Fahrenheit). The M3000 can be deployed by itself for test and development purposes — or it can be deployed in production for use with other datacenter servers in the M-Series server line.

The M3000 addresses all of the top 3 datacenter concerns of IT managers:

- Power/cooling efficiency
- Availability for workloads running on the system
- Efficient use of datacenter space

Its small form factor provides IT with flexibility regarding how to deploy the system while supporting business continuity through high availability and reliability — reducing the possibility of unplanned downtime for the system.

Specifications for the system include the following:

- ☒ **Data-error checkpoints on the SPARC64 VII processor.** This is built into the SPARC64 VII processor and provides a hardware-based check for data errors. It supports instruction retry at the processor level.
- ☒ **Protection for L1 and L2 cache memory.** Many system errors can occur in L1 or L2 cache memory. The SPARC64 VII processors provide protection against these kinds of failures.
- ☒ **Extended error checking and correction (ECC) for memory.** Extended ECC uses an ECC-like mechanism to rescue data left on faulty memory chips.
- ☒ **Solaris containers.** These are software "containers" that, through virtualization, isolate applications from one another within a single Solaris operating system domain. This prevents interference between applications and supports business continuity.
- ☒ **Redundant hardware components.** This provides alternate computing resources at the system level in case any single hardware component fails.

The M3000 provides nearly twice the performance, and roughly 25% to 50% of the space, of some of the previous generation SPARC-based server products from both Sun and Fujitsu that it will replace. This will be especially important for space- and power-constrained operations, where there is a limit on how much power can be consumed within the datacenter. All of this has been accomplished while lowering the noise levels significantly. Importantly, system reliability has been enhanced significantly across all aspects of the server system. This is deemed to be very important in today's highly competitive global marketplace, where enterprises run real-time operations on a 24 x 7 x 365 basis, supporting business activities across multiple geographies.

CHALLENGES/OPPORTUNITIES

Any time that a new server system is introduced, customers need to become familiar with its feature/functions and its use cases. This generally takes some time to accomplish, especially for a new system design. In the case of the M3000, the operational benefits of using a system that supports green IT and energy efficiency will be attractive to many IT managers, and adoption of this new system will likely be faster, for these reasons, than that of servers that do not bring these operational benefits.

As they sell the M-Series servers worldwide, Fujitsu and Sun must highlight these green IT features and energy-efficient characteristics, along with the system's support for highly available workloads.

The worldwide Unix server market is a very competitive market space, with four top vendors, resulting in intense price competition, especially in the volume server space, which benefits customers. In addition, customers buy new servers based upon the value associated with reducing operational costs for the business. The value is seen in acquiring new Unix servers that support green IT, reduced power/cooling costs, and RAS features for availability — all of which combine to show the business benefits of deploying a new generation of energy-efficient servers in datacenter IT infrastructure. Demand for these types of servers will likely increase, going forward, due to their ability to impact and reduce operational costs. Importantly, these types of products can maintain the drive toward dynamic IT — which is the alignment between sustaining IT requirements for flexible and robust data processing and supporting ever-changing business requirements.

CONCLUSION

Energy efficiency is a top concern for IT managers and datacenter managers worldwide. Power/cooling costs have been climbing in recent years, driven by rising energy costs and the crowding of closely packed servers, storage, and networking devices within the datacenter. In some cities, there are limits on how much power can be consumed within a datacenter or an office building. The introduction of the M3000 server addresses many of these top-of-mind concerns while driving the business benefits of reduced operational costs by lowering energy use, improving IT staff productivity, and improving uptime for applications and data.

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 2008 IDC. Reproduction without written permission is completely forbidden.